

Analyzing the BrowserID SSO System with Primary Identity Providers Using an Expressive Model of the Web

Daniel Fett, Ralf Küsters, and Guido Schmitz

University of Trier, Germany
{fett,kuesters,schmitzg}@uni-trier.de

Abstract. BrowserID is a complex, real-world Single Sign-On (SSO) System for web applications recently developed by Mozilla. It employs new HTML5 features (such as web messaging and web storage) and cryptographic assertions to provide decentralized login, with the intent to respect users' privacy. It can operate in a primary and a secondary identity provider mode. While in the primary mode BrowserID runs with arbitrary identity providers, in the secondary mode there is one identity provider only, namely Mozilla's default identity provider.

We recently proposed an expressive general model for the web infrastructure and, based on this web model, analyzed the security of the secondary identity provider mode of BrowserID. The analysis revealed several severe vulnerabilities, which have been fixed by Mozilla.

In this paper, we complement our prior work by analyzing the even more complex primary identity provider mode of BrowserID. We do not only study authentication properties as before, but also privacy properties. During our analysis we discovered new and practical attacks that do not apply to the secondary mode: an identity injection attack, which violates a central authentication property of SSO systems, and attacks that break an important privacy promise of BrowserID and which do not seem to be fixable without a major redesign of the system. Interestingly, some of our attacks on privacy make use of a browser side channel that, to the best of our knowledge, has not gained a lot of attention so far.

For the authentication bug, we propose a fix and formally prove in a slight extension of our general web model that the fixed system satisfies all the authentication requirements we consider. This constitutes the most complex formal analysis of a web application based on an expressive model of the web infrastructure so far.

As another contribution, we identify and prove important security properties of generic web features in the extended web model to facilitate future analysis efforts of web standards and web applications.

Table of Contents

1	Introduction	4
2	The Web Model	6
	2.1 Communication Model	6
	2.2 Scripting Processes	8
	2.3 Web System	8
	2.4 HTTP Messages	9
	2.5 Web Browsers	9
3	General Security Properties	12
4	The BrowserID System	12
	4.1 Overview	12
	4.2 Implementation Details	14
5	Analysis of BrowserID: Authentication Properties	16
	5.1 Modeling of BrowserID with Primary IdPs	17
	5.2 Authentication Properties of the BrowserID System	18
	5.3 Identity Injection Attack on BrowserID with Primary IdPs	18
	5.4 Security of the Fixed System	19
6	Privacy of BrowserID	19
	6.1 Privacy Attacks on BrowserID	20
	6.2 Fixing the Privacy of BrowserID	21
7	Related Work	22
8	Conclusion	22
	References	23
A	Communication Model	24
	A.1 Terms, Messages and Events	24
	A.2 Atomic Processes, Systems and Runs	26
	A.3 Atomic Dolev-Yao Processes	26
	A.4 Scripting Processes	27
B	Message and Data Formats	27
	B.1 Notations	27
	B.2 URLs	28
	B.3 Origins	29
	B.4 Cookies	29
	B.5 HTTP Messages	29
	B.6 DNS Messages	30
C	Detailed Description of the Browser Model	31
	C.1 Notation and Terminology (Web Browser State)	31
	C.2 Description of the Web Browser Atomic Process	34
D	General Security Properties of the Web Model	42
E	Step-By-Step Description of BrowserID (Primary IdP)	45
	E.1 LPO Sessions	45
	E.2 Step-By-Step Description	46

F	Model of BrowserID with Primary IdPs	48
F.1	Outline	48
F.2	Addresses and Domain Names	48
F.3	Keys and Secrets	49
F.4	Identities	49
F.5	Corruption	50
F.6	Processes in \mathcal{W} (Overview)	50
F.7	Attacker	51
F.8	Browsers	51
F.9	LPO	51
F.10	Relying Parties	54
F.11	Identity Providers	56
F.12	BrowserID Scripts	58
G	Formal Security Properties	77
H	Proof of Theorem 1	77
H.1	Overview	77
H.2	Condition A	78
H.3	Condition B	82
I	BrowserID Login Flow Overviews	85

1 Introduction

Single sign-on (SSO) systems have become an important building block for authentication in the web. Over the last years, many different SSO systems have been developed, for example, OpenID, OAuth, and proprietary solutions such as Facebook Connect. These systems usually allow a user to identify herself to a so-called relying party (RP), which provides some service, using an identity that is managed by an identity provider (IdP), such as Facebook or Google.

Given their role as brokers between IdPs and RPs, the security of SSO systems is particularly crucial: numerous attacks have shown that vulnerabilities in SSO systems compromise the security of many services and users at once (see, e.g., [3, 7, 23–26]).

BrowserID [21] is a relatively new complex SSO system which allows users to utilize any of their existing email addresses as an identity. BrowserID, which is also known by its marketing name *Persona*, has been developed by Mozilla and provides decentralized and federated login, with the intent to respect users' privacy: While in other SSO systems (such as OpenID), by design, IdPs can always see when and where their users log in, Mozilla's intention behind the design of BrowserID was that such tracking should not be possible. Several web applications support BrowserID authentication. For example, popular content management systems, such as Drupal and WordPress allow users to log in using BrowserID. Also Mozilla uses this SSO system on critical web sites, e.g., their bug tracker Bugzilla and their developer network MDN.

The BrowserID implementation is based solely on native web technologies. It uses many new HTML5 web features, such as web messaging and web storage. For example, BrowserID uses the `postMessage` mechanism for cross-origin inter-frame communication (i.e., communication within a browser between different windows) and the web storage concept of modern browsers to store user data on the client side.

There are two modes for BrowserID: For the best user experience, email providers (IdPs) can actively support BrowserID; they are then called *primary IdPs*. For all other email providers that do not support BrowserID, the user can register her email address at a default IdP, namely Mozilla's `login.persona.org`, the so-called *secondary IdP*.

In [13], we proposed a general and expressive Dolev-Yao style model for the web infrastructure. This web model is designed independently of a specific web application and closely mimics published (de-facto) standards and specifications for the web, for instance, the HTTP/1.1 and HTML5 standards and associated (proposed) standards (mainly RFCs). It is the most comprehensive web model to date. Among others, HTTP(S) requests and responses, including several headers, such as cookie, location, strict transport security (STS), and origin headers, are modeled. The model of web browsers captures the concepts of windows, documents, and iframes, including the complex navigation rules, as well as new technologies, such as web storage and cross-document messaging (`postMessages`). JavaScript is modeled in an abstract way by so-called scripting processes which can be sent around and, among others, can create iframes and initiate XMLHttpRequests (XHRs). Browsers may be corrupted dynamically by the adversary.

Based on this general web model, we analyzed the security of the secondary IdP mode of BrowserID [13]. The analysis revealed several severe vulnerabilities, which have since been fixed by Mozilla.

Contributions of this Paper. The main contributions of this paper are that we i) analyze authentication and privacy properties for the primary mode of BrowserID, where in both cases the analysis revealed new attacks, ii) identify generic web security properties to ease future analysis efforts, and iii) slightly extend our web model.

As mentioned before, in [13], we studied the simpler secondary mode of BrowserID only. The primary model studied here is much more complex than the secondary mode (see also the remarks in Section 4.2). It involves more components (such as an arbitrary set of IdPs, more iframes), a much more complex communication structure, and requires weaker trust assumptions (for example, some IdPs, and hence, the JavaScript they deliver, might be malicious). Also, in our previous work, we have not considered privacy properties, but authentication properties only.

More specifically, the contributions of this paper can be summarized as follows.

Extension of the Web Model. We slightly extend our web model proposed in [13]. We complement the modeling of the web storage concept of modern browsers by adding `sessionStorage` [27], which is (besides the already modeled `localStorage`) heavily used by BrowserID in its primary mode. We also extend the model to include a set of user identities (e.g., user names or email addresses) in addition to user secrets.

Authentication Attack and Security Proof for BrowserID. The authentication properties we analyze are central to any SSO system and correspond to those considered in our previous work: i) the attacker should not be able to log in at an RP as an honest user and ii) the attacker should not be able to authenticate an honest user/browser to an RP with an ID not owned by the user (identity injection). While trying to prove these authentication properties for the primary mode of BrowserID, we discovered a new attack which violates property ii). Depending on the service provided by the RP, this could allow the attacker to track the honest user or to obtain user secrets. We confirmed the attack on the actual implementation and reported it to Mozilla, who acknowledged the attack. We note that this attack does not apply to the secondary mode.

We propose a fix and provide a detailed formal proof based on the (extended) web model which shows that the fixed system satisfies the mentioned authentication properties. This constitutes the most complex formal analysis of a web application based on an expressive model of the web infrastructure, in fact, as mentioned, the most comprehensive one to date. We note that other web models are too limited to be applied to BrowserID (see also Section 7).

Privacy Attacks on BrowserID. As pointed out before, BrowserID was designed by Mozilla with the explicit intention to respect users' privacy. Unlike in other SSO systems, when using BrowserID, IdPs should not learn to which RP a user logs in. When trying to formally prove this property, we discovered attacks that show that BrowserID cannot live up to this claim. Our attacks allow malicious IdPs to check whether or not a user is logged in at a specific RP with little effort. Interestingly, one variant of these attacks exploits a browser side channel which, to our knowledge, has not received much attention in the literature so far. Just as for authentication, we have confirmed the attacks on the actual implementation and reported them to Mozilla [10], who acknowledged the attacks. Unfortunately, the attacks exploit a design flaw of BrowserID that does not seem to be easily fixable without a major redesign.

Generic Web Security Properties. Our security analysis of BrowserID and the case study in [13] show that certain security properties of the web model need to be established in most security proofs for web standards and web applications. As another contribution, we therefore identify and summarize central security properties of generic web features in our extension of our model and formalize them in a general way such that they can be used in and facilitate future analysis efforts of web standards and web applications.

Structure of this Paper. In Section 2, we present the basic communication model and the web model, including our extensions. We deduce general properties of this model, which are independent of specific web applications, in Section 3. For our security analysis, we first, in Section 4, provide a description of the BrowserID system, focusing on the primary mode. We then, in Section 5, present our attack and the formal analysis of the authentication properties of the (fixed) BrowserID system in primary mode. In Section 6, we present our attacks on privacy of BrowserID. Related work is discussed in Section 7. We conclude in Section 8. Full details can be found in the appendix.

2 The Web Model

In this section, we present the model of the web infrastructure as proposed in [13], along with our extensions (sessionStorage and user identities) mentioned in the introduction, with full details, most of which taken from [13], provided in Appendices A to C. We first present the generic Dolev-Yao style communication model which the model is based on.

2.1 Communication Model

The main entities in the communication model are *atomic processes*, which will be used to model web browsers, web servers, DNS servers as well as web and network attackers. Each atomic process has a list of addresses (representing IP addresses) it listens to. A set of atomic processes forms what is called a *system*. The different atomic processes in such a system can communicate via events, which consist of a message as well as a receiver and a sender address. In every step of a run, one event is chosen non-deterministically from the current “pool” of events and is delivered to an atomic process that listens to the receiver address of that event; if different atomic processes can listen to the same address, the atomic process to which the event is delivered is chosen non-deterministically among the possible processes. The (chosen) atomic process can then process the event and output new events, which are added to the pool of events, and so on. More specifically, messages, processes, etc. are defined as follows.

Terms, Messages and Events. As usual in Dolev-Yao models (see, e.g., [1]), messages are expressed as formal terms over a signature. Later messages may, for instance, represent HTTP(S) requests and responses.

The signature Σ for the terms and messages considered in this work is the union of the following pairwise disjoint sets of function symbols: (1) constants $C = \text{IPs} \cup \mathbb{S} \cup \{\top, \perp, \diamond\} \cup \mathcal{N}$ (IPs for (IP) addresses, \mathbb{S} for ASCII strings, and \mathcal{N} for an infinite set of nonces) where the four sets are pairwise disjoint, (2) function symbols for public keys,

asymmetric/symmetric encryption/decryption, and digital signatures: $\text{pub}(\cdot)$, $\text{enc}_a(\cdot, \cdot)$, $\text{dec}_a(\cdot, \cdot)$, $\text{enc}_s(\cdot, \cdot)$, $\text{dec}_s(\cdot, \cdot)$, $\text{sig}(\cdot, \cdot)$, $\text{checksig}(\cdot, \cdot)$, $\text{extractmsg}(\cdot)$, (3) n -ary sequences $\langle \cdot \rangle$, $\langle \cdot, \cdot \rangle$, $\langle \cdot, \cdot, \cdot \rangle$, etc., and (4) projection symbols $\pi_i(\cdot)$ for all $i \in \mathbb{N}$. *Ground terms* over this signature are terms that do not contain variables. These terms represent messages. By \mathcal{M} we denote the set of messages. An *event (over IPs and \mathcal{M})* is of the form $(a:f:m)$, for $a, f \in \text{IPs}$ and $m \in \mathcal{M}$, where a is interpreted to be the receiver address and f is the sender address.

For example, $k \in \mathcal{K}$ and $\text{pub}(k)$ are messages, where k typically models a private key and $\text{pub}(k)$ the corresponding public key. For strings $a, b \in \mathbb{S}$ and the nonce $k \in \mathcal{K}$, the message $\text{enc}_a(\langle a, b \rangle, \text{pub}(k))$ is interpreted to be the message $\langle a, b \rangle$ (the sequence of strings a and b) encrypted under the public key $\text{pub}(k)$.

The *equational theory* associated with the signature Σ is defined as usual in Dolev-Yao models. It captures the meaning of the function symbols in Σ . For instance, one equation is $\text{dec}_a(\text{enc}_a(x, \text{pub}(y)), y) = x$ and another $\pi_i(\langle x_1, \dots, x_n \rangle) = x_i$ for $1 \leq i \leq n$. We have that $\pi_1(\text{dec}_a(\text{enc}_a(\langle a, b \rangle, \text{pub}(k)), k)) \equiv a$.

Atomic Processes, Systems and Runs. Atomic Dolev-Yao processes, systems, and runs of systems are defined as follows.

A (*generic*) *atomic process* is a tuple $p = (I^p, Z^p, R^p, s_0^p)$ where I^p is a set of addresses (the set of address the process listens to), Z^p is a set of states (formally, terms), $s_0^p \in Z^p$ is an initial state, and R^p is a relation that takes an event and a state as input and (non-deterministically) returns a new state and a set of events. This relation models a non-deterministic computation step of the process, which upon receiving an event in a given state non-deterministically moves to a new state and outputs a set of messages (events).

In the web model, we consider *atomic Dolev-Yao (DY) processes* only. For these processes it is required that the events and states that they output can be computed (more formally, derived in the usual Dolev-Yao style) from the current input event and state (see Appendix A). The rest of this paper will consider DY processes only.

The so-called *attacker process* is an atomic DY process which records all messages it receives and outputs all messages it can possibly derive from its recorded messages. Hence, an attacker process is the maximally powerful DY process. It carries out all attacks any DY process could possibly perform and is parametrized by the set of sender addresses it may use.

A *system* is a (possibly infinite) set of atomic processes. Its state (i.e., the states of all atomic processes in the system) together with a multi-set of waiting events is called a *configuration*.

A *run* of a system for an initial set E_0 of events is a sequence of configurations, where each configuration (except for the first one, which consists of E_0 and the initial states of the atomic processes) is obtained by delivering one of the waiting events of the preceding configuration to an atomic process p (which listens to the receiver address of the event), and which in turn performs a computation step according to its relation R^p .

2.2 Scripting Processes

For the web model, we also define scripting processes, which model client-side scripting technologies, such as JavaScript.

A *scripting process* (or simply, a *script*) is defined similarly to a DY process. It is called by the browser in which it runs. The browser provides it with a (fresh, infinite) set N of nonces and state information s . The script then outputs a term s' , which represents the new internal state and some command which is interpreted by the browser (see Section 2.5 for details). Again, it is required that a script's output s' is derivable from its input (s, N) .

Similarly to an attacker process, the so-called *attacker script* R^{att} may output everything that is derivable from the input.

2.3 Web System

In [13], we formalize the web infrastructure and web applications by what they call a web system. A web system contains, among others, a (possibly infinite) set of DY processes, modeling web browsers, web servers, DNS servers, and attackers (which may corrupt other entities, such as browsers).

Web System. A *web system* is a tuple $(\mathcal{W}, \mathcal{S}, \text{script}, E_0)$ with its components defined as follows:

The first component, \mathcal{W} , denotes a system (a set of DY processes) and is partitioned into the sets Hon, Web, and Net, where in Hon the set of honest DY processes and in Web and Net attacker processes (see Section 2.1) are specified, *web attacker* and *network attacker processes*, respectively. While a web attacker can listen to and send messages from its own addresses only, a network attacker may listen to and spoof all addresses. Hence, it is the maximally powerful attacker. Attackers may corrupt other parties. In the analysis of a concrete web system, we typically have one network attacker only and no web attackers (as they are subsumed by the network attacker) or one or more web attackers but then no network attacker. Honest processes (in Hon) can either be web servers, web browsers, or DNS servers. In our security analysis of authentication properties, DNS servers will be subsumed by the attacker, and hence, we do not need to model them for the analysis of these properties. Our attacks on privacy work with honest DNS servers. As the details of the modeling of these servers is not essential to understand these attacks, we refer to [13] for the model of DNS servers. The modeling of a *web server* heavily depends on the specific web application. Our concrete models for the web servers of the BrowserID system are provided in Sections 4 and following. Below, we present the modeling of web browsers, including our extensions, which is independent of a specific web application, with full details provided in Appendices B and C.

The second component, \mathcal{S} , is a finite set of scripts, which include the attacker script $R^{\text{att}} \in \mathcal{S}$. In a concrete model of a web application, such as our BrowserID model, the set $\mathcal{S} \setminus \{R^{\text{att}}\}$ typically describes the set of honest scripts used in the considered application. Malicious scripts are modeled by the “worst-case” malicious script, R^{att} .

The third component, *script*, is an injective mapping from \mathcal{S} to \mathbb{S} , i.e., by script every $s \in \mathcal{S}$ is assigned its string representation *script*(s). Finally, E_0 is a multi-set of events,

containing an infinite number of events of the form $(a:a:\text{TRIGGER})$ for every process a in the web system. A *run* of the web system is a run of \mathcal{W} initiated by E_0 .

2.4 HTTP Messages

HTTP requests are represented as ground terms containing a nonce, a method (e.g., GET or POST), a domain name, a path, URL parameters, request headers (such as Cookie), and a message body. For example, a GET request for the URL <http://ex.com/show?p=1> can be modeled as the term

$$r := \langle \text{HTTPReq}, n_1, \text{GET}, \text{ex.com}, /show, \langle \langle p, 1 \rangle \rangle, \langle \rangle, \langle \rangle \rangle$$

where headers and body are empty. A response contains a nonce (the same as in the request), a status code, headers, and a body. A response to r would be

$$s := \langle \text{HTTPResp}, n_1, 200, \langle \langle \text{Set-Cookie}, \langle \text{SID}, \langle n_2, \perp, \top, \perp \rangle \rangle \rangle \rangle, \langle \text{script1}, \text{init} \rangle \rangle$$

where $\langle \text{SID}, \langle n_2, \perp, \top, \perp \rangle \rangle$ is a cookie with the name/value pair $\text{SID} = n_2$ and the attributes `httpOnly`, `secure`, `session` set or not set, and $\langle \text{script1}, \text{init} \rangle$ is the body, in this case an HTML document that is to be delivered to the browser (modeled by the string representation of a script and its initial state, see below).

A corresponding HTTPS request for r as above would be $\text{enc}_a(\langle r, k' \rangle, \text{pub}(k_{\text{ex.com}}))$, where k' is a fresh symmetric key (a nonce) generated by the sender of the request. The responder is supposed to use this key to encrypt the response, which, hence, is of the form $\text{enc}_s(s, k')$.

2.5 Web Browsers

An honest browser is thought to be used by one honest user. The honest user is modeled as part of the browser. User actions are modeled as non-deterministic actions of the web browser. For example, the web browser itself can non-deterministically follow the links provided by a web page. User data (i.e., passwords and identities) is stored in the initial state of the browser (see below) and is given to a web page when needed, similar to the AutoFill feature in browsers. As detailed below, browsers can be corrupted, i.e., taken over by web and network attackers.

A web browser p is modeled as a DY process $(I^p, Z^p, R^p, s_0^p, N^p)$ where $I^p \subseteq \text{IPs}$ is a finite set of addresses p may listen to and $N^p \subseteq \mathcal{N}$ is an infinite set of nonces p may use. The set of states Z^p , the initial state s_0^p , and the relation R^p are defined next.

Browser State: Z^p and s_0^p . The set Z^p of states of a browser consists of terms of the form

$$\langle \text{windows}, \text{ids}, \text{secrets}, \text{cookies}, \text{localStorage}, \text{sessionStorage}, \text{keyMapping}, \text{sts}, \text{DNSaddress}, \text{nonces}, \text{pendingDNS}, \text{pendingRequests}, \text{isCorrupted} \rangle.$$

Windows and documents. The most important part of the state are windows and documents, both stored in the subterm *windows*. A browser may have several windows open

at any time (resembling the tabs and windows in a real browser), each containing a list of documents (the history of visited web pages) of which one is “active”, namely the one currently presented to the user in that window. A window may be navigated forward and backward (modeling navigation buttons), deactivating one document and activating its successor or predecessor. Intuitively, a document represents a loaded HTML page. More formally, a document contains (the string representation of) a script, which is meant to model both the static HTML code (e.g., links and forms) as well as JavaScript code. When called by the browser, a script outputs a command which is then interpreted by the browser, such as following a link or issuing an XHR (see below). Documents may also contain iframes, which are represented as windows (*subwindows*) nested inside of document terms. This creates a tree of windows and documents.

Secrets and IDs. This subterm holds the secrets and the identities of the user of the web browser. Secrets (such as passwords) are modeled as nonces and they are indexed by origins (where an origin is a domain name plus the information whether the connection to this domain is via HTTP or HTTPS). Secrets are only released to documents (scripts) with the corresponding origin, similarly to the AutoFill mechanism in browsers. Identities are arbitrary terms that model public information of the user’s identity, such as email addresses. Identities are released to any origin. As mentioned in the introduction, identities were not considered in [13].

Cookies, localStorage, and sessionStorage. These subterms contain the cookies (indexed by domains), localStorage data (indexed by origins), and sessionStorage data (indexed by origins and top-level window references) stored in the browser. As mentioned in the introduction, sessionStorage was not modeled in [13].

KeyMapping. This term is the equivalent to a certificate authority (CA) certificate store in the browser. Since, for simplicity, the model currently does not formalize CAs, this term simply encodes a mapping assigning domains $d \in \text{Doms}$ to their respective public keys $\text{pub}(k_d)$.

STS. Domains that are listed in this term are contacted by the web browser over HTTPS only. Connection attempts over HTTP are transparently rewritten to HTTPS requests. Servers can employ the Strict-Transport-Security header to add their domain to this list.

DNSAddress. This term defines the address of the DNS server used by the browser.

Nonces, pendingDNS, and pendingRequests. These terms are used for bookkeeping purposes, recording the nonces that have been used by the browser so far, the HTTP(S) requests that await successful DNS resolution, and HTTP(S) requests that await a response, respectively.

IsCorrupted. This term indicates whether the browser is corrupted ($\neq \perp$) or not ($= \perp$). A corrupted browser behaves like a web attacker.

Initial state s_0^p of a web browser. In the browser’s initial state, *keyMapping*, *DNSAddress*, *secrets*, and *ids* are defined as needed, *isCorrupted* is set to \perp , and all other subterms are $\langle \rangle$.

Web Browser Relation R^p . This relation, outlined in Figure 1, specifies how the web browser processes incoming messages. The browser may receive special messages that

PROCESSING INPUT MESSAGE m

$m = \text{FULLCORRUPT}$: $isCorrupted := \text{FULLCORRUPT}$
 $m = \text{CLOSECORRUPT}$: $isCorrupted := \text{CLOSECORRUPT}$
 $m = \text{TRIGGER}$: non-deterministically choose $action$ from $\{1, 2\}$
 $action = 1$: Call script of some active document. Outputs new state and $command$.
 $command = \text{HREF}$: \rightarrow *Initiate request*
 $command = \text{IFRAME}$: Create subwindow, \rightarrow *Initiate request*
 $command = \text{FORM}$: \rightarrow *Initiate request*
 $command = \text{SETSCRIPT}$: Change script in given document.
 $command = \text{SETSCRIPTSTATE}$: Change state of script in given document.
 $command = \text{XMLHTTPREQUEST}$: \rightarrow *Initiate request*
 $command = \text{BACK}$ or FORWARD : Navigate given window.
 $command = \text{CLOSE}$: Close given window.
 $command = \text{POSTMESSAGE}$: Send postMessage to specified document.
 $action = 2$: \rightarrow *Initiate request to some URL in new window*
 $m = \text{DNS response}$: send corresponding HTTP request
 $m = \text{HTTP(S) response}$: (decrypt,) find reference.
reference to window: create document in window
reference to document: add response body to document's script input

Fig. 1. The basic structure of the web browser relation R^p with an extract of the most important processing steps, in the case that the browser is not already corrupted.

cause it to become corrupted (first two lines in Figure 1), in which case it acts like the attacker process. There are two types of corruption: If the browser gets fully corrupted, the attacker learns the entire current state of the browser. If it gets close-corrupted, any open windows, documents and used nonces (in particular, HTTPS encryption keys) are discarded from the browser's state before it is handed over to the attacker. This models that a user closed the browser, but a malicious user now uses the browser (and all information left in the browser's state).

The browser can receive a special trigger message TRIGGER, upon which the browser non-deterministically chooses one of two actions: i) Select one of the current documents, trigger its JavaScript, and evaluate the output of the script. Scripts can change the state of the browser (e.g., by setting cookies) and can trigger specific actions (e.g., following a link or creating an iframe), which are modeled as *commands* issued by the script (see the list in Figure 1). ii) Follow some URL, with the intuition that it was entered by the user.

As mentioned, some of the above actions can cause the browser to generate new HTTP(S) requests. In this case, the browser first asks the configured DNS server for the IP address belonging to the domain name in the HTTP(S) request. As soon as the DNS response arrives, the browser sends the HTTP(S) request to the respective IP address.

If the HTTP(S) response arrives, its headers are evaluated and the body of the request becomes the script of a newly created document that is then inserted at an appropriate place in the window/document tree. However, if the HTTP(S) response is a response to an XHR (triggered by a script in a document), the body of the response is added to the corresponding document and can later be processed by the script of that document.

3 General Security Properties

We have identified central application independent security properties of web features in the web model and formalized them in a general way such that they can be used in and facilitate future analysis efforts of web standards and web applications. In this section, we provide a brief overview of these properties, with precise formulations and proofs presented in Appendix D.

The first set of properties concerns encrypted connections (HTTPS): We show that HTTP requests that were encrypted by an honest browser for an honest receiver cannot be read or altered by the attacker (or any other party). This, in particular, implies correct behavior on the browser’s side, i.e., that browsers that are not fully corrupted never leak a symmetric key used for an HTTPS connection to any other party. We also show that honest browsers set the host header in their requests properly, i.e., the header reflects an actual domain name of the receiver, and that only the designated receiver can successfully respond to HTTPS requests.

The second set of properties concerns origins and origin headers. Using the properties stated above, we show that browsers cannot be fooled about the origin of an (HTTPS) document in their state: If the origin of a document in the browser’s state is a secure origin (HTTPS), then the document was actually sent by that origin. Moreover, for requests which contain an origin header with a secure origin we prove that such requests were actually initiated by a script that was sent by that origin to the browser. In other words, in this case, the origin header works as expected.

4 The BrowserID System

BrowserID [22] is a decentralized single sign-on (SSO) system developed by Mozilla for user authentication on web sites. It is a complex full-fledged web application deployed in practice, with currently ~47k LOC (excluding some libraries). It allows web sites to delegate user authentication to email providers, identifying users by their email addresses. BrowserID makes use of a broad variety of browser features, such as XHRs, postMessage, local- and sessionStorage, cookies, various headers, etc.

We first, in Section 4.1, provide a high-level overview of the BrowserID system. A more detailed description of the BrowserID implementation is then given in Section 4.2. The description of the BrowserID system presented in the following as well as our BrowserID model (see Section 5.1) is extracted mainly from the BrowserID source code [20] and the (very high-level) official BrowserID documentation [22].

4.1 Overview

The BrowserID system knows three distinct parties: the user, who wants to authenticate herself using a browser, the relying party (RP) to which the user wants to authenticate (log in) with one of her email addresses (say, `user@idp.com`), and the identity/email address provider, the IdP. If the IdP (`idp.com`) supports BrowserID directly, it is called a *primary IdP*. Otherwise, a Mozilla-provided service, the so-called *secondary IdP*, takes the role of the IdP. As mentioned before, here we concentrate on the primary IdP

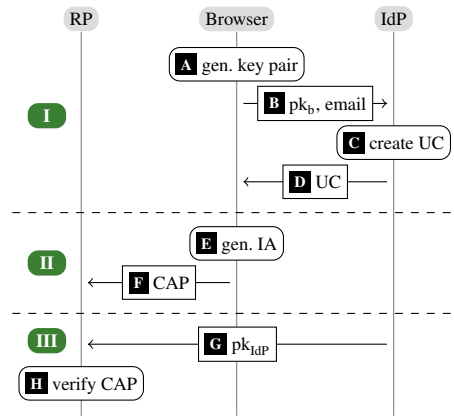


Fig. 2. BrowserID login: basic overview

mode as the secondary IdP mode was described in detail in [13]. However, we briefly discuss the differences between the two modes at the end of Section 4.2.

A primary IdP provides information about its setup in a so-called *support document*, which it provides at a fixed URL derivable from the email domain, e.g., <https://idp.com/.well-known/browserid>.

A user who wants to log in at an RP with an email address for some IdP has to present two signed documents to the RP: A *user certificate* (UC) and an *identity assertion* (IA). The UC contains the user’s email address and the user’s public key. It is signed by the IdP. The IA contains the origin of the RP and is signed with the user’s private key. Both documents have a limited validity period. A pair consisting of a UC and a matching IA is called a *certificate assertion pair* (CAP) or a *backed identity assertion*. Intuitively, the UC in the CAP tells the RP that (the IdP certified that) the owner of the email address is (or at least claims to be) the owner of the public key. By the IA contained in the CAP the RP is ensured that the owner of the given public key (i.e., the one who knows the corresponding private key) wants to log in. Altogether, given a valid CAP, RP would consider the user (identified by the email address in the CAP) to be logged in.

The BrowserID authentication process (with a primary IdP) consists of three phases (see Figure 2): (I) UC provisioning, (II) CAP creation, and (III) CAP verification.

In Phase (I), (the browser of) the user creates a public/private key pair [A]. She then sends her public key as well as the email address she wants to use to log in at some RP to the respective IdP [B]. The IdP now creates the UC [C], which is then sent to the user [D]. The above requires the user to be logged in at IdP.

With the user having received the UC, Phase (II) can start. The user wants to authenticate to an RP, so she creates the IA [E]. The UC and the IA are concatenated to a CAP, which is then sent to the RP [F].

In Phase (III), the RP checks the authenticity of the CAP. For this purpose, the RP could use an external verification service provided by Mozilla or check the CAP itself as follows: First, the RP fetches the public key of the IdP [G], which is contained in the support document. Afterwards, the RP checks the signatures of the UC, and the IA [H]. If

this check is successful, the RP can, as mentioned before, consider the user to be logged in with the given email address and send her some token (e.g., a cookie with a session ID), which we refer to as an *RP service token*.

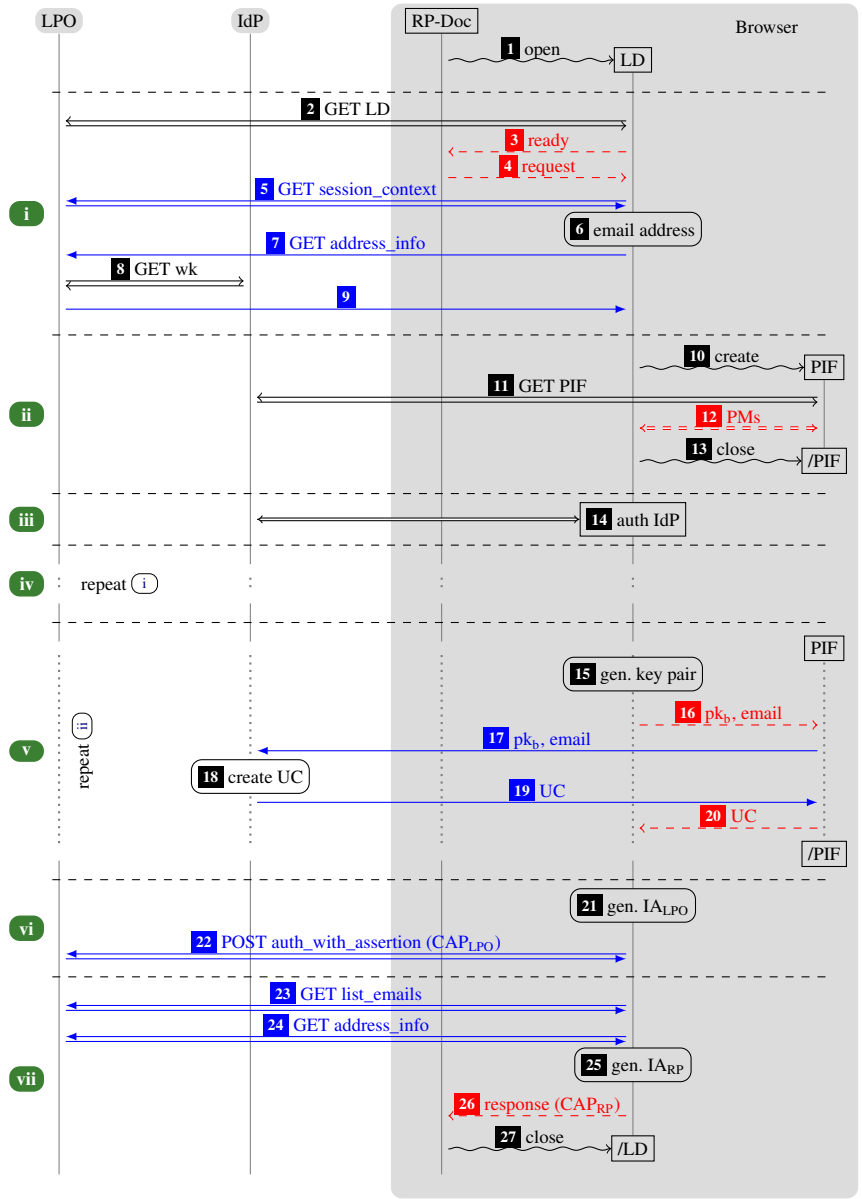
4.2 Implementation Details

We now provide a more detailed description of the BrowserID implementation. Since the system is very complex, with many HTTPS requests, XHRs, and postMessages sent between different entities (servers as well as windows and iframes within the browser), we here describe mainly the phases of the login process without explaining every single message exchange done in the implementation. A more detailed step-by-step description can be found in Appendix E. Note that BrowserID’s specification of IdPs fixes the interface to BrowserID only, but otherwise does not further detail the specification of IdPs. Therefore, in what follows, we consider a typical IdP, namely the example implementation provided by Mozilla [20].

In addition to the parties mentioned so far, the actual BrowserID implementation uses another party, Mozilla’s `login.persona.org` (LPO). Among others, LPO provides HTML and JavaScript files that, for security and privacy reasons, cannot be delivered by either IdP or RP. An overview of the implementation is given in Figure 3. For brevity of presentation, several messages and components, such as the CIF (see below), are omitted in the figure (see Figure 8 on Pages 85 and 86 for a detailed version of Figure 3).

Windows and iframes in the Browser. By *RP-Doc* we denote the window containing the document loaded from some RP, at which the user wants to log in with an email address hosted by some IdP. RP-Doc typically includes JavaScript from LPO and contains a button “Login with BrowserID”. The LPO JavaScript running in RP-Doc opens an auxiliary window called the *login dialog* (LD). Its content is provided by LPO and it handles the interaction with the user. During the login process, a temporary invisible iframe called the *provisioning iframe* (PIF) can be created in the LD. The PIF is loaded from IdP. It is used by LD to communicate (cross-origin) with the IdP via postMessages: As the BrowserID implementation mainly runs under the origin of LPO, it cannot directly communicate with the IdP, thus it uses the PIF as a proxy. Temporarily, the LD may navigate itself to a web page at IdP to allow for direct user interaction with the IdP. We then call this window the *authentication dialog* (AD).

Login Process. To describe the login process, for the sake of presentation we assume for now that the user uses a “fresh” browser, i.e., the user has not been logged in before. As mentioned, the process starts by the user visiting a web site of some RP. After the user has clicked on the login button in RP-Doc, the LD is opened and the interactive login flow is started. We can divide this login flow into seven phases: In Phase (i), the LD is initialized and the user is prompted to provide her email address. Also, LD fetches the support document (see Section 4.1) of the IdP via LPO. In Phase (ii), LD creates the PIF from the *provisioning URL* provided in the support document. As (by our assumption) the user is not logged in yet, the PIF notifies LD that the user is not authenticated to the IdP. In Phase (iii), LD navigates itself away to the *authentication URL* which is also provided in the support document and links to the IdP. Usually, this document will show a login form in which the user enters her password to authenticate to the IdP. After the



→ HTTPS messages, → XHRs (over HTTPS), - → postMessages, ~ → browser commands

Fig. 3. Simplified BrowserID implementation overview. CIF omitted for brevity.

user has been authenticated to IdP (which typically implies that the IdP sets a session cookie in the browser), the window is navigated back to LPO.

Now, the login flow continues in Phase (iv), which basically repeats Phase (i). However, the user is not prompted for her email address (it has previously been saved in the localStorage under the origin of LPO along with a nonce, where the nonce is stored in the sessionStorage). In Phase (v), which essentially repeats Phase (ii), the PIF detects that the user is now authenticated to the IdP and the provisioning phase is started (I in Figure 2): The user’s keys are created by LD and stored in the localStorage under the origin of LPO. The PIF forwards the certification request to the IdP, which then creates the UC and sends it back to the PIF. The PIF in turn forwards it to the LD, which stores it in the localStorage under the origin of LPO.

In Phases (vi) and (vii), mainly the IA is generated by LD for the origin of RP-Doc and sent (together with the UC) to RP-Doc (II in Figure 2). In the localStorage, LD stores that the user’s email address is logged in at RP. Moreover, to log the user in at LPO, LD generates an IA for the origin of LPO and sends the UC and IA to LPO.

Automatic CAP Creation. In addition to the interactive login presented above, BrowserID also contains an automatic, non-interactive way for RPs to obtain a freshly generated CAP: During initialization within RP-Doc, an invisible iframe called the *communication iframe* (CIF) is created inside RP-Doc. The CIF’s JavaScript is loaded from LPO and behaves similar to LD, but without user interaction. The CIF automatically issues a fresh CAP and sends it to RP-Doc under specific conditions: among others, the email address must be marked as logged in at RP in the localStorage. If necessary, a new key pair is created and a corresponding new UC is requested at the IdP. For this purpose, a PIF is created inside the CIF.

Differences to the Secondary IdP Mode. In the secondary IdP mode there are three parties involved only: RP, Browser, and LPO, where LPO also takes the role of an IdP; LPO is the only IdP that is present, rather than an arbitrary set of (external) IdPs. Consequently, in the secondary IdP mode the PIF and the AD do not exist. Moreover, in the primary mode, the behavior of the CIF and the LD is more complex than in the secondary mode. For example, in the primary mode, just like the LD, the CIF might contain a PIF (iframe in iframe) and interact with it via postMessages. Altogether, the secondary IdP case requires much less communication between parties/components and trust assumptions are simpler: in the secondary IdP mode LPO (which is the only IdP in this mode) has to be trusted, in the primary IdP mode some external IdPs might be malicious (and hence, also the scripts they deliver for the PIF and the AD). To illustrate the difference between the secondary and the primary IdP mode, in the appendix both modes are illustrated in more detail, see Figure 9 on Page 87 for the secondary IdP mode and Figure 8 on Pages 85 and 86 for the primary IdP mode.

5 Analysis of BrowserID: Authentication Properties

In this section, we present the analysis of the BrowserID system with primary IdPs and with respect to authentication properties. As already mentioned, in [13], we analyzed the simpler case with a secondary IdP. We first, in Section 5.1, describe our model of BrowserID with primary IdPs, with two central authentication properties one would

expect any SSO system to satisfy formalized in Section 5.2. Due to the many differences between the secondary and primary mode as described above, the model had to be written from scratch in most parts. As mentioned in the introduction, during the analysis of BrowserID it turned out that one of the security properties is not satisfied and that in fact there is an attack on BrowserID. We confirmed that this attack, which was acknowledged by Mozilla, works on the actual implementation of BrowserID. In Section 5.3, the attack is presented along with a fix. (Our BrowserID model presented in Appendix F contains this fix.) In Section 5.4, we prove that the fixed BrowserID system with primary IdPs satisfies both authentication properties.

5.1 Modeling of BrowserID with Primary IdPs

We model the BrowserID system with primary IdPs as a web system (in the sense of Section 2). Note that while in Section 4 we give only a brief overview of the BrowserID system, our modeling and analysis considers the complete system with primary IdPs, where we have extracted the model from the BrowserID source code [20].

We call a web system $\mathcal{BID} = (\mathcal{W}, \mathcal{S}, \text{script}, E_0)$ a *BrowserID web system* if it is of the form described in Appendix F and briefly outlined here.

The system $\mathcal{W} = \text{Hon} \cup \text{Web} \cup \text{Net}$ consists of the (network) attacker process attacker, the web server for LPO, a finite set B of web browsers, a finite set RP of web servers for the relying parties, and a finite set IDP of web servers for the identity providers, with $\text{Hon} := B \cup RP \cup IDP \cup \{\text{LPO}\}$, $\text{Web} := \emptyset$, and $\text{Net} := \{\text{attacker}\}$. DNS servers are assumed to be dishonest, and hence, are subsumed by attacker. IdPs and RPs can become corrupted (similar to browsers, by a special message); LPO is assumed to be honest.

The set IPs of IP addresses (see Section 2.1) contains one address for each party in \mathcal{W} . The set $\text{Doms} \subseteq \mathbb{S}$ contains one or more domains for each party in \mathcal{W} , except for browsers.

The definition of the processes in \mathcal{W} follows the description in Section 4.2. For RP , we explicitly follow the security considerations in [22] (Cross-site Request Forgery protection, e.g., by checking origin headers and HTTPS only with STS enabled). When RP receives a valid CAP (see below), RP responds with a fresh *RP service token for ID i* where i is the ID (email address) for which the CAP was issued. Intuitively, a client having such a token can use the service of the RP .

Each browser $b \in B$ owns a set of email addresses (identities) of the form $\langle name, d \rangle$ with $name \in \mathbb{S}$ and $d \in \text{Doms}$ (belonging to an IdP) and associated passwords (i.e., nonces).

A UC uc for a user u with email address $\langle name, d \rangle$ and public key (verification key) $\text{pub}(k_u)$, where $d \in \text{dom}(y)$ is a domain of the IdP y that issued the UC and k_u is the private (signing) key of u , is a term of the form $uc = \text{sig}(\langle \langle name, d \rangle, \text{pub}(k_u) \rangle, \text{signkey}(y))$, with $\text{signkey}(y)$ being the signing key of y . An IA ia for an origin o is a message of the form $ia = \text{sig}(o, k_u)$. A CAP is of the form $\langle uc, ia \rangle$. Note that time stamps are omitted both from the UC and the IA, modeling that UC and IA never expire. In reality, as explained in Section 4, they are valid for a certain period of time. So our modeling is a safe overapproximation.

The set \mathcal{S} of BID contains six scripts, with their string representations defined by script: the honest scripts running in RP-Doc, CIF, LD, AD, and PIF, respectively, and the malicious script R^{att} . The scripts for CIF and LD (issued by LPO) are defined in a straightforward way following the implementation outlined in Section 4. The script for RP-Doc (issued by RP) also includes the script that is (in reality) loaded from LPO. In particular, this script creates the CIF and the LD (sub)windows, whose contents (scripts) are loaded from LPO. The scripts for the AD and PIF are modeled following the example implementation provided by Mozilla [20]. Full formal specifications of all the above mentioned scripts are provided in Appendix F.

5.2 Authentication Properties of the BrowserID System

While the documentation of BrowserID does not contain explicit security goals, here we state two fundamental authentication properties every SSO system should satisfy. These properties are adapted from [13].

Informally, these properties can be stated as follows: **(A)** *The attacker should not be able to use a service of RP as an honest user.* In other words, the attacker should not get hold of (be able to derive from his current knowledge) an RP service token for an ID of an honest user (browser), even if the browser was closed and then later used by a malicious user (i.e., after a `CLOSECORRUPT`). **(B)** *The attacker should not be able to authenticate an honest browser to an RP with an ID that is not owned by the browser (identity injection).* We refer the reader to Appendix G for the formal definition of these properties.

We call a BrowserID web system BID secure (w.r.t. authentication) if the above conditions are satisfied in all runs of the system.

5.3 Identity Injection Attack on BrowserID with Primary IdPs

While trying to prove the above mentioned authentication properties of BrowserID with primary IdPs in our model, we discovered a serious attack, which is sketched below and does not apply to the case with secondary IdPs. We confirmed the attack on the actual implementation and reported it to Mozilla [9], who acknowledged it.

During the provisioning phase (v) (see Figure 3), the IdP issues a UC for the user's identity and public key provided in [16]. This UC is sent to the LD by the PIF in [20].

If the IdP is malicious, it can issue a UC with different data. In particular, it could replace the email address by a different one, but keep the original public key. This (malicious) UC is then later included in the CAP by LD. The CAP will still be valid, because the public key is unchanged. Now, as the RP determines the user's identity by the UC contained in the CAP, RP issues a service token for the spoofed email address. As a result, the honest user will use RP's service (and typically will be logged in to RP) under an ID that belongs to the attacker, which, for example, could allow the attacker to track actions of the honest user or obtain user secrets. This violates Condition **(B)**.

To fix this problem, upon receipt of the UC in [20], LD should check whether it contains the correct email address and public key, i.e., the one requested by LD in [16]. The same is true for the CIF, which behaves similarly to the LD. The formal model of BrowserID presented in Appendix F contains these fixes.

5.4 Security of the Fixed System

For the fixed BrowserID system with primary IdPs, we have proven the following theorem, which says that a fixed BrowserID web system (i.e., the system where the above described fix is applied) satisfies the security properties (A) and (B).

Theorem 1. *Let BID be a fixed BrowserID web system. Then, BID is secure (w.r.t. authentication).*

We prove Conditions (A) and (B) separately. For both conditions, we assume that they are not satisfied and lead this to a contradiction. In our proofs, we make use of the general security properties of the web model presented in Section 3, which helped a lot in making the proof for the primary IdP model more modular and concise. The complete proof with all details is provided in Appendix H.

6 Privacy of BrowserID

In this section, we study the privacy guarantees of the BrowserID system with primary IdPs. Regarding privacy, Mozilla states that “...the BrowserID protocol never leaks tracking information back to the Identity Provider.” [5] and “Unlike other sign-in systems, BrowserID does not leak information back to any server [...] about which sites a user visits.” [19].¹ While this is not a formal definition of the level of privacy that BrowserID is supposed to provide, these and other statements² make it certainly clear that, unlike for other SSO systems, IdPs should not be able to learn to which RPs their users log in.

In the process of formalizing this intuition in our model of BrowserID and trying to prove this property, we found severe attacks against the privacy of BrowserID which made clear that BrowserID does not provide even a rather weak privacy property in the presence of a malicious IdP. Intuitively, the property says that a malicious IdP (which acts as a web attacker) should not be able to tell whether a user logs in at an honest RP r or some other honest RP r' . In other words, a run in which the user logs in at r at some point should be indistinguishable (from the point of view of the IdP) from the run in which the user logs in at r' instead. Indistinguishability means that the two sequences of messages received by the web attacker in the two runs are statically equivalent in the usual sense of Dolev-Yao models (see [1]), i.e., a Dolev-Yao attacker cannot distinguish between the two sequences. Details of the privacy definition are not important here since our attacks clearly show that privacy is broken for any reasonable definition of privacy. Unfortunately, our attacks are not caused by a simple implementation error, but rather a fundamental design flaw in the BrowserID protocol. Fixes for this flaw are conceivable, but not without major changes to the design of BrowserID as discussed in Section 6.2. Such a redesign of BrowserID and a proof of privacy of the redesigned system are therefore out of the scope of this paper, which focuses on the existing and deployed version of BrowserID.

¹Clearly, in the current state of BrowserID a malicious LPO server could gather information about users' log in history. However, an integration of the code currently delivered by LPO into the browser, as envisioned, would avoid this issue. Currently, Mozilla's LPO needs to be trusted.

²see, e.g., https://developer.mozilla.org/en-US/Persona/Why_Persona and <https://identity.mozilla.com/post/7669886219>.




Fig. 4. The three main steps of the privacy attack. Using a specially crafted PIF document, a malicious IdP can notify itself via `postMessage` when the user is logged in at some RP r .

6.1 Privacy Attacks on BrowserID

For our attacks to work, it suffices that the IdP is a web attacker. They work even if all DNS servers, RPs, and LPO are honest, and all parties use encrypted connections. In what follows, we present several variants of attacks on privacy.

PostMessage-Based Attack. The adversary is a malicious IdP that is interested to learn whether a user is logged in at RP r . Figure 4 illustrates the main steps:

Step a. First, the victim visits her IdP. In BrowserID, email providers serve as IdPs, and therefore it is not unlikely that a user visits this web site (e.g., for checking email or to use other services). As the IdP usually has some cookie set at the user’s browser, it learns the identity of the victim. The IdP now creates a hidden iframe containing the login page of r .

Step b. The login page of r (now loaded as an iframe within IdP’s web site) includes and runs the BrowserID script. As defined in the BrowserID protocol, the script creates the communication iframe (see “Automatic CAP Creation” in Section 4.2), which in turn checks whether the email address is marked as logged in at r in the `localStorage` of the user’s browser. Only then it will try to create a new CAP, for which it needs a PIF (the same as in Phase (ii) in Figure 3).

Step c. The PIF is loaded from the IdP. Note that from this action alone, the IdP does not learn where the user wants to log in. However, instead of the original (honest) PIF document, the IdP can send a modified one that sends a `postMessage` to the parent of the parent of its own window, which in this setting is the IdP document that was opened by the user in Step a. When the IdP receives this message in the document from Step a, it knows that the PIF was loaded, and therefore, that the user is currently logged in at r .

Note that the IdP can repeatedly apply the above as long as the user stays on the IdP’s web site. During this period, the IdP can see whether or not the user is logged in at the targeted RP. Clearly, the IdP can simultaneously run the attack for different RPs in order to track the user’s login status for all such RPs. In particular, the IdP can distinguish whether a user is logged in at RP r or r' , which violates the privacy property sketched above. In our formal model, the malicious IdP would run the attacker script

R^{att} in `idp.com/index` and in `idp.com/pif` (see Figure 4) in order to carry out the attack.

Variant 1: Waiting for UC requests. The IdP first acts as in Step [a](#). Now, it could passively wait for incoming requests for the PIF document or UC requests on its server, which tell the IdP that a provisioning flow (probably initiated by Step [a](#)) was started. This variant cannot be executed in parallel and is less reliable in practice, though.

Variant 2: PIF as Attack Source. Step [a](#) can also be launched from within a PIF itself (i.e., the PIF also takes the role of `idp.com/index` above). This way, while the user logs in at some r_1 , the IdP could check whether the user is logged in at r_2 , for any r_2 .

Variant 3: Scanning the Window Structure (I). Instead of using a `postMessage` to alert the IdP's outer document about the existence of the inner PIF document, the outer document could as well repeatedly scan the window tree of the `iframe` containing r 's web site: While the IdP sees almost no information about r 's document in the `iframe` (as it is not same origin), it can see the list of subwindows (i.e., the CIF, and possibly other `iframes`). For these frames, again, it would see the subwindows, especially the PIF, which it could identify uniquely by checking whether it is same origin with the IdP's outer window.

Variant 4: Scanning the Window Structure (II). In Variant 2, using a same-origin check, the malicious IdP can uniquely identify the PIF in the window structure. This same-origin check could be skipped and it could only be checked whether a PIF is generated, based on the window structure alone. While this is less reliable, this attack could be launched by *any* third party web attacker (not only the IdP to which the user's email address belongs) to check whether the victim is logged in at r or not.

We verified (all variants of) the attacks in our model as well as in a real-world BrowserID setup. Implementing proofs-of-concept required only a few lines of (trivial) JavaScript. In most attack variants, we directly or indirectly use the structure of the windows inside the web browser as a side channel. To our knowledge, this is the first description of this side channel for breaking privacy in browsers. The attacks have been reported to and confirmed by Mozilla [10].

6.2 Fixing the Privacy of BrowserID

Fixing the privacy of BrowserID seems to require a substantial redesign of the system. Regarding the presented attacks, BrowserID's main weakness is the window structure. The most obvious mitigation, modifying the CIF such that it always creates the PIF (even if the user has not logged in before), does not work: To open the PIF, the CIF looks up (in the `localStorage`) the user's identity at the current RP to derive the address of the PIF. If the user has not logged in before, this information is not available.

Another approach would be to use cross-origin XHRs to replace the features of the PIF. This solution would require a major revision in the inner workings of BrowserID and would not protect against Variant 1.

7 Related Work

The formal treatment of the security of the web infrastructure and web applications based on this infrastructure is a young discipline. Of the few works in this area even less are based on a general model that incorporates essential mechanisms of the web.

Early works in formal web security analysis (see, e.g., [3, 11, 16, 17, 25]) are based on very limited models developed specifically for the application under scrutiny. The first work to consider a general model of the web, written in the finite-state model checker Alloy, is the work by Akhawe et al. [2]. Inspired by this work, Bansal et al. [6, 7] built a more expressive model, called WebSpi, in ProVerif [8], a tool for symbolic cryptographic protocol analysis. These models have successfully been applied to web standards and applications. Recently, Kumar [18] presented a high-level Alloy model and applied it to SAML single sign-on. However, compared to our model in [13] and its extensions considered here, on the one hand, all above mentioned models are formulated in the specification languages of specific analysis tools, and hence, are tailored towards automation (while we perform manual analysis). On the other hand, the models considered in these works are much less expressive and precise. For example, these models do not incorporate a precise handling of windows, documents, or iframes; cross-document messaging (postMessages) or session storage are not included at all. In fact, several general web features and technologies that have been crucial for the analysis of BrowserID are not supported by these models, and hence, these models cannot be applied to BrowserID. Moreover, the complexity of BrowserID exceeds that of the systems analyzed in these other works in terms of the use of web technologies and the complexity of the protocols. For example, BrowserID in primary mode is a protocol consisting of 48 different (network and inter-frame) messages compared to typically about 10–15 in the protocols analyzed in other models.

The BrowserID system in the primary mode has been analyzed before using the AuthScan tool developed by Bai et al. [4]. Their work focusses on the automated extraction of a model from a protocol implementation. This tool-based analysis did not reveal the identity injection attack, though; privacy properties have not been studied there. Dietz and Wallach demonstrated a technique to secure BrowserID when specific flaws in TLS are considered [12].

8 Conclusion

In this paper, we slightly extended our existing web model, resulting in the most comprehensive model of the web so far. It contains many security-relevant features and is designed to closely mimic standards and specifications for the web. As such, it constitutes a solid basis for the analysis of a broad range of web standards and applications.

Based on this model, we presented a detailed analysis of the BrowserID SSO system in the primary IdP mode. During the security proof of the fundamental authentication requirements **(A)** and **(B)**, we found a flaw in BrowserID that does not apply to its secondary mode and leads to an identity injection attack, and hence, violates property **(B)**. We confirmed the attack on the actual BrowserID implementation and reported it to Mozilla, who acknowledged it. We proposed a fix and formally proved that the fixed

system fulfills both **(A)** and **(B)**. Among the so far very few efforts on formally analyzing web applications and standards in expressive web models, our analysis constitutes the most complex formal analysis of a web application to date. It illustrates that (manual) security analysis of complex real-world web applications in a detailed web model, while laborious, is feasible and yields meaningful and practically relevant results.

During an attempt to formally analyze the privacy promise of the BrowserID system, we again found practical attacks. These attacks have been reported to and confirmed by Mozilla and, unfortunately, show that BrowserID would have to undergo a substantial redesign in order to fulfill its privacy promise. Interestingly, for our attacks we use a side channel that exploits information about the structure of windows in a browser. To the best of our knowledge, such side channel attacks have not gained much attention so far in the literature.

Finally, we have identified and proven important security properties of general application independent web features in order to facilitate future analysis efforts of web standards and web applications in the web model.

References

- [1] M. Abadi and C. Fournet. Mobile Values, New Names, and Secure Communication. In *POPL 2001*, pages 104–115. ACM Press, 2001.
- [2] D. Akhawe, A. Barth, P. E. Lam, J. Mitchell, and D. Song. Towards a Formal Foundation of Web Security. In *CSF 2010*, pp. 290–304. IEEE Computer Society, 2010.
- [3] A. Armando, R. Carbone, L. Compagna, J. Cuéllar, and M. L. Tobarra. Formal Analysis of SAML 2.0 Web Browser Single Sign-on: Breaking the SAML-based Single Sign-on for Google Apps. In *FMSE 2008*, pp. 1–10. ACM, 2008.
- [4] G. Bai, J. Lei, G. Meng, S. S. Venkatraman, P. Saxena, J. Sun, Y. Liu, and J. S. Dong. AUTHSCAN: Automatic Extraction of Web Authentication Protocols from Implementations. In *NDSS'13*. The Internet Society, 2013.
- [5] W. Bamberg et al. Persona FAQ. Mozilla Developer Network Wiki. Last edited Sept. 29, 2013. <https://developer.mozilla.org/en-US/Persona/FAQ>.
- [6] C. Bansal, K. Bhargavan, A. Delignat-Lavaud, and S. Maffei. Keys to the Cloud: Formal Analysis and Concrete Attacks on Encrypted Web Storage. In *POST 2013*, vol. 7796 of *LNCIS*, pp. 126–146. Springer, 2013.
- [7] C. Bansal, K. Bhargavan, and S. Maffei. Discovering Concrete Attacks on Website Authorization by Formal Analysis. In *CSF 2012*, pp. 247–262. IEEE Computer Society, 2012.
- [8] B. Blanchet. An Efficient Cryptographic Protocol Verifier Based on Prolog Rules. In *CSFW-14*, pp. 82–96. IEEE Computer Society, 2001.
- [9] Bugzilla@Mozilla. Bug 1064254 - Identity Injection Attack on Persona by Malicious IdP, September 2014. https://bugzilla.mozilla.org/show_bug.cgi?id=1064254 (access restricted).
- [10] Bugzilla@Mozilla. Bug 1120255 - Privacy leak in Persona, January 2015. https://bugzilla.mozilla.org/show_bug.cgi?id=1120255 (access restricted).
- [11] S. Chari, C. S. Jutla, and A. Roy. Universally Composable Security Analysis of OAuth v2.0. *IACR Cryptology ePrint Archive*, 2011:526, 2011.
- [12] M. Dietz and D. S. Wallach. Hardening Persona – Improving Federated Web Login. In *NDSS 2014*. The Internet Society, 2014.
- [13] D. Fett, R. Küsters, and G. Schmitz. An Expressive Model for the Web Infrastructure: Definition and Application to the BrowserID SSO System. In *S&P 2014*, pp. 673–688. IEEE Computer Society, 2014.

- [14] D. Fett, R. Küsters, and G. Schmitz. An Expressive Model for the Web Infrastructure: Definition and Application to the BrowserID SSO System. Technical Report arXiv:1403.1866, arXiv, 2014. Available at <http://arxiv.org/abs/1403.1866>.
- [15] HTML5, W3C Recommendation. Oct. 28, 2014.
- [16] D. Jackson. Alloy: A New Technology for Software Modelling. In *TACAS 2002*, vol. 2280 of *LNCS*, p. 20. Springer, 2002.
- [17] F. Kerschbaum. Simple Cross-Site Attack Prevention. In *SecureComm 2007*, pp. 464–472. IEEE Computer Society, 2007.
- [18] A. Kumar. A Lightweight Formal Approach for Analyzing Security of Web Protocols. In *RAID 2014*, vol. 8688 of *LNCS*, pp. 192–211. Springer, 2014.
- [19] C. Mills. Introducing BrowserID: A better way to sign in. Identity at Mozilla. Jul. 14, 2011. <http://identity.mozilla.com/post/7616727542/>.
- [20] Mozilla Identity Team. BrowserID Source Code. BrowserID Repository. <https://github.com/mozilla/browserid>.
- [21] Mozilla Identity Team. Persona. <https://login.persona.org>.
- [22] Mozilla Identity Team. Persona. Mozilla Developer Network. Last visited October 15, 2014. <https://developer.mozilla.org/en/docs/persona>.
- [23] J. Somorovsky, A. Mayer, J. Schwenk, M. Kampmann, and M. Jensen. On Breaking SAML: Be Whoever You Want to Be. In *USENIX Security '12*, pp. 397–412. USENIX Association, 2012.
- [24] S.-T. Sun and K. Beznosov. The Devil is in the (Implementation) Details: An Empirical Analysis of OAuth SSO Systems. In *CCS'12*, pp. 378–390. ACM, 2012.
- [25] S.-T. Sun, K. Hawkey, and K. Beznosov. Systematically Breaking and Fixing OpenID Security: Formal Analysis, Semi-Automated Empirical Evaluation, and Practical Countermeasures. *Computers & Security*, 31(4):465–483, 2012.
- [26] R. Wang, S. Chen, and X. Wang. Signing Me onto Your Accounts through Facebook and Google: A Traffic-Guided Security Study of Commercially Deployed Single-Sign-On Web Services. In *S&P 2012*, pp. 365–379. IEEE Computer Society, 2012.
- [27] Web Storage - W3C Recommendation 30 July 2013. <http://www.w3.org/TR/webstorage/>.

A Communication Model

Extending Section 2.1, we here present details and definitions on the basic concepts of the communication model. For readability, some parts from Section 2.1 are repeated.

A.1 Terms, Messages and Events

The signature Σ for the terms and messages considered in this work is the union of the following pairwise disjoint sets of function symbols:

- constants $C = \text{IPs} \cup \mathbb{S} \cup \{\top, \perp, \diamond\}$ where the three sets are pairwise disjoint, \mathbb{S} is interpreted to be the set of ASCII strings (including the empty string ε), and IPs is interpreted to be a set of (IP) addresses,³

³For brevity of presentation, in Section 2.1 the set C contained also the set of nonces \mathcal{N} . Here nonces are considered separately (see Definition 1).

- function symbols for public keys, (a)symmetric encryption/decryption, and signatures: $\text{pub}(\cdot)$, $\text{enc}_a(\cdot, \cdot)$, $\text{dec}_a(\cdot, \cdot)$, $\text{enc}_s(\cdot, \cdot)$, $\text{dec}_s(\cdot, \cdot)$, $\text{sig}(\cdot, \cdot)$, $\text{checksig}(\cdot, \cdot)$, and $\text{extractmsg}(\cdot)$,
- n -ary sequences $\langle \cdot \rangle$, $\langle \cdot, \cdot \rangle$, $\langle \cdot, \cdot, \cdot \rangle$, etc., and
- projection symbols $\pi_i(\cdot)$ for all $i \in \mathbb{N}$.

Definition 1. Let $X = \{x_0, x_1, \dots\}$ be a set of variables and \mathcal{N} be an infinite set of constants (nonces) such that Σ , X , and \mathcal{N} are pairwise disjoint. For $N \subseteq \mathcal{N}$, we define the set $\mathcal{T}_N(X)$ of terms over $\Sigma \cup N \cup X$ inductively as usual: (1) If $t \in N \cup X$, then t is a term. (2) If $f \in \Sigma$ is an n -ary function symbol in Σ for some $n \geq 0$ and t_1, \dots, t_n are terms, then $f(t_1, \dots, t_n)$ is a term.

By $\mathcal{T}_N = \mathcal{T}_N(\emptyset)$, we denote the set of all terms over $\Sigma \cup N$ without variables, called *ground terms*. The set \mathcal{M} of messages (over \mathcal{N}) is defined to be the set of ground terms $\mathcal{T}_{\mathcal{N}}$.

Example 1. For example, $k \in \mathcal{N}$ and $\text{pub}(k)$ are messages, where k typically models a private key and $\text{pub}(k)$ the corresponding public key. For constants a, b, c and the nonce $k \in \mathcal{N}$, the message $\text{enc}_a(\langle a, b, c \rangle, \text{pub}(k))$ is interpreted to be the message $\langle a, b, c \rangle$ (the sequence of constants a, b, c) encrypted by the public key $\text{pub}(k)$.

For strings (elements in \mathbb{S}), we use a specific font. For example, `HTTPReq` and `HTTPRes` are strings. We denote by $\text{Doms} \subseteq \mathbb{S}$ the set of domains, e.g., `example.com` \in Doms . We denote by $\text{Methods} \subseteq \mathbb{S}$ the set of methods used in HTTP requests, e.g., `GET`, `POST` \in Methods .

The equational theory associated with the signature Σ is given in Figure 5.

$$\begin{aligned} \text{dec}_a(\text{enc}_a(x, \text{pub}(y)), y) &= x & (1) \\ \text{dec}_s(\text{enc}_s(x, y), y) &= x & (2) \\ \text{extractmsg}(\text{sig}(x, y)) &= x & (3) \\ \text{checksig}(\text{sig}(x, y), \text{pub}(y)) &= \top & (4) \\ \pi_i(\langle x_1, \dots, x_n \rangle) &= x_i \text{ if } 1 \leq i \leq n & (5) \\ \pi_j(\langle x_1, \dots, x_n \rangle) &= \diamond \text{ if } j \notin \{1, \dots, n\} & (6) \\ \pi_j(t) &= \diamond \text{ if } t \text{ is not a sequence} & (7) \end{aligned}$$

Fig. 5. Equational theory for Σ .

By \equiv we denote the congruence relation on $\mathcal{T}_{\mathcal{N}}(X)$ induced by this theory. For example, we have that $\pi_1(\text{dec}_a(\text{enc}_a(\langle \mathbf{a}, \mathbf{b} \rangle, \text{pub}(k)), k)) \equiv \mathbf{a}$.

Definition 2. An event (over IPs and \mathcal{M}) is of the form $(a:f:m)$, for $a, f \in \text{IPs}$ and $m \in \mathcal{M}$, where a is interpreted to be the receiver address and f is the sender address. We denote by \mathcal{E} the set of all events.

A.2 Atomic Processes, Systems and Runs

We here define atomic processes, systems, and runs of systems.

An atomic process takes its current state and an event as input, and then (non-deterministically) outputs a new state and a set of events.

Definition 3. A (generic) atomic process is a tuple $p = (I^p, Z^p, R^p, s_0^p)$ where $I^p \subseteq \text{IPs}$, Z^p is a set of states, $R^p \subseteq (\mathcal{E} \times Z^p) \times (2^{\mathcal{E}} \times Z^p)$, and $s_0^p \in Z^p$ is the initial state of p . We write $(e, z)R(E, z')$ instead of $((e, z), (E, z')) \in R$.

A system \mathcal{P} is a (possibly infinite) set of atomic processes.

Definition 4. A configuration of a system \mathcal{P} is a tuple (S, E) where S maps every atomic process $p \in \mathcal{P}$ to its current state $S(p) \in Z^p$ and E is a (possibly infinite) multi-set of events waiting to be delivered.

Definition 5. A processing step of the system \mathcal{P} is of the form

$$(S, E) \xrightarrow[p \rightarrow E_{out}]{e \rightarrow p} (S', E')$$

such that (1) there exists an event $e = (a:f:m) \in E$, $E_{out} \subseteq E'$, and a process $p \in \mathcal{P}$ with $(e, S(p))R^p(E_{out}, S'(p))$ and $a \in I^p$, (2) $S'(p') = S(p')$ for all $p' \neq p$, and (3) $E' = (E \setminus \{e\}) \cup E_{out}$ (multi-set operations). We may omit the superscript and/or subscript of the arrow.

Definition 6. Let \mathcal{P} be a system and E_0 be a multi-set of events. A run ρ of a system \mathcal{P} initiated by E_0 is a finite sequence of configurations $(S_0, E_0), \dots, (S_n, E_n)$ or an infinite sequence of configurations $(S_0, E_0), \dots$ such that $S_0(p) = s_0^p$ for all $p \in \mathcal{P}$ and $(S_i, E_i) \rightarrow (S_{i+1}, E_{i+1})$ for all $0 \leq i < n$ (finite run) or for all $i \geq 0$ (infinite run).

A.3 Atomic Dolev-Yao Processes

We next define atomic Dolev-Yao processes, for which we require that the messages and states that they output can be computed (more formally, derived) from the current input event and state. For this purpose, we first define what it means to derive a message from given messages.

Definition 7. Let $N \subseteq \mathcal{N}$, $\tau \in \mathcal{T}_N(\{x_1, \dots, x_n\})$, and $t_1, \dots, t_n \in \mathcal{T}_N$. By $\tau[t_1/x_1, \dots, t_n/x_n]$ we denote the (ground) term obtained from τ by replacing all occurrences of x_i in τ by t_i , for all $i \in \{1, \dots, n\}$.

Definition 8. Let $M \subseteq \mathcal{M}$ be a set of messages. We say that a message m can be derived from M with nonces N if there exist $n \geq 0$, $m_1, \dots, m_n \in M$, and $\tau \in \mathcal{T}_N(\{x_1, \dots, x_n\})$ such that $m \equiv \tau[m_1/x_1, \dots, m_n/x_n]$. We denote by $d_N(M)$ the set of all messages that can be derived from M with nonces N .

For example, $a \in d_{\{k\}}(\{\text{enc}_a(\langle a, b, c \rangle, \text{pub}(k))\})$.

Definition 9. An atomic Dolev-Yao process (or simply, a DY process) is a tuple $p = (I^p, Z^p, R^p, s_0^p, N^p)$ such that (I^p, Z^p, R^p, s_0^p) is an atomic process and (1) $N^p \subseteq \mathcal{N}$ is an (initial) set of nonces, (2) $Z^p \subseteq \mathcal{T}_{\mathcal{N}}$ (and hence, $s_0^p \in \mathcal{T}_{\mathcal{N}}$), and (3) for all $a, a', f, f' \in \text{IPs}$, $m, m', s, s' \in \mathcal{T}_{\mathcal{N}}$, set of events E with $((a:f:m), s)R(E, s')$ and $(a':f':m') \in E$ it holds true that $m', s' \in d_N(\{m, s\})$. (Note that $a', f' \in d_N(\{m, s\})$.)

Definition 10. An (atomic) attacker process for a set of sender addresses $A \subseteq \text{IPs}$ is an atomic DY process $p = (I, Z, R, s_0, N)$ such that for all $a, f \in \text{IPs}$, $m \in \mathcal{T}_{\mathcal{N}}$, and $s \in Z$ we have that $((a:f:m), s)R(E, s')$ iff $s' = \langle \langle a, f, m \rangle, s \rangle$ and $E = \{(a':f':m') \mid a' \in \text{IPs}, f' \in A, m' \in d_N(\{m, s\})\}$.

A.4 Scripting Processes

We define scripting processes, which model client-side scripting technologies, such as JavaScript. Scripting processes are defined similarly to DY processes.

Definition 11. A scripting process (or simply, a script) is a relation $R \subseteq (\mathcal{T}_{\mathcal{N}} \times 2^{\mathcal{N}}) \times \mathcal{T}_{\mathcal{N}}$ such that for all $s, s' \in \mathcal{T}_{\mathcal{N}}$ and $N \subseteq \mathcal{N}$ with $(s, N)R s'$ it follows that $s' \in d_N(s)$.

A script is called by the browser which provides it with a (fresh, infinite) set N of nonces and state information s . The script then outputs a term s' , which represents the new internal state and some command which is interpreted by the browser.

Similarly to an attacker process, we define the *attacker script* R^{att} . This script outputs everything that is derivable from the input, i.e., $R^{\text{att}} = \{((s, N), s') \mid s \in \mathcal{T}_{\mathcal{N}}, N \subseteq \mathcal{N}, s' \in d_N(s)\}$.

B Message and Data Formats

We now provide some more details about data and message formats that are needed for the formal treatment of the web model and the analysis of BrowserID presented in the rest of the appendix.

B.1 Notations

Definition 12 (Sequence Notations). For a sequence $t = \langle t_1, \dots, t_n \rangle$ and a set s we use $t \subset^{\diamond} s$ to say that $t_1, \dots, t_n \in s$. We define $x \in^{\diamond} t \iff \exists i : t_i = x$. We write $t +^{\diamond} y$ to denote the sequence $\langle t_1, \dots, t_n, y \rangle$. For a sequence $t = \langle t_1, \dots, t_n \rangle$ we define $|t| = n$. If t is not a sequence, we set $|t| = \diamond$. For a finite set M with $M = \{m_1, \dots, m_n\}$ we use $\langle M \rangle$ to denote the term of the form $\langle m_1, \dots, m_n \rangle$. (The order of the elements does not matter; one is chosen arbitrarily.)

Definition 13. A dictionary over X and Y is a term of the form

$$\langle \langle k_1, v_1 \rangle, \dots, \langle k_n, v_n \rangle \rangle$$

where $k_1, \dots, k_n \in X$, $v_1, \dots, v_n \in Y$, and the keys k_1, \dots, k_n are unique, i.e., $\forall i \neq j : k_i \neq k_j$. We call every term $\langle k_i, v_i \rangle$, $i \in \{1, \dots, n\}$, an element of the dictionary with key k_i and value v_i . We often write $[k_1 : v_1, \dots, k_i : v_i, \dots, k_n : v_n]$ instead of $\langle \langle k_1, v_1 \rangle, \dots, \langle k_n, v_n \rangle \rangle$. We denote the set of all dictionaries over X and Y by $[X \times Y]$.

We note that the empty dictionary is equivalent to the empty sequence, i.e., $[] = \langle \rangle$. Figure 6 shows the short notation for dictionary operations that will be used when describing the browser atomic process. For a dictionary $z = [k_1 : v_1, k_2 : v_2, \dots, k_n : v_n]$ we write $k \in z$ to say that there exists i such that $k = k_i$. We write $z[k_j] := v_j$ to extract elements. If $k \notin z$, we set $z[k] := \langle \rangle$.

$$[k_1 : v_1, \dots, k_i : v_i, \dots, k_n : v_n][k_i] = v_i \quad (8)$$

$$\begin{aligned} [k_1 : v_1, \dots, k_{i-1} : v_{i-1}, k_i : v_i, k_{i+1} : v_{i+1}, \dots, k_n : v_n] - k_i = \\ [k_1 : v_1, \dots, k_{i-1} : v_{i-1}, k_{i+1} : v_{i+1}, \dots, k_n : v_n] \end{aligned} \quad (9)$$

Fig. 6. Dictionary operators with $1 \leq i \leq n$.

Given a term $t = \langle t_1, \dots, t_n \rangle$, we can refer to any subterm using a sequence of integers. The subterm is determined by repeated application of the projection π_i for the integers i in the sequence. We call such a sequence a *pointer*:

Definition 14. A pointer is a sequence of non-negative integers. We write $\tau.\bar{p}$ for the application of the pointer \bar{p} to the term τ . This operator is applied from left to right. For pointers consisting of a single integer, we may omit the sequence braces for brevity.

Example 2. For the term $\tau = \langle a, b, \langle c, d, \langle e, f \rangle \rangle \rangle$ and the pointer $\bar{p} = \langle 3, 1 \rangle$, the subterm of τ at the position \bar{p} is $c = \pi_1(\pi_3(\tau))$. Also, $\tau.3.\langle 3, 1 \rangle = \tau.3.\bar{p} = \tau.3.3.1 = e$.

To improve readability, we try to avoid writing, e.g., $o.2$ or $\pi_2(o)$ in this document. Instead, we will use the names of the components of a sequence that is of a defined form as pointers that point to the corresponding subterms. E.g., if an *Origin* term is defined as $\langle \text{host}, \text{protocol} \rangle$ and o is an Origin term, then we can write $o.\text{protocol}$ instead of $\pi_2(o)$ or $o.2$. See also Example 3.

In the pseudocode, we will write, for example,

let x, y **such that** $\langle \text{Constant}, x, y \rangle \equiv t$ **if possible; otherwise** doSomethingElse

for some variables x, y , a string Constant , and some term t to express that $x := \pi_2(t)$, and $y := \pi_3(t)$ if $\text{Constant} \equiv \pi_1(t)$ and if $|\langle \text{Constant}, x, y \rangle| = |t|$, and that otherwise x and y are not set and doSomethingElse is executed.

B.2 URLs

Definition 15. A URL is a term of the form $\langle \text{URL}, \text{protocol}, \text{host}, \text{path}, \text{params} \rangle$ with $\text{protocol} \in \{\text{P}, \text{S}\}$ (for *plain* (HTTP) and *secure* (HTTPS)), $\text{host} \in \text{Doms}$, $\text{path} \in \mathbb{S}$ and $\text{params} \in [\mathbb{S} \times \mathcal{T}_{\mathcal{N}}]$. The set of all valid URLs is URLs.

Example 3. For the URL $u = \langle \text{URL}, a, b, c, d \rangle$, $u.\text{protocol} = a$. If, in the algorithm described later, we say $u.\text{path} := e$ then $u = \langle \text{URL}, a, b, c, e \rangle$ afterwards.

B.3 Origins

Definition 16. An origin is a term of the form $\langle \text{host}, \text{protocol} \rangle$ with $\text{host} \in \text{Doms}$ and $\text{protocol} \in \{\text{P}, \text{S}\}$. We write *Origins* for the set of all origins. See Example 6 for an example of an origin.

B.4 Cookies

Definition 17. A cookie is a term of the form $\langle \text{name}, \text{content} \rangle$ where $\text{name} \in \mathcal{T}_{\mathcal{N}}$, and content is a term of the form $\langle \text{value}, \text{secure}, \text{session}, \text{httpOnly} \rangle$ where $\text{value} \in \mathcal{T}_{\mathcal{N}}$, $\text{secure}, \text{session}, \text{httpOnly} \in \{\top, \perp\}$. We write *Cookies* for the set of all cookies.

If the *secure* attribute of a cookie is set, the browser will not transfer this cookie over unencrypted HTTP connections. If the *session* flag is set, this cookie will be deleted as soon as the browser is closed. The *httpOnly* attribute controls whether JavaScript has access to this cookie.

Note that cookies of the form described here are only contained in HTTP(S) requests. In responses, only the components *name* and *value* are transferred as a pairing of the form $\langle \text{name}, \text{value} \rangle$.

B.5 HTTP Messages

Definition 18. An HTTP request is a term of the form shown in (10). An HTTP response is a term of the form shown in (11).

$$\langle \text{HTTPReq}, \text{nonce}, \text{method}, \text{host}, \text{path}, \text{parameters}, \text{headers}, \text{body} \rangle \quad (10)$$

$$\langle \text{HTTPResp}, \text{nonce}, \text{status}, \text{headers}, \text{body} \rangle \quad (11)$$

The components are defined as follows:

- $\text{nonce} \in \mathcal{N}$ serves to map each response to the corresponding request
- $\text{method} \in \text{Methods}$ is one of the HTTP methods.
- $\text{host} \in \text{Doms}$ is the host name in the *HOST* header of HTTP/1.1.
- $\text{path} \in \mathbb{S}$ is a string indicating the requested resource at the server side
- $\text{status} \in \mathbb{S}$ is the HTTP status code (i.e., a number between 100 and 505, as defined by the HTTP standard)
- $\text{parameters} \in [\mathbb{S} \times \mathcal{T}_{\mathcal{N}}]$ contains URL parameters
- $\text{headers} \in [\mathbb{S} \times \mathcal{T}_{\mathcal{N}}]$, containing request/response headers. The dictionary elements are terms of one of the following forms:
 - $\langle \text{Origin}, o \rangle$ where o is an origin
 - $\langle \text{Set-Cookie}, c \rangle$ where c is a sequence of cookies
 - $\langle \text{Cookie}, c \rangle$ where $c \in [\mathbb{S} \times \mathcal{T}_{\mathcal{N}}]$ (note that in this header, only names and values of cookies are transferred)
 - $\langle \text{Location}, l \rangle$ where $l \in \text{URLs}$
 - $\langle \text{Strict-Transport-Security}, \top \rangle$
- $\text{body} \in \mathcal{T}_{\mathcal{N}}$ in requests and responses.

We write HTTPRequests/HTTPResponses for the set of all HTTP requests or responses, respectively.

Example 4 (HTTP Request and Response).

$$r := \langle \text{HTTPReq}, n_1, \text{POST}, \text{example.com}, / \text{show}, \langle \langle \text{index}, 1 \rangle \rangle, \\ [\text{Origin} : \langle \text{example.com}, S \rangle], \langle \text{foo}, \text{bar} \rangle \rangle \quad (12)$$

$$s := \langle \text{HTTPResp}, n_1, 200, \langle \langle \text{Set-Cookie}, \langle \langle \text{SID}, \langle n_2, \perp, \perp, \top \rangle \rangle \rangle \rangle, \langle \text{somecript}, x \rangle \rangle \quad (13)$$

An HTTP GET request for the URL <http://example.com/show?index=1> is shown in (12), with an Origin header and a body that contains $\langle \text{foo}, \text{bar} \rangle$. A possible response is shown in (13), which contains an httpOnly cookie with name SID and value n_2 as well as the string representation somecript of the scripting process $\text{script}^{-1}(\text{somecript})$ (which should be an element of S) and its initial state x .

Encrypted HTTP Messages. For HTTPS, requests are encrypted using the public key of the server. Such a request contains an (ephemeral) symmetric key chosen by the client that issued the request. The server is supported to encrypt the response using the symmetric key.

Definition 19. An encrypted HTTP request is of the form $\text{enc}_a(\langle m, k' \rangle, k)$, where $k, k' \in \mathcal{K}$ and $m \in \text{HTTPRequests}$. The corresponding encrypted HTTP response would be of the form $\text{enc}_s(m', k')$, where $m' \in \text{HTTPResponses}$. We call the sets of all encrypted HTTP requests and responses HTTPSRequests or HTTPSResponses, respectively.

Example 5.

$$\text{enc}_a(\langle r, k' \rangle, \text{pub}(k_{\text{example.com}})) \quad (14)$$

$$\text{enc}_s(s, k') \quad (15)$$

The term (14) shows an encrypted request (with r as in (12)). It is encrypted using the public key $\text{pub}(k_{\text{example.com}})$. The term (15) is a response (with s as in (13)). It is encrypted symmetrically using the (symmetric) key k' that was sent in the request (14).

B.6 DNS Messages

Definition 20. A DNS request is a term of the form $\langle \text{DNSResolve}, \text{domain}, n \rangle$ where $\text{domain} \in \text{Doms}$, $n \in \mathcal{K}$. We call the set of all DNS requests DNSRequests.

Definition 21. A DNS response is a term of the form $\langle \text{DNSResolved}, \text{result}, n \rangle$ with $\text{result} \in \text{IPs}$, $n \in \mathcal{K}$. We call the set of all DNS responses DNSResponses.

DNS servers are supposed to include the nonce they received in a DNS request in the DNS response that they send back so that the party which issued the request can match it with the request.

C Detailed Description of the Browser Model

Following the informal description of the browser model in Section 2.5, we now present a formal model. We start by introducing some notation and terminology.

C.1 Notation and Terminology (Web Browser State)

Before we can define the state of a web browser, we first have to define windows and documents. Concrete window and document terms are shown in Example 6.

Definition 22. A window is a term of the form $w = \langle \textit{nonce}, \textit{documents}, \textit{opener} \rangle$ with $\textit{nonce} \in \mathcal{N}$, $\textit{documents} \subset^{\diamond} \text{Documents}$ (defined below), $\textit{opener} \in \mathcal{N} \cup \{\perp\}$ where $d.\textit{active} = \top$ for exactly one $d \in^{\diamond} \textit{documents}$ if $\textit{documents}$ is not empty (we then call d the active document of w). We write Windows for the set of all windows. We write $w.\textit{activedocument}$ to denote the active document inside window w if it exists and $\langle \rangle$ else.

We will refer to the window \textit{nonce} as (*window*) *reference*.

The documents contained in a window term to the left of the active document are the previously viewed documents (available to the user via the “back” button) and the documents in the window term to the right of the currently active document are documents available via the “forward” button, as will be clear from the description of web browser model (see Section C.2).

A window a may have opened a top-level window b (i.e., a window term which is not a subterm of a document term). In this case, the *opener* part of the term b is the nonce of a , i.e., $b.\textit{opener} = a.\textit{nonce}$.

Definition 23. A document d is a term of the form

$$\langle \textit{nonce}, \textit{origin}, \textit{script}, \textit{scriptstate}, \textit{scriptinput}, \textit{subwindows}, \textit{active} \rangle$$

where $\textit{nonce} \in \mathcal{N}$, $\textit{origin} \in \text{Origins}$, $\textit{script} \in \mathcal{T}_{\mathcal{N}}$, $\textit{scriptstate} \in \mathcal{T}_{\mathcal{N}}$, $\textit{scriptinput} \in \mathcal{T}_{\mathcal{N}}$, $\textit{subwindows} \subset^{\diamond} \text{Windows}$, $\textit{active} \in \{\top, \perp\}$. A limited document is a term of the form $\langle \textit{nonce}, \textit{subwindows} \rangle$ with \textit{nonce} , $\textit{subwindows}$ as above. A window $w \in^{\diamond} \textit{subwindows}$ is called a subwindow (of d). We write Documents for the set of all documents.

We will refer to the document \textit{nonce} as (*document*) *reference*.

Example 6. The following is an example of a window term with reference n_1 , two documents, and an opener (n_4):

$$\langle n_1, \langle \langle n_2, \langle \text{example.com}, \text{P} \rangle, \text{script1}, \langle \rangle, \langle \rangle, \langle \rangle, \perp \rangle, \langle n_3, \langle \text{example.com}, \text{S} \rangle, \text{script2}, \langle \rangle, \langle \rangle, \langle \rangle, \top \rangle \rangle, n_4 \rangle$$

The first document has the reference n_2 . It was loaded from the origin $\langle \text{example.com}, \text{P} \rangle$, which translates into <http://example.com>. Its scripting process has the string representation `script1`, the last state and the input history of this process are empty. The document does not have subwindows and is inactive (\perp). The second document has the reference n_3 , its origin corresponds to <https://example.com>, the scripting process is represented by `script2`, and the document is active (\top). All other components are empty.

We can now define the set of states of web browsers. Note that we use the dictionary notation that we introduced in Definition 13.

Definition 24. Let $OR := \{\langle o, r \rangle \mid o \in \text{Origins}, r \in \mathcal{N}\}$. The set of states Z^p of a web browser atomic process p consists of the terms of the form

$$\langle \text{windows}, \text{ids}, \text{secrets}, \text{cookies}, \text{localStorage}, \text{sessionStorage}, \text{keyMapping}, \\ \text{sts}, \text{DNSaddress}, \text{nonces}, \text{pendingDNS}, \text{pendingRequests}, \text{isCorrupted} \rangle$$

where

- $\text{windows} \subset^{\diamond} \text{Windows}$,
- $\text{ids} \subset^{\diamond} \mathcal{T}_{\mathcal{N}}$,
- $\text{secrets} \in [\text{Origins} \times \mathcal{N}]$,
- cookies is a dictionary over Doms and dictionaries of Cookies,
- $\text{localStorage} \in [\text{Origins} \times \mathcal{T}_{\mathcal{N}}]$,
- $\text{sessionStorage} \in [OR \times \mathcal{T}_{\mathcal{N}}]$,
- $\text{keyMapping} \in [\text{Doms} \times \mathcal{T}_{\mathcal{N}}]$,
- $\text{sts} \subset^{\diamond} \text{Doms}$,
- $\text{DNSaddress} \in \text{IPs}$,
- $\text{nonces} \subset^{\diamond} \mathcal{N}$,
- $\text{pendingDNS} \in [\mathcal{N} \times \mathcal{T}_{\mathcal{N}}]$,
- $\text{pendingRequests} \in \mathcal{T}_{\mathcal{N}}$,
- and $\text{isCorrupted} \in \{\perp, \text{FULLCORRUPT}, \text{CLOSECORRUPT}\}$.

Definition 25. For two window terms w and w' we write $w \xrightarrow{\text{childof}} w'$ if

$$w \in^{\diamond} w'.\text{activedocument.subwindows}.$$

We write $\xrightarrow{\text{childof}^+}$ for the transitive closure.

In the following description of the web browser relation R^p we will use the helper functions Subwindows, Docs, Clean, CookieMerge and AddCookie.

Given a browser state s , Subwindows(s) denotes the set of all pointers⁴ to windows in the window list $s.\text{windows}$, their active documents, and (recursively) the subwindows of these documents. We exclude subwindows of inactive documents and their subwindows. With Docs(s) we denote the set of pointers to all active documents in the set of windows referenced by Subwindows(s).

Definition 26. For a browser state s we denote by Subwindows(s) the minimal set of pointers that satisfies the following conditions: (1) For all windows $w \in^{\diamond} s.\text{windows}$ there is a $\bar{p} \in \text{Subwindows}(s)$ such that $s.\bar{p} = w$. (2) For all $\bar{p} \in \text{Subwindows}(s)$, the active document d of the window $s.\bar{p}$ and every subwindow w of d there is a pointer $\bar{p}' \in \text{Subwindows}(s)$ such that $s.\bar{p}' = w$.

Given a browser state s , the set Docs(s) of pointers to active documents is the minimal set such that for every $\bar{p} \in \text{Subwindows}(s)$, there is a pointer $\bar{p}' \in \text{Docs}(s)$ with $s.\bar{p}' = s.\bar{p}.\text{activedocument}$.

⁴Recall the definition of a pointer in Definition 14.

The function `Clean` will be used to determine which information about windows and documents the script running in the document d has access to.

Definition 27. Let s be a browser state and d a document. By `Clean`(s, d) we denote the term that equals $s.\text{windows}$ but with all inactive documents removed (including their subwindows etc.) and all subterms that represent non-same-origin documents w.r.t. d replaced by a limited document d' with the same nonce and the same subwindow list. Note that non-same-origin documents on all levels are replaced by their corresponding limited document.

The function `CookieMerge` merges two sequences of cookies together: When used in the browser, oldcookies is the sequence of existing cookies for some origin, newcookies is a sequence of new cookies that was output by some script. The sequences are merged into a set of cookies using an algorithm that is based on the *Storage Mechanism* algorithm described in RFC6265.

Definition 28. For a sequence of cookies (with pairwise different names) oldcookies and a sequence of cookies newcookies , the set `CookieMerge`($\text{oldcookies}, \text{newcookies}$) is defined by the following algorithm: From newcookies remove all cookies c that have $c.\text{content.httpOnly} \equiv \top$. For any $c, c' \in \langle \rangle \text{newcookies}$, $c.\text{name} \equiv c'.\text{name}$, remove the cookie that appears left of the other in newcookies . Let m be the set of cookies that have a name that either appears in oldcookies or in newcookies , but not in both. For all pairs of cookies ($c_{\text{old}}, c_{\text{new}}$) with $c_{\text{old}} \in \langle \rangle \text{oldcookies}$, $c_{\text{new}} \in \langle \rangle \text{newcookies}$, $c_{\text{old}}.\text{name} \equiv c_{\text{new}}.\text{name}$, add c_{new} to m if $c_{\text{old}}.\text{content.httpOnly} \equiv \perp$ and add c_{old} to m otherwise. The result of `CookieMerge`($\text{oldcookies}, \text{newcookies}$) is m .

The function `AddCookie` adds a cookie c received in an HTTP response to the sequence of cookies contained in the sequence oldcookies . It is again based on the algorithm described in RFC6265 but simplified for the use in the browser model.

Definition 29. For a sequence of cookies (with different names) oldcookies and a cookie c , the sequence `AddCookie`($\text{oldcookies}, c$) is defined by the following algorithm: Let $m := \text{oldcookies}$. Remove any c' from m that has $c.\text{name} \equiv c'.\text{name}$. Append c to m and return m .

The function `NavigableWindows` returns a set of windows that a document is allowed to navigate. We closely follow [15], Section 5.1.4 for this definition.

Definition 30. The set `NavigableWindows`(\bar{w}, s') is the set $\bar{W} \subseteq \text{Subwindows}(s')$ of pointers to windows that the active document in \bar{w} is allowed to navigate. The set \bar{W} is defined to be the minimal set such that for every $\bar{w}' \in \text{Subwindows}(s')$ the following is true:

- If $s'.\bar{w}'.\text{activedocument.origin} \equiv s'.\bar{w}.\text{activedocument.origin}$ (i.e., the active documents in \bar{w} and \bar{w}' are same-origin), then $\bar{w}' \in \bar{W}$, and
- If $s'.\bar{w} \xrightarrow{\text{childof}^*} s'.\bar{w}' \wedge \nexists \bar{w}'' \in \text{Subwindows}(s') \text{ with } s'.\bar{w}' \xrightarrow{\text{childof}^*} s'.\bar{w}''$ (\bar{w}' is a top-level window and \bar{w} is an ancestor window of \bar{w}'), then $\bar{w}' \in \bar{W}$, and

- If $\exists \bar{p} \in \text{Subwindows}(s')$ such that $s'.\bar{w}' \xrightarrow{\text{childof}^+} s'.\bar{p}$
 $\wedge s'.\bar{p}.\text{activedocument}.\text{origin} = s'.\bar{w}.\text{activedocument}.\text{origin}$ (\bar{w}' is not a top-level window but there is an ancestor window \bar{p} of \bar{w}' with an active document that has the same origin as the active document in \bar{w}), then $\bar{w}' \in \bar{W}$, and
- If $\exists \bar{p} \in \text{Subwindows}(s')$ such that $s'.\bar{w}'.\text{opener} = s'.\bar{p}.\text{nonce} \wedge \bar{p} \in \bar{W}$ (\bar{w}' is a top-level window—it has an opener—and \bar{w} is allowed to navigate the opener window of \bar{w}' , \bar{p}), then $\bar{w}' \in \bar{W}$.

C.2 Description of the Web Browser Atomic Process

We will now describe the relation R^p of a standard HTTP browser p . For a tuple $r = ((a:f:m), s), (M, s')$ we define r to belong to R^p iff the non-deterministic algorithm presented in Section 7, when given $((a:f:m), s)$ as input, terminates with **stop** M, s' , i.e., with output M and s' . Recall that $(a:f:m)$ is an (input) event and s is a (browser) state, M is a set of (output) events, and s' is a new (browser) state.

The notation **let** $n \leftarrow N$ is used to describe that n is chosen non-deterministically from the set N . We write **for each** $s \in M$ **do** to denote that the following commands (until **end for**) are repeated for every element in M , where the variable s is the current element. The order in which the elements are processed is chosen non-deterministically.

We first define some functions which will be used in the main algorithm presented in Section 7.

Functions. In the description of the following functions we use a, f, m, s and N^p as read-only global input variables. Also, the functions use the set N^p as a read-only set. All other variables are local variables or arguments.

TAKENONCE returns a nonce from the set of unused nonces and modifies the browser state such that the nonce is added to the sequence of used nonces. Note that this function returns two values, the nonce n and the modified state s' .

Algorithm 1 Non-deterministically choose a fresh nonce.

```

1: function TAKENONCE( $s'$ )
2:   let  $n \leftarrow \{x \mid x \in N^p \wedge x \notin s'.\text{nonces}\}$ 
3:   let  $s'.\text{nonces} := s'.\text{nonces} + n$ 
4:   return  $n, s'$ 
5: end function

```

The following function, GETNAVIGABLEWINDOW, is called by the browser to determine the window that is *actually* navigated when a script in the window $s'.\bar{w}$ provides a window reference for navigation (e.g., for opening a link). When it is given a window reference (nonce) $window$, GETNAVIGABLEWINDOW returns a pointer to a selected window term in s' :

- If $window$ is the string `_BLANK`, a new window is created and a pointer to that window is returned.
- If $window$ is a nonce (reference) and there is a window term with a reference of that value in the windows in s' , a pointer \bar{w}' to that window term is returned, as

long as the window is navigable by the current window's document (as defined by NavigableWindows above).

In all other cases, \bar{w} is returned instead (the script navigates its own window).

Algorithm 2 Determine window for navigation.

```

1: function GETNAVIGABLEWINDOW( $\bar{w}$ , window,  $s'$ )
2:   if window  $\equiv$  _BLANK then                                 $\triangleright$  Open a new window when _BLANK is used
3:     let  $n, s' :=$  TAKENONCE( $s'$ )
4:     let  $w' := \langle n, \langle \rangle, s'.\bar{w}.nonce \rangle$ 
5:     let  $s'.windows := s'.windows + \langle \rangle w'$ 
         $\hookrightarrow$  and let  $\bar{w}'$  be a pointer to this new element in  $s'$ 
6:     return ( $\bar{w}', s'$ )
7:   end if
8:   let  $\bar{w}' \leftarrow$  NavigableWindows( $\bar{w}, s'$ ) such that  $s'.\bar{w}'.nonce \equiv window$ 
         $\hookrightarrow$  if possible; otherwise return ( $\bar{w}, s'$ )
9:   return ( $\bar{w}', s'$ )
10: end function

```

The following function takes a window reference as input and returns a pointer to a window as above, but it checks only that the active documents in both windows are same-origin. It creates no new windows.

Algorithm 3 Determine same-origin window.

```

1: function GETWINDOW( $\bar{w}$ , window,  $s'$ )
2:   let  $\bar{w}' \leftarrow$  Subwindows( $s'$ ) such that  $s'.\bar{w}'.nonce \equiv window$ 
         $\hookrightarrow$  if possible; otherwise return ( $\bar{w}, s'$ )
3:   if  $s'.\bar{w}'.activedocument.origin \equiv s'.\bar{w}.activedocument.origin$  then
4:     return ( $\bar{w}', s'$ )
5:   end if
6:   return ( $\bar{w}, s'$ )
7: end function

```

The next function is used to stop any pending requests for a specific window. From the pending requests and pending DNS requests it removes any requests with the given window reference n .

Algorithm 4 Cancel pending requests for given window.

```

1: function CANCELNAV( $n, s'$ )
2:   remove all  $\langle n, req, key, f \rangle$  from  $s'.pendingRequests$  for any  $req, key, f$ 
3:   remove all  $\langle x, \langle n, message, protocol \rangle \rangle$  from  $s'.pendingDNS$ 
         $\hookrightarrow$  for any  $x, message, protocol$ 
4:   return  $s'$ 
5: end function

```

The following function takes an HTTP request *message* as input, adds cookie and origin headers to the message, creates a DNS request for the hostname given in the request and stores the request in $s'.pendingDNS$ until the DNS resolution finishes. For normal HTTP requests, *reference* is a window reference. For XHRs, *reference* is a value of the form $\langle document, nonce \rangle$ where *document* is a document reference and *nonce* is

some nonce that was chosen by the script that initiated the request. *protocol* is either P or S. *origin* is the origin header value that is to be added to the HTTP request.

Algorithm 5 Prepare headers, do DNS resolution, save message.

```

1: function SEND(reference, message, protocol, origin, s')
2:   if message.host  $\in \langle \rangle s'.sts$  then
3:     let protocol := S
4:   end if
5:   let cookies :=  $\langle \{ \langle c.name, c.content.value \rangle \mid c \in \langle \rangle s'.cookies[message.host] \} \rangle$ 
    $\hookrightarrow \wedge (c.content.secure \implies (protocol = S))$ 
6:   let message.headers[Cookie] := cookies
7:   if origin  $\neq \perp$  then
8:     let message.headers[Origin] := origin
9:   end if
10:  let n, s' := TAKENONCE(s')
11:  let s'.pendingDNS[n] :=  $\langle reference, message, protocol \rangle$ 
12:  stop  $\{ \langle s'.DNSaddress : a : \langle DNSResolve, host, n \rangle \rangle, s' \}$ 
13: end function

```

The following two functions have informally been described in Section 2.5.

The function RUNSCRIPT performs a script execution step of the script in the document $s'.\bar{d}$ (which is part of the window $s'.\bar{w}$). A new script and document state is chosen according to the relation defined by the script and the new script and document state is saved. Afterwards, the *command* that the script issued is interpreted. Note that **for each** (Line 13) works in a non-deterministic order.

Algorithm 6 Execute a script.

```

1: function RUNSCRIPT( $\bar{w}, \bar{d}, s'$ )
2:   let n, s' := TAKENONCE(s')
3:   let tree := Clean(s', s'.\bar{d})
4:   let cookies :=  $\langle \{ \langle c.name, c.content.value \rangle \mid c \in \langle \rangle s'.cookies[s'.\bar{d}.origin.host] \} \rangle$ 
    $\hookrightarrow \wedge c.content.httpOnly = \perp$ 
    $\hookrightarrow \wedge (c.content.secure \implies (s'.\bar{d}.origin.protocol \equiv S))$ 
5:   let tlw  $\leftarrow s'.windows$  such that tlw is the top-level window containing  $\bar{d}$ 
6:   let sessionStorage := s'.sessionStorage [ $\langle s'.\bar{d}.origin, tlw.nonce \rangle$ ]
7:   let localStorage := s'.localStorage [s'.\bar{d}.origin]
8:   let secret := s'.secrets [s'.\bar{d}.origin]
9:   let nonces be an infinite subset of  $\{ x \mid x \in N^P \wedge x \notin \langle \rangle s'.nonces \}$ 
10:  let R  $\leftarrow \text{script}^{-1}(s'.\bar{d}.script)$ 
11:  let in :=  $\langle tree, s'.\bar{d}.nonce, s'.\bar{d}.scriptstate, s'.\bar{d}.scriptinput, cookies, \rangle$ 
    $\hookrightarrow \langle localStorage, sessionStorage, s'.ids, secret \rangle$ 
12:  let state'  $\leftarrow \mathcal{T}_{\mathcal{N}}$ ,
    $\hookrightarrow \langle cookies' \leftarrow Cookies, \rangle$ 
    $\hookrightarrow \langle localStorage' \leftarrow \mathcal{T}_{\mathcal{N}}, \rangle$ 
    $\hookrightarrow \langle command \leftarrow \mathcal{T}_{\mathcal{N}}, \rangle$ 
    $\hookrightarrow \langle out := \langle state', cookies', localStorage', sessionStorage', command \rangle \rangle$ 
    $\hookrightarrow \langle \text{such that } ((in, nonces), out) \in R. \rangle$ 
13:  for each n  $\in d_{\mathcal{N}}(\langle in, out \rangle) \cap N^P$  do

```

```

14:   let  $s'.nonces := s'.nonces + \langle \rangle n$ 
15: end for
16: let  $s'.cookies [s'.\bar{d}.origin.host]$ 
    $\hookrightarrow := \langle CookieMerge(s'.cookies [s'.\bar{d}.origin.host], cookies') \rangle$ 
17: let  $s'.localStorage [s'.\bar{d}.origin] := localStorage'$ 
18: let  $s'.sessionStorage [\langle s'.\bar{d}.origin, tlw.nonce \rangle] := sessionStorage'$ 
19: let  $s'.\bar{d}.scriptstate := state'$ 
20: switch command do
21:   case  $\langle \rangle$ 
22:     stop {},  $s'$ 
23:   case  $\langle HREF, url, hrefwindow \rangle^5$ 
24:     let  $\bar{w}', s' := GETNAVIGABLEWINDOW(\bar{w}, hrefwindow, s')$ 
25:     let  $req := \langle HTTPReq, n, GET, url.host, url.path, \langle \rangle, url.params, \langle \rangle \rangle$ 
26:     let  $s' := CANCELNAV(s'.\bar{w}'.nonce, s')$ 
27:     SEND( $s'.\bar{w}'.nonce, req, url.protocol, \perp, s'$ )
28:   case  $\langle IFRAME, url, window \rangle$ 
29:     let  $\bar{w}', s' := GETWINDOW(\bar{w}, window, s')$ 
30:     let  $req := \langle HTTPReq, n, GET, url.host, url.path, \langle \rangle, url.params, \langle \rangle \rangle$ 
31:     let  $n, s' := TAKENONCE(s')$ 
32:     let  $w' := \langle n, \langle \rangle, \perp \rangle$ 
33:     let  $s'.\bar{w}'.activedocument.subwindows$ 
    $\hookrightarrow := s'.\bar{w}'.activedocument.subwindows + \langle \rangle w'$ 
34:     SEND( $n, req, url.protocol, \perp, s'$ )
35:   case  $\langle FORM, url, method, data, hrefwindow \rangle$ 
36:     if  $method \notin \{GET, POST\}$  then 6
37:       stop {},  $s'$ 
38:     end if
39:     let  $\bar{w}', s' := GETNAVIGABLEWINDOW(\bar{w}, hrefwindow, s')$ 
40:     if  $method = GET$  then
41:       let  $body := \langle \rangle$ 
42:       let  $params := data$ 
43:       let  $origin := \perp$ 
44:     else
45:       let  $body := data$ 
46:       let  $params := url.params$ 
47:       let  $origin := s'.\bar{d}.origin$ 
48:     end if
49:     let  $req := \langle HTTPReq, n, method, url.host, url.path, \langle \rangle, params, body \rangle$ 
50:     let  $s' := CANCELNAV(s'.\bar{w}'.nonce, s')$ 
51:     SEND( $s'.\bar{w}'.nonce, req, url.protocol, origin, s'$ )
52:   case  $\langle SETSCRIPT, window, script \rangle$ 
53:     let  $\bar{w}', s' := GETWINDOW(\bar{w}, window, s')$ 
54:     let  $s'.\bar{w}'.activedocument.script := script$ 
55:     stop {},  $s'$ 

```

⁵See the definition of URLs in Appendix B.2.

⁶The working draft for HTML5 allowed for DELETE and PUT methods in HTML5 forms. However, these have since been removed. See <http://www.w3.org/TR/2010/WD-html5-diff-20101019/#changes-2010-06-24>.

```

56:   case ⟨SETSCRIPTSTATE, window, scriptstate⟩
57:     let  $\bar{w}'$ ,  $s'$  := GETWINDOW( $\bar{w}$ , window,  $s'$ )
58:     let  $s'.\bar{w}'$ .activedocument.scriptstate := scriptstate
59:     stop {},  $s'$ 
60:   case ⟨XMLHTTPREQUEST, url, method, data, xhrreference⟩
61:     if method ∈ {CONNECT, TRACE, TRACK} then
62:       stop {},  $s'$ 
63:     end if
64:     if url.host ≠  $s'.\bar{d}$ .origin.host
65:       ↪ ∨ url.protocol ≠  $s'.\bar{d}$ .origin.protocol then
66:         stop {},  $s'$ 
67:       end if
68:       if method ∈ {GET, HEAD} then
69:         let data := ⟨⟩
70:         let origin := ⊥
71:       else
72:         let origin :=  $s'.\bar{d}$ .origin
73:       end if
74:       let req := ⟨HTTPReq, n, method, url.host, url.path, url.params, data⟩
75:       SEND(⟨ $s'.\bar{d}$ .nonce, xhrreference⟩, req, url.protocol, origin,  $s'$ )
76:   case ⟨BACK, window⟩7
77:     let  $\bar{w}'$ ,  $s'$  := GETNAVIGABLEWINDOW( $\bar{w}$ , window,  $s'$ )
78:     if ∃  $\bar{j} ∈ \mathbb{N}, \bar{j} > 1$  such that  $s'.\bar{w}'$ .documents. $\bar{j}$ .active ≡ ⊤ then
79:       let  $s'.\bar{w}'$ .documents. $\bar{j}$ .active := ⊥
80:       let  $s'.\bar{w}'$ .documents. $(\bar{j}-1)$ .active := ⊤
81:       let  $s'$  := CANCELNAV( $s'.\bar{w}'$ .nonce,  $s'$ )
82:     end if
83:     stop {},  $s'$ 
84:   case ⟨FORWARD, window⟩
85:     let  $\bar{w}'$ ,  $s'$  := GETNAVIGABLEWINDOW( $\bar{w}$ , window,  $s'$ )
86:     if ∃  $\bar{j} ∈ \mathbb{N}$  such that  $s'.\bar{w}'$ .documents. $\bar{j}$ .active ≡ ⊤
87:       ↪ ∧  $s'.\bar{w}'$ .documents. $(\bar{j}+1) ∈ Documents$  then
88:         let  $s'.\bar{w}'$ .documents. $\bar{j}$ .active := ⊥
89:         let  $s'.\bar{w}'$ .documents. $(\bar{j}+1)$ .active := ⊤
90:         let  $s'$  := CANCELNAV( $s'.\bar{w}'$ .nonce,  $s'$ )
91:       end if
92:     stop {},  $s'$ 
93:   case ⟨CLOSE, window⟩
94:     let  $\bar{w}'$ ,  $s'$  := GETNAVIGABLEWINDOW( $\bar{w}$ , window,  $s'$ )
95:     remove  $s'.\bar{w}'$  from the sequence containing it
96:     stop {},  $s'$ 
97:   case ⟨POSTMESSAGE, window, message, origin⟩
98:     let  $\bar{w}' ← Subwindows(s')$  such that  $s'.\bar{w}'$ .nonce ≡ window

```

⁷Note that navigating a window using the back/forward buttons does not trigger a reload of the affected documents. While real world browser may chose to refresh a document in this case, we assume that the complete state of a previously viewed document is restored.

```

97:   if  $\exists \bar{j} \in \mathbb{N}$  such that  $s'.\bar{w}'.documents.\bar{j}.active \equiv \top$ 
       $\hookrightarrow \wedge (origin \neq \perp \implies s'.\bar{w}'.documents.\bar{j}.origin \equiv origin)$  then
98:     let  $s'.\bar{w}'.documents.\bar{j}.scriptinput$ 
       $\hookrightarrow := s'.\bar{w}'.documents.\bar{j}.scriptinput$ 
       $\hookrightarrow + \langle \rangle (POSTMESSAGE, s'.\bar{w}.nonce, s'.\bar{d}.origin, message)$ 
99:   end if
100: end function

```

The function PROCESSRESPONSE is responsible for processing an HTTP response (*response*) that was received as the response to a request (*request*) that was sent earlier. In *reference*, either a window or a document reference is given (see explanation for Algorithm 5 above). Again, *protocol* is either P or S.

The function first saves any cookies that were contained in the response to the browser state, then checks whether a redirection is requested (Location header). If that is not the case, the function creates a new document (for normal requests) or delivers the contents of the response to the respective receiver (for XHR responses).

Algorithm 7 Process an HTTP response.

```

1: function PROCESSRESPONSE(response, reference, request, protocol,  $s'$ )
2:   let  $n, s' := TAKENONCE(s')$ 
3:   if Set-Cookie  $\in response.headers$  then
4:     for each  $c \in \langle \rangle response.headers[Set-Cookie]$ ,  $c \in Cookies$  do
5:       let  $s'.cookies[request.url.host]$ 
       $\hookrightarrow := AddCookie(s'.cookies[request.url.host], c)$ 
6:     end for
7:   end if
8:   if Strict-Transport-Security  $\in response.headers \wedge protocol \equiv S$  then
9:     let  $s'.sts := s'.sts + \langle \rangle request.host$ 
10:  end if
11:  if Location  $\in response.headers \wedge response.status \in \{303, 307\}$  then8
12:    let  $url := response.headers[Location]$ 
13:    let  $method' := request.method$ 9
14:    let  $body' := request.body$ 10
15:    if Origin  $\in request.headers$  then
16:      let  $origin := \langle request.headers[Origin], \langle request.host, protocol \rangle \rangle$ 
17:    else
18:      let  $origin := \perp$ 

```

⁸The RFC for HTTPbis (currently in draft status), which obsoletes RFC 2616, does not specify whether a POST/DELETE/etc. request that was answered with a status code of 301 or 302 should be rewritten to a GET request or not (“for historic reasons” that are detailed in Section 7.4.). As the specification is clear for the status codes 303 and 307 (and most browsers actually follow the specification in this regard), we focus on modeling these.

⁹While the standard demands that users confirm redirections of non-safe-methods (e.g., POST), we assume that users generally confirm these redirections.

¹⁰If, for example, a GET request is redirected and the original request contained a body, this body is preserved, as HTTP allows for payloads in messages with all HTTP methods, except for the TRACE method (a detail which we omit). Browsers will usually not send body payloads for methods that do not specify semantics for such data in the first place.

```

19:   end if
20:   if  $response.status \equiv 303 \wedge request.method \notin \{GET, HEAD\}$  then
21:     let  $method' := GET$ 
22:     let  $body' := \langle \rangle$ 
23:   end if
24:   if  $\nexists \bar{w} \in Subwindows(s')$  such that  $s'.\bar{w}.nonce \equiv reference$  then
25:     stop  $\{\}, s$ 
26:   end if
27:   let  $req := \langle HTTPReq, n, method', url.host, url.path, \langle \rangle, url.params, body' \rangle$ 
28:   SEND( $reference, req, url.protocol, origin, s'$ )
29: end if
30: if  $\exists \bar{w} \in Subwindows(s')$  such that  $s'.\bar{w}.nonce \equiv reference$  then
31:   let  $script := \pi_1(response.body)$ 
32:   let  $scriptstate := \pi_2(response.body)$ 
33:   let  $d := \langle n, \langle request.host, request.protocol \rangle, script, scriptstate, \langle \rangle, \langle \rangle, \top \rangle$ 
34:   if  $s'.\bar{w}.documents \equiv \langle \rangle$  then
35:     let  $s'.\bar{w}.documents := \langle d \rangle$ 
36:   else
37:     let  $\bar{i} \leftarrow \mathbb{N}$  such that  $s'.\bar{w}.documents.\bar{i}.active \equiv \top$ 
38:     let  $s'.\bar{w}.documents.\bar{i}.active := \perp$ 
39:     remove  $s'.\bar{w}.documents.(\bar{i} + 1)$  and all following documents
40:      $\hookrightarrow$  from  $s'.\bar{w}.documents$ 
41:     let  $s'.\bar{w}.documents := s'.\bar{w}.documents + \langle \rangle d$ 
42:   end if
43:   stop  $\{\}, s'$ 
44:   else if  $\exists \bar{w} \in Subwindows(s'), \bar{d}$  such that  $s'.\bar{d}.nonce \equiv \pi_1(reference)$ 
45:      $\hookrightarrow \wedge s'.\bar{d} = s'.\bar{w}.activedocument$  then
46:     let  $s'.\bar{d}.scriptinput := s'.\bar{d}.scriptinput + \langle \rangle$ 
47:      $\langle P, response.body, \pi_2(reference) \rangle$ 
48:   end if
49: end function

```

Main Algorithm. This is the main algorithm of the browser relation. It was already presented informally in Section 2.5 and follows the structure presented there. It receives the message m as input, as well as a , f and s as above.

Algorithm 8 Main Algorithm

```

Input:  $(a:f:m), s$ 
1: let  $s' := s$ 
2: if  $s.isCorrupted \equiv FULLCORRUPT$  then
3:   let  $s'.pendingRequests := \langle m, s.pendingRequests \rangle$ 
4:   let  $m' \leftarrow d_{NP}(s')$ 
5:   let  $a' \leftarrow IPs$ 
6:   stop  $\{(a':a:m')\}, s'$ 
7: else if  $s.isCorrupted \equiv CLOSECORRUPT$  then

```

\triangleright Collect incoming messages


```

8:   let  $s'.pendingRequests := \langle m, s.pendingRequests \rangle$ 
9:   let  $N^{clean} := NP \setminus \{n | n \in \langle s.nonces \rangle\}$ 
10:  let  $m' \leftarrow d_{N^{clean}}(s')$ 
11:  let  $a' \leftarrow IPs$ 
12:  let  $s'.nonces := s.nonces$ 
13:  stop  $\{(a':a:m')\}, s'$ 
14: end if
15: let  $n, s' := TAKENONCE(s')$ 
16: if  $m \equiv TRIGGER$  then
17:   let  $switch \leftarrow \{1, 2\}$ 
18:   if  $switch \equiv 1$  then
19:     let  $\bar{w} \leftarrow Subwindows(s')$  such that  $s'.\bar{w}.documents \neq \langle \rangle$ 
20:      $\hookrightarrow$  if possible; otherwise stop  $\{\}, s'$ 
21:     let  $\bar{d} := \bar{w} + \langle \rangle$  activedocument
22:     RUNSCRIPT( $\bar{w}, \bar{d}, s'$ )
23:   else if  $switch \equiv 2$  then
24:     let  $w' := \langle n, \langle \rangle, \perp \rangle$ 
25:     let  $s'.windows := s'.windows + \langle \rangle w'$ 
26:     let  $protocol \leftarrow \{P, S\}$ 
27:     let  $host \leftarrow Doms$ 
28:     let  $path \leftarrow \mathbb{S}$ 
29:     let  $parameters \leftarrow [\mathbb{S} \times \mathbb{S}]$ 
30:     let  $n', s' := TAKENONCE(s')$ 
31:     let  $req := \langle HTTPReq, n', GET, host, path, \langle \rangle, parameters, \langle \rangle \rangle$ 
32:     SEND( $n, req, protocol, \perp, s'$ )
33:   end if
34: else if  $m \equiv FULLCORRUPT$  then
35:   let  $s'.isCorrupted := FULLCORRUPT$ 
36:   stop  $\{\}, s'$ 
37: else if  $m \equiv CLOSECORRUPT$  then
38:   let  $s'.secrets := \langle \rangle$ 
39:   let  $s'.windows := \langle \rangle$ 
40:   let  $s'.pendingDNS := \langle \rangle$ 
41:   let  $s'.pendingRequests := \langle \rangle$ 
42:   let  $s'.sessionStorage := \langle \rangle$ 
43:   let  $s'.cookies \subset \langle \rangle$  Cookies such that
44:    $\hookrightarrow (c \in \langle \rangle s'.cookies) \iff (c \in \langle \rangle s.cookies \wedge c.content.session \equiv \perp)$ 
45:   let  $s'.isCorrupted := CLOSECORRUPT$ 
46:   stop  $\{\}, s'$ 
47: else if  $\exists \langle reference, request, key, f \rangle \in \langle \rangle s'.pendingRequests$ 
48:    $\hookrightarrow$  such that  $\pi_1(dec_s(m, key)) \equiv HTTPResp$  then
49:     let  $m' := dec_s(m, key)$ 
50:     if  $m'.nonce \neq request.nonce$  then
51:       stop  $\{\}, s$ 
52:     end if
53:     remove  $\langle reference, request, key, f \rangle$  from  $s'.pendingRequests$ 
54:     PROCESSRESPONSE( $m', reference, request, S, s'$ )
55:   else if  $\pi_1(m) \equiv HTTPResp \wedge \exists \langle reference, request, \perp, f \rangle \in \langle \rangle s'.pendingRequests$ 
56:      $\hookrightarrow$  such that  $m'.nonce \equiv request.key$  then

```

```

53:   remove  $\langle \text{reference}, \text{request}, \perp, f \rangle$  from  $s'.\text{pendingRequests}$ 
54:   PROCESSRESPONSE( $m, \text{reference}, \text{request}, P, s'$ )
55: else if  $m \in \text{DNSResponses}$  then ▷ Successful DNS response
56:   if  $m.\text{nonce} \notin s.\text{pendingDNS}$  then
57:     stop  $\{\}, s$ 
58:   end if
59:   let  $\langle \text{reference}, \text{message}, \text{protocol} \rangle := s.\text{pendingDNS}[m.\text{nonce}]$ 
60:   if  $\text{protocol} \equiv \text{S}$  then
61:     let  $k, s' := \text{TAKENONCE}(s')$ 
62:     let  $s'.\text{pendingRequests} := s'.\text{pendingRequests}$ 
63:      $\hookrightarrow + \langle \text{reference}, \text{message}, k, m.\text{result} \rangle$ 
64:     let  $\text{message} := \text{enc}_a(\langle \text{message}, k \rangle, s'.\text{keyMapping}[\text{message}.\text{host}])$ 
65:   else
66:     let  $s'.\text{pendingRequests} := s'.\text{pendingRequests}$ 
67:      $\hookrightarrow + \langle \text{reference}, \text{message}, \perp, m.\text{result} \rangle$ 
68:   end if
69:   let  $s'.\text{pendingDNS} := s'.\text{pendingDNS} - m.\text{nonce}$ 
70:   stop  $\{(m.\text{result}:a:\text{message})\}, s'$ 
71: else
72:   stop  $\{\}, s$ 
73: end if

```

D General Security Properties of the Web Model

We now formally state and prove the general application independent security properties of the web which in Section 3 have been sketched only.

Let $\text{Web} = (\mathcal{W}, \mathcal{S}, \text{script}, E_0)$ be a web system. In the following, we write $s_x = (S_x, E_x)$ for the states of a web system.

Definition 31. *In what follows, given an atomic process p and a message m , we say that p emits m in a run $\rho = s_0, s_1, \dots$ if there is a processing step of the form*

$$s_{u-1} \xrightarrow[p \rightarrow E]{} s_u$$

for some $u \in \mathbb{N}$, a set of events E and some addresses x, y with $(x:y:m) \in E$.

Definition 32. *We say that a term t is derivably contained in (a term) t' for (a set of DY processes) P (in a processing step $s_i \rightarrow s_{i+1}$ of a run $\rho = s_0, s_1, \dots$) if t is derivable from t' with the knowledge available to P , i.e.,*

$$t \in d_\eta(\{t'\} \cup \zeta) \text{ with } \eta := \bigcup_{p \in P} N^p \text{ and } \zeta := \bigcup_{p \in P, j \leq i} S_j(p).$$

Definition 33. *We say that a set of processes P leaks a term t (in a processing step $s_i \rightarrow s_{i+1}$) to a set of processes P' if there exists a message m that is emitted (in $s_i \rightarrow s_{i+1}$) by some $p \in P$ and t is derivably contained in m for P' in the processing step $s_i \rightarrow s_{i+1}$. If we omit P' , we define $P' := \mathcal{W} \setminus P$. If P is a set with a single element, we omit the set notation.*

Definition 34. We say that an DY process p created a message m (at some point) in a run if m is derivably contained in a message emitted by p in some processing step and if there is no earlier processing step where m is derivably contained in a message emitted by some DY process p' .

Definition 35. We say that a browser b accepted a message (as a response to some request) if the browser decrypted the message (if it was an HTTPS message) and called the function PROCESSRESPONSE, passing the message and the request (see Algorithm 7).

Definition 36. We say that an atomic DY process p knows a term t in some state $s = (S, E)$ of a run if it can derive the term from its knowledge, i.e., $t \in d_{NP}(S(p))$.

Definition 37. We say that a script initiated a request r if a browser triggered the script (in Line 12 of Algorithm 6) and the first component of the command output of the script relation is either HREF, IFRAME, FORM, or XMLHTTPREQUEST such that the browser issues the request r in the same step as a result.

For a run $\rho = s_0, s_1, \dots$ of any \mathcal{Web} , we state the following lemmas:

Lemma 1. If in the processing step $s_i \rightarrow s_{i+1}$ of a run ρ of \mathcal{Web} an honest browser b (I) emits an HTTPS request of the form

$$m = \text{enc}_a(\langle req, k \rangle, \text{pub}(k'))$$

(where req is an HTTP request, k is a nonce (symmetric key), and k' is the private key of some other DY process u), and (II) in the initial state s_0 the private key k' is only known to u , and (III) u never leaks k' , then all of the following statements are true:

- (1) There is no state of \mathcal{Web} where any party except for u knows k' , thus no one except for u can decrypt req .
- (2) If there is a processing step $s_j \rightarrow s_{j+1}$ where the browser b leaks k to $\mathcal{W} \setminus \{u, b\}$ there is a processing step $s_h \rightarrow s_{h+1}$ with $h < j$ where u leaks the symmetric key k to $\mathcal{W} \setminus \{u, b\}$ or the browser is fully corrupted in s_j .
- (3) The value of the host header in req is the domain that is assigned the public key $\text{pub}(k')$ in the browsers' keymapping $s_0.\text{keymapping}$ (in its initial state).
- (4) If b accepts a response (say, m') to m in a processing step $s_j \rightarrow s_{j+1}$ and b is honest in s_j and u did not leak the symmetric key k to $\mathcal{W} \setminus \{u, b\}$ prior to s_j , then u created the HTTPS response m' to the HTTPS request m , i.e., the nonce of the HTTP request req is not known to any atomic process p , except for the atomic process b and u .

Proof. (1) follows immediately from the condition. If k' is initially only known to u and u never leaks k' , i.e., even with the knowledge of all nonces (except for those of u), k' can never be derived from any network output of u , k' cannot be known to any other party. Thus, nobody except for u can derive req from m .

(2) We assume that b leaks k to $\mathcal{W} \setminus \{u, b\}$ in the processing step $s_j \rightarrow s_{j+1}$ without u prior leaking the key k to anyone except for u and b and that the browser is not fully corrupted in s_j , and lead this to a contradiction.

The browser is honest in s_i . From the definition of the browser b , we see that the key k is always chosen from a fresh set of nonces (Line 61 of Algorithm 7) that are not used

anywhere else. Further, the key is stored in the browser's state in *pendingRequests*. The information from *pendingRequests* is not extracted or used anywhere else (in particular it is not accessible by scripts). If the browser becomes closecorrupted prior to s_j (and after s_i), the key cannot be used anymore (compare Line 9 of Algorithm 8). Hence, b does not leak k to any other party in s_j (except for u and b). This proves (2).

(3) Per the definition of browsers (Algorithm 8), a host header is always contained in HTTP requests by browsers. From Line 63 of Algorithm 8 we can see that the encryption key for the request req was chosen using the host header of the message. It is chosen from the *keymapping* in the browser's state, which is never changed during ρ . This proves (3).

(4) An HTTPS response m' that is accepted by b as a response to m has to be encrypted with k . The nonce k is stored by the browser in the *pendingRequests* state information. The browser only stores freshly chosen nonces there (i.e., the nonces are not used twice, or for other purposes than sending one specific request). The information cannot be altered afterwards (only deleted) and cannot be read except when the browser checks incoming messages. The nonce k is only known to u (which did not leak it to any other party prior to s_j) and b (which did not leak it either, as u did not leak it and b is honest, see (2)). The browser b cannot send responses. This proves (4). \square

Corollary 1. *In the situation of Lemma 1, as long as u does not leak the symmetric key k to $\mathcal{W} \setminus \{u, b\}$ and the browser does not become fully corrupted, k is not known to any DY process $p \notin \{b, u\}$ (i.e., $\nexists s' = (S', E') \in \rho : k \in d_{N^p}(S'(p))$).*

Lemma 2. *If for some $s_i \in \rho$ an honest browser b has a document d in its state $S_i(b).windows$ with the origin $\langle dom, S \rangle$ where $dom \in \text{Domain}$, and $S_i(b).keyMapping[dom] \equiv \text{pub}(k)$ with $k \in \mathcal{K}$ being a private key, and there is only one DY process p that knows the private key k in all s_j , $j \leq i$, then b extracted (in Line 33 in Algorithm 7) the script in that document from an HTTPS response that was created by p .*

Proof. The origin of the document d is set only once: In Line 33 of Algorithm 7. The values (domain and protocol) used there stem from the information about the request (say, req) that led to loading of d . These values have been stored in *pendingRequests* between the request and the response actions. The contents of *pendingRequests* are indexed by freshly chosen nonces and can never be altered or overwritten (only deleted when the response to a request arrives). The information about the request req was added to *pendingRequests* in Line 62 (or Line 65 which we can exclude as we will see later) of Algorithm 8. In particular, the request was an HTTPS request iff a (symmetric) key was added to the information in *pendingRequests*. When receiving the response to req , it is checked against that information and accepted only if it is encrypted with the proper key and contains the same nonce as the request (say, n). Only then the protocol part of the origin of the newly created document becomes S . The domain part of the origin (in our case dom) is taken directly from the *pendingRequests* and is thus guaranteed to be unaltered.

From Line 63 of Algorithm 8 we can see that the encryption key for the request req was actually chosen using the host header of the message which will finally be the value of the origin of the document d . Since b therefore selects the public key

$S_i(b).\text{keyMapping}[dom] = S_0(b).\text{keyMapping}[dom] \equiv \text{pub}(k)$ for p (the key mapping cannot be altered during a run), we can see that req was encrypted using a public key that matches a private key which is only (if at all) known to p . With Lemma 1 we see that the symmetric encryption key for the response, k , is only known to b and the respective web server. The same holds for the nonce n that was chosen by the browser and included in the request. Thus, no other party than p can encrypt a response that is accepted by the browser b and which finally defines the script of the newly created document. \square

Lemma 3. *If in a processing step $s_i \rightarrow s_{i+1}$ of a run ρ of Web an honest browser b issues an HTTP(S) request with the Origin header value $\langle dom, S \rangle$ where and $S_i(b).\text{keyMapping}[dom] \equiv \text{pub}(k)$ with $k \in \mathcal{K}$ being a private key, and there is only one DY process p that knows the private key k in all s_j , $j \leq i$, then that request was initiated by a script that b extracted (in Line 33 in Algorithm 7) from an HTTPS response that was created by p .*

Proof. First, we can see that the request was initiated by a script: As it contains an origin header, it must have been a POST request (see the browser definition in Appendix C.2). POST requests can only be initiated in Lines 51, 74 of Algorithm 6 and Line 28 of Algorithm 7. In the latter instance (Location header redirect), the request contains at least two different origins, therefore it is impossible to create a request with exactly the origin $\langle dom, S \rangle$ using a redirect. In the other two cases (FORM and XMLHttpRequest), the request was initiated by a script.

The Origin header of the request is defined by the origin of the script’s document. With Lemma 2 we see that the content of the document, in particular the script, was indeed provided by p . \square

E Step-By-Step Description of BrowserID (Primary IdP)

We now present additional details of the implementation of BrowserID. While the basic steps have been shown in Section 4.2, we will now again refer to Figure 3 and provide a step-by-step description. As above, for brevity of presentation, we focus on the main login flow without the CIF, and we leave out steps for fetching additional resources (like JavaScript files) and some less relevant postMessages and XHRs. Also, we assume that a typical IdP implementation like the example implementation provided by Mozilla is used.

We emphasize, however, that our formal model of BrowserID with primary IdPs (cf. Appendix F) closely follows the full BrowserID implementation (see also Figure 8 on Pages 85 and 86, which is an extended version of Figure 3).

E.1 LPO Sessions

Before we describe the login flow step-by-step, we first introduce LPO sessions.

LPO establishes a session with the browser by setting a cookie `browserid_state` (Step 5 in Figure 3) on the client-side. LPO considers such a session authenticated after having received a valid CAP (Step 22 in Figure 3). In future runs, the user is presented a list of her email addresses (which is fetched from LPO) in order to choose one address.

Then, she is asked if she trusts the computer she is using and is given the option to be logged in for one month or “for this session only” (*ephemeral* session). In order to use any of the email addresses, the user is required to authenticate to the IdP responsible for that address to get an UC issued. If the localStorage (under the origin LPO) already contains a valid UC, then, however, authentication at the IdP is not necessary.

E.2 Step-By-Step Description

We (again) assume that the user uses a “fresh” browser, i.e., the user has not been logged in before. The user has already opened a document of some RP (RP-Doc) in her browser. RP-Doc includes a JavaScript file, which provides the BrowserID API. The user is now about to click on a login button in order to start a BrowserID login.

Phase (i). After the user has clicked on the login button, RP-Doc opens a new browser window, the *login dialog* (LD) [1]. The document of LD is loaded from LPO [2]. Now, LD sends a *ready* postMessage [3] to its opener, which is RP-Doc. RP-Doc then responds by sending a *request* postMessage [4]. This postMessage may contain additional information like a name or a logo of RP-Doc. LD then fetches the so-called *session context* from LPO using [5]. The session context contains information about whether the user is already logged in at LPO, which, by our assumption, is not the case at this point. The session context also contains an XSRF protection token which will be sent in all subsequent POST requests to LPO. Also, an `httpOnly` cookie called `browserid_state` is set, which contains an LPO session identifier. Now, the user is prompted to enter her email address (*login email address*), which she wants to use to log in at RP [6]. LD sends the login email address to LPO via an XHR [7], in order to get information about the IdP the email address belongs to. The information from this so-called *support document* may be cached at LPO for further use. LPO extracts the domain part of the login email address and fetches an information document [8] from a fixed path (`/.well-known/browserid`) at the IdP. This document contains the public key of IdP, and two paths, the provisioning path and the authentication path at IdP. These paths will be used later in the login process by LD. LPO converts these paths into URLs and sends them in its response [9] to the requesting XHR [7].

Phase (ii). As there is no record about the login email address in the localStorage under the origin of LPO, the LD now tries to get a UC for this identity. For that to happen, the LD creates a new *iframe*, the *provisioning iframe* (PIF) [10]. The PIF’s document is loaded [11] from the provisioning URL LD has just received before in [9]. The PIF now interacts with the LD via postMessages [12]. As the user is currently not logged in, the PIF tells the LD that the user is not authenticated yet. This also indicates to the LD that the PIF has finished operation. The LD then closes the PIF [13].

Phase (iii). Now, the LD saves the login email address in the localStorage indexed by a fresh nonce. This nonce is stored in the sessionStorage to retrieve the email address later from the localStorage again. Next, the LD navigates itself to the authentication URL it has received in [9]. The loaded document now interacts with the user and the IdP [14] in order to establish some authenticated session depending on the actual IdP

implementation, which is out of scope of the BrowserID standard. For example, during this authentication procedure, the IdP may issue some session cookie.

Phase (iv). After the authentication to the IdP has been completed, the authentication document navigates the LD to the LD URL again. The LD's document is fetched again from LPO and the login process starts over. The following steps are similar to Phase (i): The ready and request postMessages are exchanged and the session context is fetched. As the user has not been authenticated to LPO yet, the session context still contains the same information as above in [5]. Now, the user is not prompted to enter her email address again. The email address is fetched from the localStorage under the index of the nonce stored in the sessionStorage. Now, the address information is requested again from LPO.

Phase (v). As there still is no UC belonging to the login email address in the localStorage, the PIF is created again. As the user now has established an authenticated session with the IdP, the PIF asks the LD to generate a fresh key pair. After the LD has generated the key pair [15], it stores the key pair in the localStorage (under the origin of LPO) and sends the public key to the PIF as a postMessage [16]. The following steps [17]–[19] are not specified in the BrowserID protocol. Typically, the PIF would send the public key to IdP (via an XHR) [17]. The IdP would create the UC [18] and send it back to the PIF [19]. The PIF then sends the UC to the LD [20], which stores it in the localStorage. Now, the LD closes the PIF.

Phase (vi). The LD is now able to create a CAP, as it has access to a UC and the corresponding private key in its localStorage. First, LD creates an IA for LPO [21]. The IA and the UC is then combined to a CAP, which is then sent to LPO in an XHR POST message [22]. LPO is now able to verify this CAP with the public key of IdP, which LPO has already fetched and cached before in [8]. If the CAP is valid, LPO considers its session with the user's browser to be authenticated for the email address the UC in the CAP is issued for.

Phase (vii). Now, in [23], the LD fetches a list of email addresses, which LPO considers to be owned by the user. If the login email address would not appear in this list, LD would abort the login process. After this, the LD fetches the address information about the login email address again in [24]. Using this information, LD validates if the UC is signed by the correct party (primary/secondary IdP). Now, LD generates an IA for the sender's origin of the request postMessage [4] (which was repeated in Phase (iv)) using the private key from the localStorage [25] (the IA is generated for the login email address). Also, it is recorded in the localStorage that the user is now logged in at RP with this email address. The LD then combines the IA with the UC stored in the localStorage to the CAP, which is then sent to RP-Doc in the *response* postMessage [26].

This concludes the login process that runs in LD. Afterwards, RP-Doc closes LD [27].

F Model of BrowserID with Primary IdPs

We now present the full details of our formal model of BrowserID with primary IdPs and the fixes discussed in Section 5.3 applied. We consider ephemeral sessions (the default), which are supposed to last until the browser is closed.

We model the BrowserID system as a web system (in the sense of Section 2). We call a web system $BID = (\mathcal{W}, \mathcal{S}, \text{script}, E_0)$ a *BrowserID web system* if it is of the form described in what follows.

F.1 Outline

The system $\mathcal{W} = \text{Hon} \cup \text{Web} \cup \text{Net}$ consists of the (network) attacker process attacker, the web server for LPO, a finite set B of web browsers, a finite set RP of web servers for the relying parties, and a finite set IDP of web servers for the identity providers, with $\text{Hon} := B \cup RP \cup IDP \cup \{\text{LPO}\}$, $\text{Web} := \emptyset$, and $\text{Net} := \{\text{attacker}\}$. DNS servers are assumed to be dishonest, and hence, are subsumed by attacker. More details on the processes in \mathcal{W} are provided below. Figure 7 shows the set of scripts \mathcal{S} and their respective string representations that are defined by the mapping script . The set E_0 contains only the trigger events as specified in Section 2.3.

$s \in \mathcal{S}$	$\text{script}(s)$
R^{att}	att_script
script_{rp_index}	script_rp_index
script_{lpo_cif}	script_lpo_cif
script_{lpo_ld}	script_lpo_ld
script_{idp_pif}	script_idp_pif
script_{idp_ad}	script_idp_ad

Fig. 7. List of scripts in \mathcal{S} and their respective string representations.

This outlines BID . We will now define the DY processes in BID and their addresses, domain names, and secrets in more detail.

F.2 Addresses and Domain Names

The set IPs contains for LPO, attacker, every relying party in RP , every identity provider in IDP , and every browser in B one address each. By addr we denote the corresponding assignment from a process to its address. The set $Doms$ contains one domain for LPO, one for every relying party in RP , a finite set of domains for every identity provider in IDP , and a finite set of domains for attacker. Browsers (in B) do not have a domain.

By addr and dom we denote the assignments from atomic processes to sets of IPs and $Doms$, respectively. If dom or addr returns a set with only one element, we often write $\text{dom}(x)$ or $\text{addr}(x)$ to refer to the element.

F.3 Keys and Secrets

The set \mathcal{N} of nonces is partitioned into four sets, an infinite set $N^{\mathcal{W}}$, an infinite set K_{SSL} , an infinite set K_{sign} , and a finite set Secrets. We thus have

$$\mathcal{N} = \underbrace{N^{\mathcal{W}}}_{\text{infinite}} \dot{\cup} \underbrace{K_{\text{SSL}}}_{\text{finite}} \dot{\cup} \underbrace{K_{\text{sign}}}_{\text{finite}} \dot{\cup} \underbrace{\text{Secrets}}_{\text{finite}} .$$

The set $N^{\mathcal{W}}$ contains the nonces that are available for each DY process in \mathcal{W} . It is partitioned into infinite sets of nonces, one set $N^p \subseteq N^{\mathcal{W}}$ for every $p \in \mathcal{W}$.

The set K_{SSL} contains the keys that will be used for SSL encryption. Let $\text{sslkey}: \text{Doms} \rightarrow K_{\text{SSL}}$ be an injective mapping that assigns a (different) private key to every domain.

The set K_{sign} contains the keys that will be used by IdPs for signing UCs. Let $\text{signkey}: \text{IdPs} \rightarrow K_{\text{sign}}$ be an injective mapping that assigns a (different) private key to every identity provider.

The set $\text{Secrets} \subseteq \mathcal{N}$ is the set of passwords (secrets) the browsers share with the identity providers.

F.4 Identities

Identities are email addresses, which consist of a user name and a domain part. For our model, this is defined as follows:

Definition 38. An identity (email address) i is a term of the form $\langle \text{name}, \text{domain} \rangle$ with $\text{name} \in \mathbb{S}$ and $\text{domain} \in \text{Doms}$.

Let ID be the finite set of identities. By ID^y we denote the set $\{\langle \text{name}, \text{domain} \rangle \in \text{ID} \mid \text{domain} \in \text{dom}(y)\}$.

We say that an ID is governed by the DY process to which the domain of the ID belongs. Formally, we define the mapping $\text{governor}: \text{ID} \rightarrow \mathcal{W}$, $\langle \text{name}, \text{domain} \rangle \mapsto \text{dom}^{-1}(\text{domain})$.

The governor of an ID will usually be an IdP, but could also be the attacker. Note that we omit delegation of authority over domains.

We further define UCs, IAs and CAPs formally:

Definition 39. A (valid) user certificate (UC) uc for a user u with email address $id = \langle \text{name}, d \rangle$ and public key (verification key) $\text{pub}(k_u)$, where $d \in \text{dom}(y)$ is a domain of the governor y of id and k_u is the private key (signing key) of u , is a message of the form $uc = \text{sig}(\langle \langle \text{name}, d \rangle, \text{pub}(k_u) \rangle, \text{signkey}(y))$.

An (valid) identity assertion (IA) ia for an origin o (e.g., $\langle \text{example.com}, \mathbb{S} \rangle$) signed with the key k_u is a message of the form $ia = \text{sig}(o, k_u)$.

A certificate assertion pair (CAP) is of the form $\langle uc, ia \rangle$, with uc and ia as above.¹¹

¹¹Note that the time stamps are omitted both from the UC and the IA. This models that both certificates are valid indefinitely. In reality, they are valid for a certain period of time, as indicated by the time stamps. So our modeling is a safe overapproximation.

Each browser $b \in \mathcal{B}$ owns a set of secrets ($\in \text{Secrets}$). Each secret is assigned a set S of IDs for a specific IdP y such that $S \subseteq \text{ID}^y$. Browsers have disjoint secrets and secrets have disjoint sets of IDs. The IdPs of the secrets of a browser are disjoint. An ID i is owned by a browser b if the identity associated with i belongs to b :

Let $\text{ownerOfSecret} : \text{Secrets} \rightarrow \mathcal{B}$ denote the mapping that assigns to each secret the browser that owns this secret. Let $\text{secretOfID} : \text{ID} \rightarrow \text{Secrets}$ denote the mapping that assigns to each identity the associated secret. Now, we define the mapping $\text{ownerOfID} : \text{ID} \rightarrow \mathcal{B}$, $i \mapsto \text{ownerOfSecret}(\text{secretOfID}(i))$, which assigns to each identity the browser that owns this identity (we say that the identity belongs to the browser).

F.5 Corruption

RPs and IdPs can become corrupted: If they receive the message `CORRUPT`, they start collecting all incoming messages in their state and (upon triggering) send out all messages that are derivable from their state and collected input messages, just like the attacker process. We say that an RP or IdP is *honest* if the according part of their state ($s.\text{corrupt}$) is \perp , and that they are corrupted otherwise.

We are now ready to define the processes in \mathcal{W} as well as the scripts in \mathcal{S} in more detail.

F.6 Processes in \mathcal{W} (Overview)

We first provide an overview of the processes in \mathcal{W} . All processes in \mathcal{W} contain in their initial states all public keys and the private keys of their respective domains (if any). We define $I^p = \{\text{addr}(p)\}$ for all $p \in \text{Hon}$.

Attacker. The attacker process is a network attacker (see Section 2.3), who uses all addresses for sending and listening. All parties use the attacker as a DNS server. See Appendix F.7 for details.

Browsers. Each $b \in \mathcal{B}$ is a web browser as defined in Section 2.5. The initial state contains all secrets owned by b , stored under the origin of the respective IdP. See Appendix F.8 for details.

LPO. LPO is a web server that serves important scripts (`script_lpo_cif` and `script_lpo_ld`) and manages user sessions. See Appendix F.9 for details.

IdPs. Each IdP is a web server. IdPs are modeled following the example implementation provided by Mozilla. As outlined in Section 4, users can authenticate to the IdP with their credentials. IdP tracks the state of the users with sessions. Authenticated users can receive signed UCs from the IdP. When receiving a special message (`CORRUPT`) IdPs can become corrupted. Similar to the definition of corruption for the browser, IdPs then start sending out all messages that are derivable from their state. See Appendix F.11 for details.

Relying Parties. A relying party $r \in \text{RP}$ is a web server. The definition of R^r follows the description in Section 4 and the security considerations in [22] (Cross-site Request Forgery protection, e.g., by checking origin headers, and HTTPS only with STS enabled).

RP answers any GET request with the script `script_rp_index` (see below). When receiving an HTTPS POST message, RP checks (among others) if the message contains a valid CAP. For this purpose, all signing keys of the identity providers (see below) are contained in the initial state of all RPs. If successful, RP responds with an *RP service token for ID i* of the form $\langle n, i \rangle$, where $i \in \text{ID}$ is the ID for which the CAP was issued and n is a freshly chosen nonce. The RP r keeps a list of such tokens in its state. Intuitively, a client having such a token can use the service of r for ID i . See Appendix F.10 for details. Just like IdPs, RPs can become corrupted.

F.7 Attacker

As mentioned, the attacker `attacker` is modeled to be a network attacker as specified in Section 2.3. We allow it to listen to/spoof all available IP addresses, and hence, define $I^{\text{attacker}} = \text{IPs}$. His initial state is $s_0^{\text{attacker}} = \langle \text{attdoms}, \text{sslkeys}, \text{signkeys} \rangle$, where *attdoms* is a sequence of all domains along with the corresponding private keys owned by the attacker, *sslkeys* is a sequence of all domains and the corresponding public keys, and *signkeys* is a sequence containing all public signing keys for all IdPs. All other parties use the attacker as a DNS server.

F.8 Browsers

Each $b \in \text{B}$ is a web browser as defined in Section 2.5, with $I^b := \{\text{addr}(b)\}$ being its address.

To define the initial state, first let $ID^b := \text{ownerOfID}^{-1}(b)$ be the set of all IDs of b , $ID^{b,d} := \{i \mid \exists x : i = \langle x, d \rangle \in ID^b\}$ be the set of IDs of b for a domain d , and $\text{SecretDomains}^b := \{d \mid ID^{b,d} \neq \emptyset\}$ be the set of all domains that b owns identities for.

Then, the initial state s_0^b is defined as follows: the key mapping maps every domain to its public (ssl) key, according to the mapping *sslkey*; the DNS address is `addr(attacker)`; the list of secrets contains an entry $\langle \langle d, \mathcal{S} \rangle, s \rangle$ for each $d \in \text{SecretDomains}^b$ and $s = \text{secretOfID}(i)$ for some $i \in ID^{b,d}$ (s is the same for all i); *ids* is $\langle ID^b \rangle$; *sts* is empty.

F.9 LPO

LPO is a an atomic DY process $(I^{\text{LPO}}, Z^{\text{LPO}}, R^{\text{LPO}}, s_0^{\text{LPO}}, N^{\text{LPO}})$ with the IP address $I^{\text{LPO}} = \{\text{addr}(\text{LPO})\}$. The initial state s_0^{LPO} of LPO contains the private key of its domain, and the signing keys of all IdPs (LPO does not need the public ssl keys of other parties, which is why we omit them from LPO's initial state.). The definition of R^{LPO} follows the description of LPO in Appendix E.

HTTP responses by LPO can contain strings representing scripts, namely the script `script_lpo_cif` run in the CIF and the script `script_lpo_ld` run in the LD. These scripts are defined in Appendix F.12.

Client sessions at LPO. Any party can establish a *session* at LPO. Such a session can either be authenticated or unauthenticated. Roughly speaking, a session becomes authenticated if a client has provided a valid CAP (for the origin of LPO) to LPO during the session. LPO manages groups of IDs, i.e., lists of email addresses. If a user

authenticates a session using any ID in the group, she is authenticated for all IDs in the group. An authenticated session can (non-deterministically) *expire*, i.e. the authenticated session can get unauthenticated or it is removed completely. Such an expiration is used to model a user logout or a session expiration caused by a timeout.

More specifically, a session is identified by a nonce, which is issued by LPO. Each session is associated with some `xsrftoken`, which is also a nonce issued by LPO. LPO stores all information about established sessions in its state as a dictionary indexed by the session identifier. In this dictionary, for every session LPO stores a pair containing the `xsrftoken` and, in authenticated sessions, the sequence of all IDs associated with the secret provided in the session, or, in unauthenticated sessions, the empty sequence $\langle \rangle$ of IDs. On the receiver side (typically a browser) LPO places, by appropriate headers in its HTTPS responses, a cookie named `browserid_state` whose value is the session identifier (a nonce). This cookie is flagged to be a session, `httpOnly`, and secure cookie.

Before we provide a detailed formal specification of LPO, we first provide an informal description.

HTTPSRequests to LPO. LPO answers only to certain requests (listed below). In reality, all such requests have to be over HTTPS, and all responses send by LPO contain the `Strict-Transport-Security` header. We overapproximate safely here in omitting these two requirements from the model.

GET `/cif`. LPO replies to this request by providing the script `script_lpo_cif`.
 GET `/ld`. LPO replies to this request by providing the script `script_lpo_ld`.
 GET `/ctx`. This requests the session context information from LPO. The response body is of the form $\langle \text{loggedIn}, \text{xsrftoken} \rangle$, where *loggedIn* is \top or \perp , depending on whether the user is logged in at LPO or not, and *xsrftoken* is the token that the client is supposed to include into the auth request (see below).
 POST `/auth`. With this request, a client can log into LPO. The client has to provide a sequence of a CAP and an XSRF token. The CAP must be valid and issued for the origin of LPO.

We define LPO formally as an atomic DY process $(I^{\text{LPO}}, Z^{\text{LPO}}, R^{\text{LPO}}, s_0^{\text{LPO}}, N^{\text{LPO}})$. As already mentioned, we define $I^{\text{LPO}} = \{\text{addr}(\text{LPO})\}$.

In order to define the set Z^{LPO} of states of LPO, we first define the terms describing the session context of a session.

Definition 40. A term of the form $\langle \text{ids}, \text{xsrftoken} \rangle$ with $\text{ids} \subset \langle \rangle \text{ID}$ and $\text{xsrftoken} \in \mathcal{N}$ is called an LPO session context. We denote the set of all LPO session contexts by `LPOSessionCTXs`.

Now, we define the set Z^{LPO} of states of LPO as well as the initial state s_0^{LPO} of LPO.

Definition 41. A state $s \in Z^{\text{LPO}}$ of LPO is a term of the form $\langle \text{nonces}, \text{sslkey}, \text{signkeys}, \text{sessions} \rangle$ where $\text{nonces} \subset \langle \rangle \mathcal{N}$ (used nonces), $\text{sslkey} = \text{sslkey}(\text{dom}(\text{LPO}))$, *signkeys* is a mapping of domain names to public signing keys of the form $\text{signkeys} = \langle \{ \langle d, \text{pub}(\text{signkey}(y)) \rangle \mid y \in \text{IdPs}, d \in \text{dom}(y) \} \rangle$, and $\text{sessions} \in [\mathcal{N} \times \text{LPOSessionCTXs}]$.¹²

¹²As mentioned before, the state of LPO does not need to contain public keys.

The initial state s_0^{LPO} of LPO is a state of LPO with $s_0^{\text{LPO}}.\text{nonces} = \langle \rangle$ and $s_0^{\text{LPO}}.\text{sessions} = \langle \rangle$.

Example 7. Let k be a private signing key for some identity provider which owns the domain `example.com`. A possible state s of LPO may look like this:

$$s = \langle \langle n_1, \dots, n_m \rangle, \text{sslkey}(\text{dom}(\text{LPO})), [\text{example.com} : \text{pub}(k)], \text{sessions} \rangle$$

with

$$\text{sessions} = \langle \langle \text{sessionId}_1, \langle \langle id'_1, \dots, id'_l \rangle, \text{xsrftoken} \rangle \rangle, \dots \rangle$$

We now specify the relation $R^{\text{LPO}} \subseteq (\mathcal{E} \times Z^{\text{LPO}}) \times (2^{\mathcal{E}} \times Z^{\text{LPO}})$ of LPO. Just like in Appendix C.2, we describe this relation by a non-deterministic algorithm.

Algorithm 9 Relation of LPO R^{LPO}

Input: $\langle a:f:m \rangle, s$

- 1: **let** $s' := s$
- 2: **let** $sts := \langle \text{Strict-Transport-Security}, \top \rangle$
- 3: **if** $m \equiv \text{TRIGGER}$ **then**
- 4: **if** $s'.\text{sessions} \equiv \langle \rangle$ **then**
- 5: **stop** $\{ \}, s$
- 6: **end if**
- 7: **let** $\text{sessionId} \leftarrow \{ id \mid id \in \langle \rangle s'.\text{sessions} \}$
- 8: **let** $\text{choice} \leftarrow \{ \text{logout}, \text{expire} \}$
- 9: **if** $\text{choice} \equiv \text{logout}$ **then**
- 10: **let** $s'.\text{sessions}[\text{sessionId}].\text{ids} := \langle \rangle$
- 11: **else**
- 12: **let** $s'.\text{sessions} := s'.\text{sessions} - \text{sessionId}$
- 13: **end if**
- 14: **stop** $\{ \}, s$
- 15: **end if**
- 16: **let** m_{dec}, k **such that** $\langle m_{\text{dec}}, k \rangle \equiv \text{dec}_a(m, s.\text{sslkey})$
 \hookrightarrow **if possible; otherwise stop** $\{ \}, s$
- 17: **let** $n, \text{method}, \text{path}, \text{params}, \text{headers}, \text{body}$ **such that**
 $\hookrightarrow \langle \text{HTTPReq}, n, \text{method}, \text{dom}(\text{LPO}), \text{path}, \text{params}, \text{headers}, \text{body} \rangle \equiv m_{\text{dec}}$
 \hookrightarrow **if possible; otherwise stop** $\{ \}, s$
- 18: **if** $\text{method} \equiv \text{GET} \wedge \text{path} \equiv / \text{cif}$ **then** ▷ Deliver CIF script
- 19: **let** $\text{scriptinit} := \langle \text{init}, \perp, \perp, \perp, \perp, \perp, \langle \rangle, \perp, \perp \rangle$
- 20: **let** $m' := \text{enc}_s(\langle \text{HTTPResp}, n, 200, \langle sts \rangle, \langle \text{script_lpo_cif}, \text{scriptinit} \rangle \rangle, k)$
- 21: **stop** $\{ (f:a:m') \}, s'$
- 22: **else if** $\text{method} \equiv \text{GET} \wedge \text{path} \equiv / \text{ld}$ **then** ▷ Deliver LD script
- 23: **let** $\text{scriptinit} := \langle \text{init}, \perp, \perp, \perp, \perp, \perp, \langle \rangle, \perp, \perp \rangle$
- 24: **let** $m' := \text{enc}_s(\langle \text{HTTPResp}, n, 200, \langle sts \rangle, \langle \text{script_lpo_ld}, \text{scriptinit} \rangle \rangle, k)$
- 25: **stop** $\{ (f:a:m') \}, s'$
- 26: **else if** $\text{method} \equiv \text{GET} \wedge \text{path} \equiv / \text{ctx}$ **then** ▷ Deliver context info.
- 27: **let** $\text{sessionId} := \text{headers}[\text{Cookie}][\text{browserid_state}]$
- 28: **if** $\text{sessionId} \notin \langle \rangle s'.\text{sessions}$ **then**
- 29: **let** $\text{sessionId}, s' := \text{TAKENONCE}(s')$

```

30:     let xsrfToken, s' := TAKENONCE(s')
31:     let s'.sessions := s'.sessions +⟨⟩ ⟨sessionid, ⟨⟩, xsrfToken⟩
32:   end if
33:   let context := ⟨⊥, s'.sessions[sessionid].xsrfToken⟩
34:   if s'.session[sessionid].ids ≠ ⟨⟩ then
35:     let context.l := ⊤
36:   end if
37:   let setCookie := ⟨Set-Cookie, ⟨⟨browserid_state, sessionid, ⊤, ⊤, ⊤⟩⟩⟩
38:   let headers := ⟨sts, setCookie⟩
39:   let m' := encs(⟨HTTPResp, n, 200, headers, context⟩, k)
40:   stop {(f:a:m')}, s'
41: else if method ≡ POST ∧ path ≡ /auth then
42:   let uc, ia, xsrfToken such that ⟨⟨uc, ia⟩, xsrfToken⟩ ≡ body
    ↪ if possible; otherwise stop {}, s
43:   let sessionid := headers[Cookie][browserid_state]
44:   if s'.sessions[sessionid].xsrfToken ≠ xsrfToken then
45:     stop {}, s
46:   end if
47:   let name, domain, userpubkey such that
    ↪ ⟨⟨name, domain⟩, userpubkey⟩ ≡ extractmsg(uc)
    ↪ if possible; otherwise stop {}, s
48:   let id := ⟨name, domain⟩
49:   let origin := extractmsg(ia)
50:   if checksig(uc, s.signkeys[domain]) ≠ ⊤ ∨ checksig(ia, userpubkey) ≠ ⊤
    ↪ ∨ origin ≠ ⟨s.domain, S⟩ then
51:     stop {}, s
52:   end if
53:   if s'.sessions[sessionid].ids ≡ ⟨⟩ then
54:     if ∃ n ∈ ℕ such that id ∈⟨⟩ s'.idgroups.n then
55:       let s'.idgroups := s'.idgroups +⟨⟩ ⟨id⟩
56:     end if
57:     let n ← ℕ such that id ∈⟨⟩ s'.idgroups.n
58:   else
59:     let n ← ℕ such that s'.idgroups.n ≡ s'.sessions[sessionid].ids
    ↪ if possible; otherwise stop {}, s
60:     if id ∉⟨⟩ s'.idgroups.n then
61:       let s'.idgroups.n := s'.idgroups.n +⟨⟩ ⟨name, domain⟩
62:     end if
63:   end if
64:   let s'.sessions[sessionid].ids := s'.idgroups.n
65:   let m' := encs(⟨HTTPResp, n, 200, ⟨sts⟩, ⊤⟩, k)
66:   stop {(f:a:m')}, s'
67: end if
68: stop {}, s

```

F.10 Relying Parties

A relying party $r \in \text{RP}$ is a web server modeled as an atomic DY process $(I^r, Z^r, R^r, s_0^r, N^r)$ with the address $I^r := \{\text{addr}(r)\}$. Its initial state s_0^r contains its domain, the private key associated with its domain, the DNS server address, and the

signing keys of all IdPs.¹³ The full state additionally contains the set of service tokens the RP has issued. The definition of R^r again follows the description in Appendix E. RP accepts only HTTPS requests.

In a typical flow with one client, r will first receive an HTTP GET request. In this case, it returns the script `script_rp_index` (see Appendix F.12 below) and sets the Strict-Transport-Security header.

Afterwards, it will receive an HTTPS POST request. Provided that the message contains a CAP, r checks that the UC and IA are valid and matching, and that the IA contains the Origin of r (with HTTPS). If the check is successful, r creates a new *RP service token for the identity* i , $\langle n, i \rangle$, and sends it to the browser. The RP keeps a list of such tokens in its state. Intuitively, a client in possession of such a token can use the service of r for ID i (e.g., access data of i at r).

We now provide the formal definition of r as an atomic DY process $(I^r, Z^r, R^r, s_0^r, N^r)$. As mentioned, we define $I^r = \{\text{addr}(r)\}$. Next, we define the set Z^r of states of r and the initial state s_0^r of r .

Definition 42. A state $s \in Z^r$ of an RP r is a term of the form $\langle \text{nonces}, \text{domain}, \text{sslkey}, \text{signkeys}, \text{serviceTokens}, \text{corrupt} \rangle$ where $\text{nonces} \subset \langle \mathcal{N} \text{ (used nonces)}, \text{domain} = \text{dom}(r), \text{sslkey} = \text{sslkey}(\text{dom}(r)), \text{signkeys} = \{ \langle d, \text{pub}(\text{signkey}(y)) \rangle \mid y \in \text{IdPs}, d \in \text{dom}(y) \}$ (same as for LPO), $\text{serviceTokens} \in [\mathcal{N} \times \mathbb{S}]$, $\text{corrupt} \in \mathcal{T}_{\mathcal{N}}$.

The initial state s_0^r of r is a state of r with $s_0^r.\text{nonces} = s_0^r.\text{serviceTokens} = \langle \rangle$ and $s_0^r.\text{corrupt} = \perp$.

We now specify the relation $R^r \subseteq (\mathcal{E} \times Z^r) \times (2^{\mathcal{E}} \times Z^r)$ of r . Just like in Appendix C.2, we describe this relation by a non-deterministic algorithm. We note that we use the function TAKENONCE introduced in Section C.2 for this purpose.

Algorithm 10 Relation of a Relying Party R^r

Input: $(a:f:m), s$

- 1: **let** $s' := s$
- 2: **if** $s'.\text{corrupt} \neq \perp \vee m \equiv \text{CORRUPT}$ **then**
- 3: **let** $s'.\text{corrupt} := \langle \langle a, f, m \rangle, s'.\text{corrupt} \rangle$
- 4: **let** $m' \leftarrow d_{N^r}(s')$
- 5: **let** $a' \leftarrow \text{IPs}$
- 6: **stop** $\{ \langle a':a:m' \rangle \}, s'$
- 7: **end if**
- 8: **let** $sts := \langle \text{Strict-Transport-Security}, \top \rangle$
- 9: **let** m_{dec}, k **such that** $\langle m_{\text{dec}}, k \rangle \equiv \text{dec}_a(m, s.\text{sslkey})$
 \hookrightarrow **if possible; otherwise stop** $\{ \}, s$
- 10: **let** $n, \text{method}, \text{path}, \text{params}, \text{headers}, \text{body}$ **such that**
 $\hookrightarrow \langle \text{HTTPReq}, n, \text{method}, s.\text{domain}, \text{path}, \text{params}, \text{headers}, \text{body} \rangle \equiv m_{\text{dec}}$
 \hookrightarrow **if possible; otherwise stop** $\{ \}, s$
- 11: **if** $\text{method} \equiv \text{GET}$ **then** \triangleright Deliver CIF script
- 12: **let** $\text{scriptinit} := \langle \text{init}, \perp, \perp, \perp, \langle \rangle, \langle \rangle, \perp \rangle$
- 13: **let** $m' := \text{enc}_s(\langle \text{HTTPResp}, n, 200, \langle sts \rangle, \langle \text{script_rp_index}, \text{scriptinit} \rangle \rangle, k)$

¹³We add the IdP verification keys to the initial status (instead of having RPs retrieve them dynamically from the IdP) in order to reduce the overall complexity.

```

14:   stop  $\{(f:a:m')\}, s'$ 
15: else if  $method \equiv POST \wedge headers \equiv \langle Origin, \langle s.domain, S \rangle \rangle$  then
16:   let  $uc, ia$  such that  $\langle uc, ia \rangle \equiv body$  if possible; otherwise stop  $\{\}, s$ 
17:   let  $name, domain, userpubkey$  such that
      $\hookrightarrow \langle \langle name, domain \rangle, userpubkey \rangle \equiv extractmsg(uc)$ 
      $\hookrightarrow$  if possible; otherwise stop  $\{\}, s$ 
18:   let  $id := \langle name, domain \rangle$ 
19:   let  $origin := extractmsg(ia)$ 
20:   if  $checksig(uc, s.signkeys[domain]) \not\equiv \top \vee checksig(ia, userpubkey) \not\equiv \top \vee$ 
      $origin \not\equiv \langle s.domain, S \rangle$  then
21:     stop  $\{\}, s$ 
22:   end if
23:   let  $n_{token}, s' := TAKENONCE(s')$ 
24:   let  $s'.serviceTokens := s'.serviceTokens + \langle n_{token}, id \rangle$ 
25:   let  $m' := enc_s(\langle HTTPResp, n, 200, \langle sts, \langle n_{token}, id \rangle \rangle, k)$ 
26:   stop  $\{(f:a:m')\}, s'$ 
27: end if

```

F.11 Identity Providers

An identity provider $i \in \text{IdPs}$ is a web server modeled as an atomic process $(I^i, Z^i, R^i, s_0^i, N^i)$ with the address $I^i := \{\text{addr}(i)\}$. Its initial state s_0^i contains a list of domains and (private) SSL keys (see below), a list of users and identities (see below), and a private key for signing UCs. Besides this, the full state of i further contains a list of used nonces, and information about active sessions.

Sessions are structured as a dictionary: For each session identifier (session ID) the dictionary contains the list of identities for which the session is authenticated.

IdPs, in our model, only accept SSL connections. Thus, after receiving a request, an IdP first decodes the message. It then checks whether a valid session ID is contained in the cookie that was sent with the request. If there is no such ID, a new session with a freshly chosen session ID is created. IdP saves this ID into its list of active sessions, along with the initial session data (an empty list of authenticated identities). A Set-Cookie header is added to IdPs response to the browser in order to add the session cookie to the client's cookie store.

The IdP then checks the method and the path of the request and acts as follows:

If the method is GET, IdP serves, depending on the path, the provisioning iframe (`script_idp_pif`) or the authentication dialog (`script_idp_ad`) defined in Appendix F.12.

If the method is POST, the IdP can either authenticate the user or sign a UC. In the first case, IdP extracts the identity of the user (an email address) and the user's secret from the request. If the secret and the identity are found in the user database, the session is considered to be logged in for all identities associated with this secret. In the second case (signing UC), the IdP extracts the user's identity and the public key of the user from the request. If the session is considered to be logged in for this identity, the IdP creates a UC and signs it with its signing key before sending it to the user.

Formal description. In the following, we will first define the (initial) state of i formally and afterwards present the definition of the relation R^i .

To define the initial state, we will need a term that represents the “user database” of the IdP i . We will call this term $userset^i$. This database defines, which secret is valid for which set of identities. It is encoded as a mapping of secrets to lists of identities for which these secrets are valid. For example, if the secret $secret_1$ is valid for the identities id_1 and id_2 and the secret $secret_2$ is valid for the identities id_3 and id_4 , the $userset^i$ looks as follows:

$$userset^i = [secret_1:\langle id_1, id_2 \rangle, secret_2:\langle id_3, id_4 \rangle]$$

To define $userset^i$ (for the identity provider i), we first define the set $Secrets^i = \bigcup_{j \in ID^i} secretOfID(j)$, the function $IDsofSecret : Secrets \rightarrow ID, s \mapsto \{j \mid j \in ID, secretOfID(j) = s\}$, and finally $userset^i = \langle \{s, \langle IDsofSecret(s) \rangle \mid s \in Secrets^i \} \rangle$.

We also need a term that represents a dictionary that maps domains to (private) SSL keys of the IdP i . We define $sslkeys^i = \langle \{ \langle d, sslkey(d) \rangle \mid d \in dom(i) \} \rangle$.

Definition 43. A state $s \in Z^i$ of an IdP i is a term of the form $\langle nonces, sslkeys, users, signkey, sessions, corrupt \rangle$ where $nonces \subset^{\diamond} \mathcal{N}$ (used nonces), $sslkeys = sslkeys^i$, $users = userset^i$, $signkey \in \mathcal{N}$ (the key used by the IdP i to sign UCs), $sessions \in [\mathcal{N} \times \mathcal{T}_{\mathcal{N}}]$, $corrupt \in \mathcal{T}_{\mathcal{N}}$.

The initial state s_0^i of i is the state $\langle \langle \rangle, sslkeys^i, userset^i, signkey(i), \langle \rangle, \perp \rangle$.

The relation R^i that defines the behavior of the IdP i is defined as follows:

Algorithm 11 Relation of IdP R^i

Input: $(a:f:m), s$

- 1: **let** $s' := s$
- 2: **if** $s'.corrupt \neq \perp \vee m \equiv \text{CORRUPT}$ **then**
- 3: **let** $s'.corrupt := \langle \langle a, f, m \rangle, s'.corrupt \rangle$
- 4: **let** $m' \leftarrow d_{N^p}(s')$
- 5: **let** $a' \leftarrow \text{IPs}$
- 6: **stop** $\{ \langle a':a:m' \rangle \}, s'$
- 7: **end if**
- 8: **let** $sts := \langle \text{Strict-Transport-Security}, \top \rangle$
- 9: **let** $m_{dec}, k, k', inDomain$ **such that**
 - $\hookrightarrow \langle m_{dec}, k \rangle \equiv dec_a(m, k') \wedge \langle inDomain, k' \rangle \in s.sslkeys$
 - \hookrightarrow **if possible; otherwise stop** $\{ \}, s$
- 10: **let** $n, method, path, params, headers, body$ **such that**
 - $\hookrightarrow \langle \text{HTTPReq}, n, method, inDomain, path, params, headers, body \rangle \equiv m_{dec}$
 - \hookrightarrow **if possible; otherwise stop** $\{ \}, s$
- 11: **if** $method \equiv \text{POST}$ **then**
- 12: **if** $path \neq /certreq$ **then** ▷ User logs in.
- 13: **let** $id, secret$ **such that** $\langle id, secret \rangle \equiv body$
 - \hookrightarrow **if possible; otherwise stop** $\{ \}, s$
- 14: **if** $headers \neq \langle \text{Origin}, \langle inDomain, S \rangle \rangle$ **then**
- 15: **stop** $\{ \}, s$
- 16: **end if**
- 17: **let** $ids := s.users[secret]$
- 18: **if** $ids \equiv \langle \rangle \vee id \equiv \langle \rangle \vee id \notin^{\diamond} ids$ **then** ▷ Check id/secret pair.

```

19:         stop {}, s
20:     end if
21:     let sessionid, s' := TAKENONCE(s')
22:     let s'.sessions[sessionid] := ids
23:     let setCookie := ⟨Set-Cookie, ⟨⟨sessionid, sessionid, T, T, T⟩⟩⟩
24:     let m' := encs(⟨HTTPResp, n, 200, ⟨sts, setCookie⟩, T⟩, k)
25:     stop {(f:a:m')}, s'
26: else                                     ▷ User wants a certificate.
27:     let id, pubkey such that ⟨id, pubkey⟩ ≡ body
        ↪ if possible; otherwise stop {}, s
28:     let sessionid := headers[Cookie][sessionid]
29:     if id ∉⟨⟩ s'.sessions[sessionid] then   ▷ Check if user is logged in.
30:         stop {}, s
31:     end if
32:     let uc := sig(⟨id, pubkey⟩, s.signkey)
33:     let m' := encs(⟨HTTPResp, n, 200, ⟨sts, uc⟩, k)
34:     stop {(f:a:m')}, s'
35: end if
36: else
37:     if path ≡ /pif then
38:         let m' := encs(⟨HTTPResp, n, 200, ⟨sts⟩, ⟨script_idp_pif,
            ↪ ⟨init, ⟨⟩, ⟨⟩, ⟨⟩, ⊥, ⊥, ⊥⟩⟩), k)
39:     else
40:         let m' := encs(⟨HTTPResp, n, 200, ⟨sts⟩, ⟨script_idp_ad, ⟨⟩⟩), k)
41:     end if
42:     stop {(f:a:m')}, s'
43: end if
44: stop {}, s

```

F.12 BrowserID Scripts

As already mentioned in Section F.1, the set \mathcal{S} of the web system $\mathcal{BID}_{\text{primary}} = (\mathcal{W}, \mathcal{S}, \text{script}, E_0)$ consists of the scripts R^{att} , $\text{script}_{\text{rp_index}}$, $\text{script}_{\text{lpo_cif}}$, $\text{script}_{\text{lpo_ld}}$, $\text{script}_{\text{idp_pif}}$, and $\text{script}_{\text{idp_ad}}$ with their string representations being att_script , $\text{script}_{\text{rp_index}}$, $\text{script}_{\text{lpo_cif}}$, $\text{script}_{\text{lpo_ld}}$, $\text{script}_{\text{idp_pif}}$, and $\text{script}_{\text{idp_ad}}$ (defined by script).

The script R^{att} is the attacker script (see Section 2.3). The formal model of the other scripts follows the description in Appendix E. The script $\text{script}_{\text{rp_index}}$ defines the script of the RP index page. In reality, this page has its own script(s) and includes a script from LPO. In our model, we combine both scripts into $\text{script}_{\text{rp_index}}$. In particular, this script is responsible for creating the CIF and the LD iframes/subwindows, whose contents are loaded from LPO.

In what follows, the scripts $\text{script}_{\text{rp_index}}$, $\text{script}_{\text{lpo_cif}}$, and $\text{script}_{\text{lpo_ld}}$ are defined formally. First, we introduce some notation and helper functions.

Notations and Helper Functions. In the formal description of the scripts we use an abbreviation for URLs at LPO. We write $\text{URL}_{\text{path}}^{\text{LPO}}$ to describe the following URL term: $\langle \text{URL}, \mathcal{S}, \text{dom}(\text{LPO}), \text{path}, \langle \rangle \rangle$. Also, we call $\text{origin}_{\text{LPO}}$ the origin of LPO which describes the following origin term: $\langle \text{dom}(\text{LPO}), \mathcal{S} \rangle$.

In order to simplify the description of the scripts, several helper functions are used.

CHOOSEINPUT. As explained in Section 2.5, the state of a document contains a term, say, *scriptinputs*, which records the input this document has obtained so far (via XHRs and postMessages). If the script of the document is activated, it will typically need to pick one input message from *scriptinputs* and record which input it has already processed. For this purpose, the function $\text{CHOOSEINPUT}(s', \text{scriptinputs})$ is used, where s' denotes the script's current state. It saves the indexes of already handled messages in the scriptstate s' and chooses a yet unhandled input message from *scriptinputs*. The index of this message is then saved in the scriptstate (which is returned to the script).

Algorithm 12 Choose an unhandled input message for a script

```

1: function CHOOSEINPUT( $s', \text{scriptinputs}$ )
2:   let  $iid$  such that  $iid \in \{1, \dots, |\text{scriptinputs}|\} \wedge iid \notin s'.\text{handledInputs}$ 
    $\hookrightarrow$  if possible; otherwise return  $(\perp, s')$ 
3:   let  $input := \pi_{iid}(\text{scriptinputs})$ 
4:   let  $s'.\text{handledInputs} := s'.\text{handledInputs} + \diamond iid$ 
5:   return  $(input, s')$ 
6: end function

```

PARENTWINDOW. To determine the nonce referencing the parent window in the browser, the function $\text{PARENTWINDOW}(\text{tree}, \text{docnonce})$ is used. It takes the term *tree*, which is the (partly cleaned) tree of browser windows the script is able to see and the document nonce *docnonce*, which is the nonce referencing the current document the script is running in, as input. It outputs the nonce referencing the window which directly contains in its subwindows the window of the document referenced by *docnonce*. If there is no such window (which is the case if the script runs in a document of a top-level window), PARENTWINDOW returns \perp .

SUBWINDOWS. This function takes a term *tree* and a document nonce *docnonce* as input just as the function above. If *docnonce* is not a reference to a document contained in *tree*, then $\text{SUBWINDOWS}(\text{tree}, \text{docnonce})$ returns $\langle \rangle$. Otherwise, let $\langle \text{docnonce}, \text{origin}, \text{script}, \text{scriptstate}, \text{scriptinput}, \text{subwindows}, \text{active} \rangle$ denote the subterm of *tree* corresponding to the document referred to by *docnonce*. Then, $\text{SUBWINDOWS}(\text{tree}, \text{docnonce})$ returns *subwindows*.

AUXWINDOW. This function takes a term *tree* and a document nonce *docnonce* as input as above. From all window terms in *tree* that have the window containing the document identified by *docnonce* as their opener, it selects one non-deterministically and returns its nonce. If there is no such window, it returns the nonce of the window containing *docnonce*.

OPENERWINDOW. This function takes a term *tree* and a document nonce *docnonce* as input as above. It returns the window nonce of the opener window of the window that contains the document identified by *docnonce*. Recall that the nonce identifying the opener of each window is stored inside the window term. If no document with nonce *docnonce* is found in the tree *tree*, \diamond is returned.

GETWINDOW. This function takes a term *tree* and a document nonce *docnonce* as input as above. It returns the nonce of the window containing *docnonce*.

GETORIGIN. The function $\text{GETORIGIN}(tree, docnonce)$ extracts the origin of a document. It searches for the document with the identifier $docnonce$ in the (cleaned) tree $tree$ of the browser's windows and documents. It returns the origin o of the document. If no document with nonce $docnonce$ is found in the tree $tree$, \diamond is returned.

Web storage under LPO's origin. The web storage under the origin of LPO used by the scripts $script_lpo_cif$ and $script_lpo_ld$ (see below) is organized as follows:

The localStorage is a dictionary. There are two types of entries in this dictionary: Under the key `siteInfo`, a dictionary is stored which has origins as keys and IDs as values. An entry in this dictionary indicates that the user is logged in at the referenced origin with a certain ID. The second type of entry has a nonce as a key. The value is an email address (ID). This models the email address a user entered in the LD before being navigated away to the AD. The nonce is also stored in the sessionStorage (see below).

Example 8.

$$\begin{aligned} &\langle\langle\text{siteInfo}, \langle\langle\text{domain}_{\text{RP1}}, \text{S}\rangle, id_1\rangle, \\ &\quad \langle\langle\text{domain}_{\text{RP2}}, \text{S}\rangle, id_1\rangle, \\ &\quad \langle\langle\text{domain}_{\text{RP3}}, \text{S}\rangle, id_2\rangle\rangle\rangle, \\ &\langle n_1, id_1\rangle, \\ &\langle n_2, id_3\rangle \end{aligned}$$

This example shows a localStorage under the origin of LPO, indicating that the user is logged in at $\text{domain}_{\text{RP1}}$ and $\text{domain}_{\text{RP2}}$ with id_1 and at $\text{domain}_{\text{RP3}}$ with id_2 (using HTTPS). Further, the nonces n_1 and n_2 each refer to an email address which the user entered in the LD.

The sessionStorage is also a dictionary. It may only contain one key, `idpnonce`. Its value is a nonce (like n_1 or n_2 in the example above) which references the latest email address entry in the localStorage (see above).

login.persona.org Communication Iframe Script ($script_lpo_cif$). As defined in Section 2.3, a script is a relation that takes as input a term and a set of nonces it may use. It outputs a new term. As specified in Section 2.5 (Triggering the Script of a Document ($m = \text{TRIGGER}$, $action = 1$)) and formally specified in Algorithm 6, the input term is provided by the browser. It contains the current internal state of the script (which we call *scriptstate* in what follows) and additional information containing all browser state information the script has access to, such as the input the script has obtained so far via XHRs and postMessages, information about windows, etc. The browser expects the output term to have a specific form, as also specified in Section 2.5 and Algorithm 6. The output term contains, among other information, the new internal scriptstate.

As for $script_lpo_cif$, this script models the script run in the CIF, as sketched in Appendix E.

We first describe the structure of the internal scriptstate of the script $script_lpo_cif$.

Definition 44. A scriptstate s of $script_lpo_cif$ is a term of the form $\langle q, requestOrigin, loggedInUser, pause, context, key, uc, handledInputs, refXHRctx, PIFindex \rangle$ where $q \in \mathbb{S}$, $requestOrigin \in \text{Origins} \cup \{\perp\}$, $loggedInUser \in \text{ID} \cup \{\langle \rangle, \perp\}$, $pause \in \{\top, \perp\}$, $context \in \mathcal{T}_{\mathcal{N}}$, $key \in \mathcal{N} \cup \{\perp\}$, $uc \in \mathcal{T}_{\mathcal{N}}$, $handledInputs \subset^{\diamond} \mathbb{N}$, $refXHRctx \in \mathcal{N} \cup \{\perp\}$,

$PIFIndex \in \mathbb{N} \cup \{\perp\}$. The initial scriptstate $initState_{cif}$ of $script_lpo_cif$ is the state $\langle init, \perp, \perp, \perp, \perp, \perp, \perp, \perp, \langle \rangle, \perp, \perp \rangle$.

Before we provide the formal specification of the relation that defines the behavior of $script_lpo_cif$, we present an informal description. The behavior mainly depends on the state q the script is in.

$q = init$ This is the initial state. Its only transition handles no input and outputs a `postMessage` `cifready` to its parent window and transitions to `default`.

$q = default$ This is the state to which `script_lpo_cif` always returns to. This state handles all `postMessages` the CIF expects to receive from its parent window. If the `postMessage` received was sent from the parent window of the CIF, it behaves as follows, depending on the first element of the received `postMessage`:

postMessage `loaded` The script records the sender's origin of the received `postMessage` as the remote origin in the scriptstate if the scriptstate did not contain any information about the remote origin yet. Also, an ID, which represents the assumption of the sender on who it believes to be logged in, is saved in the scriptstate. If the `pause` flag in the scriptstate is \top it transitions to the state `default`. Otherwise, it is checked, if the current context in the scriptstate is \perp . If the check is true, the script transitions to the state `fetchContext`, or to the state `checkAndEmit` otherwise.

postMessage `dlgRun` The script sets the `pause` flag in the scriptstate to \top and transitions to `default`.

postMessage `dlgCmplt` The script sets the `pause` flag in the scriptstate to \perp . It then transitions to the state `fetchContext`.

postMessage `loggedInUser` This message has to contain an ID. This ID is saved in the scriptstate and then the script transitions to `default`.

postMessage `logout` The script removes the entry for the RP (recorded in the scriptstate) from the `localStorage` and then transitions to the state `sendLogout`. If no remote origin is set in the script's state, it is now set to the sender's origin of the received `postMessage`.

$q = fetchContext$ In this state, the script sends an XHR to LPO with a GET request to the path `/ctx` and then transitions to the state `receiveContext`.

$q = receiveContext$ In this state, the script expects an XHR response as input containing the session context. This context is saved as the current context in the scriptstate. The script transitions to `checkAndEmit`.

$q = checkAndEmit$ This state lets the script create the provisioning iframe and transition to `startPIF` iff (1) some email address is marked as logged in at RP in the `localStorage`, (2) if an email address is recorded in the current scriptstate, this email address differs from the one recorded in the `localStorage`, and (3) the user is marked as logged in in the current context. Otherwise, if the email address recorded in the current scriptstate is $\langle \rangle$, the script transitions to `default`, else it transitions to `sendLogout`.

$q = startPIF$ In this state, the script waits for a `postMessage` from the PIF containing a `ping` message. If such a message is received and the sender's window and origin match the PIF, the script sends a `pong` message back to the PIF and transitions to the state `runPIF`.

$q = \text{runPIF}$ This is the state in which `script_lpo_cif` interacts with the PIF. This state handles all `postMessages` the CIF expects to receive from the latest PIF (as recorded in `PIFindex` in its state). If the `postMessage` received was sent from the PIF's window and the PIF's origin, it behaves as follows, depending on the first element of the `postMessage`:

postMessage `beginProvisioning` The script responds with a `postMessage` to the PIF containing the email address of the identity which is to authenticate to the relying party (as recorded in the CIF's state).

postMessage `genKeyPair` The script creates a fresh key pair (i.e. the CIF chooses a fresh nonce) and sends the public key contained in a `postMessage` to the PIF.

postMessage `registerCertificate` The script stores the UC received in this `postMessage` in the CIF's state and transitions to the state `createCAPforRP`.

postMessage `raiseProvisioningFailure` This message indicates that no one is logged in. This is recorded in the CIF's state accordingly. The script transitions to the state `sendLogout` in which the CIF's parent window will be notified that no one is logged in.

$q = \text{createCAPforRP}$ In this state, the script creates an IA for the request origin (as recorded in the script's state), combines the IA with the UC to a CAP, and sends the CAP in a `postMessage` to its parent restricting the receiver to the request origin.

$q = \text{sendLogout}$ In this state, the script sends a `logout` `postMessage` to the parent document and then transitions to the `default` state.

We now specify the relation $\text{script_lpo_cif} \subseteq (\mathcal{T}_{\mathcal{N}} \times 2^{\mathcal{N}}) \times \mathcal{T}_{\mathcal{N}}$ of the CIF's scripting process formally. Just like in Appendix C.2, we describe this relation by a non-deterministic algorithm.

Just like all scripts, as explained in Section 2.5 (see also Algorithm 6 for the formal specification), the input term this script obtains from the browser contains the cleaned tree of the browser's windows and documents `tree`, the nonce of the current document `docnonce`, its own scriptstate `scriptstate` (as defined in Definition 44), a sequence of all inputs `scriptinput` (also containing already handled inputs), a dictionary `cookies` of all accessible cookies of the document's domain, the localStorage `localStorage` belonging to the document's origin, the secrets `secret` of the document's origin, and a set `nonces` of fresh nonces as input. The script returns a new scriptstate `s'`, a new set of cookies `cookies'`, a new localStorage `localStorage'`, and a term `command` denoting a command to the browser.

Algorithm 13 Relation of `script_lpo_cif`

Input: $\langle \text{tree}, \text{docnonce}, \text{scriptstate}, \text{scriptinputs}, \text{cookies}, \text{localStorage}, \text{sessionStorage}, \text{ids}, \text{secret} \rangle, \text{nonces}$

- 1: **let** `s'` := `scriptstate`
- 2: **let** `cookies'` := `cookies`
- 3: **let** `localStorage'` := `localStorage`
- 4: **let** `sessionStorage'` := `sessionStorage`
- 5: **switch** `s'.q` **do**
- 6: **case** `init`
- 7: **let** `command` := $\langle \text{POSTMESSAGE}, \text{PARENTWINDOW}(\text{tree}, \text{docnonce}), \text{cifready}, \langle \rangle, \perp \rangle$

```

8:     let s'.q := default
9:     stop ⟨s', cookies', localStorage', sessionStorage', command⟩
10:  case default
11:    let input, s' := CHOOSEINPUT(s', scriptinputs)
12:    if  $\pi_1(input) \equiv \text{POSTMESSAGE}$  then
13:      let senderWindow :=  $\pi_2(input)$ 
14:      let senderOrigin :=  $\pi_3(input)$ 
15:      let m :=  $\pi_4(input)$ 
16:      if senderWindow  $\equiv \text{PARENTWINDOW}(tree, docnonce)$  then
17:        switch m do
18:          case ⟨loaded, id⟩
19:            if s'.requestOrigin  $\equiv \perp$  then
20:              let s'.requestOrigin := senderOrigin
21:            end if
22:            let s'.loggedInUser := id
23:            if s'.pause  $\equiv \top$  then
24:              stop ⟨s', cookies', localStorage', sessionStorage', ⟨⟩⟩
25:            else if s'.context  $\equiv \perp$  then
26:              let s'.q := fetchContext
27:              stop ⟨s', cookies', localStorage', sessionStorage', ⟨⟩⟩
28:            else
29:              let s'.q := checkAndEmit
30:              stop ⟨s', cookies', localStorage', sessionStorage', ⟨⟩⟩
31:            end if
32:          case ⟨dlgRun, ⟨⟩⟩
33:            let s'.pause :=  $\top$ 
34:            stop ⟨s', cookies', localStorage', sessionStorage', ⟨⟩⟩
35:          case ⟨dlgCmplt, ⟨⟩⟩
36:            let s'.pause :=  $\perp$ 
37:            let s'.q := fetchContext
38:            stop ⟨s', cookies', localStorage', sessionStorage', ⟨⟩⟩
39:          case ⟨loggedInUser, id⟩
40:            let s'.loggedInUser := id
41:            stop ⟨s', cookies', localStorage', sessionStorage', ⟨⟩⟩
42:          case ⟨logout, ⟨⟩⟩
43:            if s'.requestOrigin  $\equiv \perp$  then
44:              let s'.requestOrigin := senderOrigin
45:            end if
46:            let s'.loggedInUser :=  $\perp$ 
47:            remove the element with key s'.requestOrigin
              ↪ from the dictionary localStorage'[siteInfo]
48:            let s'.q := sendLogout
49:          end if
50:        end if
51:      case fetchContext
52:        let s'.refXHRctx ← nonces
53:        let command := ⟨XMLHTTPREQUEST, URLLPOctx, GET, ⟨⟩, s'.refXHRctx⟩
54:        let s'.q := receiveContext

```

```

55:   stop ⟨s', cookies', localStorage', sessionStorage', command⟩
56:   case receiveContext
57:     let input, s' := CHOOSEINPUT(s', scriptinputs)
58:     if (π1(input) ≡ XMLHTTPREQUEST) ∧ (π3(input) ≡ s'.refXHRctx) then
59:       let s'.context := π2(input)
60:       let s'.q := checkAndEmit
61:       stop ⟨s', cookies', localStorage', sessionStorage', ⟨⟩⟩
62:     end if
63:   case checkAndEmit
64:     let s'.email := localStorage'[siteInfo][s'.requestOrigin]
65:     if (s'.email ≠ ⟨⟩)
66:       ↪ ∧(s'.loggedInUser ∉ {⟨⟩, ⊥} ⇒ (s'.loggedInUser ≠ s'.email))
67:       ↪ ∧(π1(s'.context) ≡ T) then
68:         let s'.q := startPIF
69:         let url := ⟨URL, S, π2(s'.email), /pif⟩
70:         let s'.PIFindex := |subwindows| + 1
71:         ▷ Index of the next subwindow to be created.
72:         let command := ⟨IFRAME, url, _SELF⟩
73:         stop ⟨s', cookies', localStorage', sessionStorage', command⟩
74:       else if s'.loggedInUser ≡ ⟨⟩ then
75:         let s'.q := default
76:         stop ⟨s', cookies', localStorage', sessionStorage', ⟨⟩⟩
77:       else
78:         let s'.q := sendLogout
79:         stop ⟨s', cookies', localStorage', sessionStorage', ⟨⟩⟩
80:       end if
81:   case startPIF
82:     let idpOrigin := ⟨π2(s'.email), S⟩
83:     let input, s' := CHOOSEINPUT(s', scriptinputs)
84:     let pifNonce := πs'.PIFindex(subwindows).nonce
85:     if π1(input) ≡ POSTMESSAGE then
86:       let senderWindow := π2(input)
87:       let senderOrigin := π3(input)
88:       let m := π4(input)
89:       if m ≡ ping ∧ senderWindow ≡ pifNonce
90:         ↪ ∧senderOrigin ≡ idpOrigin then
91:           let command := ⟨POSTMESSAGE, pifNonce, pong, idpOrigin⟩
92:           let s'.q := runPIF
93:           stop ⟨s', cookies', localStorage', sessionStorage', command⟩
94:         end if
95:       end if
96:   case runPIF
97:     let idpOrigin := ⟨π2(s'.email), S⟩
98:     let input, s' := CHOOSEINPUT(s', scriptinputs)
99:     let pifNonce := πs'.PIFindex(subwindows).nonce
100:    if π1(input) ≡ POSTMESSAGE then
101:      let senderWindow := π2(input)
102:      let senderOrigin := π3(input)
103:      let m := π4(input)

```



```

100:   if  $senderWindow \equiv pifNonce \wedge senderOrigin \equiv idpOrigin$  then
101:     switch  $\pi_1(m)$  do
102:       case beginProvisioning
103:         let  $jschannel\_nonce := \pi_2(m)$ 
104:         let  $command := \langle POSTMESSAGE, pifNonce,$ 
            $\hookrightarrow \langle jschannel\_nonce, s'.email \rangle, idpOrigin \rangle$ 
105:         stop  $\langle s', cookies', localStorage', sessionStorage', command \rangle$ 
106:       case genKeyPair
107:         let  $jschannel\_nonce := \pi_2(m)$ 
108:         let  $s'.key \leftarrow nonces$ 
109:         let  $command := \langle POSTMESSAGE, pifNonce,$ 
            $\hookrightarrow \langle jschannel\_nonce, pub(s'.key) \rangle, idpOrigin \rangle$ 
110:         stop  $\langle s', cookies', localStorage', sessionStorage', command \rangle$ 
111:       case registerCertificate
112:         if  $\pi_1(extractmsg(\pi_2(m))) \equiv s'.email \wedge s'.email \neq \langle \rangle$  then
            $\triangleright$  This check is our fix against identity injection.
113:           let  $s'.uc := \pi_2(m)$ 
114:           let  $s'.q := createCAPforRP$ 
115:           end if
116:           stop  $\langle s', cookies', localStorage', sessionStorage', command \rangle$ 
117:       case raiseProvisioningFailure
118:         let  $s'.loggedInUser := \perp$ 
119:         let  $s'.q := sendLogout$ 
120:         stop  $\langle s', cookies', localStorage', sessionStorage', command \rangle$ 
121:     end if
122:   end if
123: case createCAPforRP
124:   let  $ia := sig(s'.requestOrigin, s'.key)$ 
125:   let  $cap := \langle s'.uc, ia \rangle$ 
126:   let  $command := \langle POSTMESSAGE, PARENTWINDOW(tree, docnonce),$ 
      $\hookrightarrow \langle response, cap \rangle, s'.requestOrigin \rangle$ 
127:   let  $s'.q := null$ 
128:   stop  $\langle s', cookies', localStorage', sessionStorage', command \rangle$ 
129: case sendLogout
130:   let  $command := \langle POSTMESSAGE, PARENTWINDOW(tree, docnonce),$ 
      $\hookrightarrow \langle logout, \langle \rangle \rangle, \perp \rangle$ 
131:   let  $s'.q := default$ 
132:   stop  $\langle s', cookies', localStorage', sessionStorage', command \rangle$ 
133: stop  $\langle scriptstate, cookies, localStorage, sessionStorage, \langle \rangle \rangle$ 

```

login.persona.org Login Dialog Script (script_lpo_ld). This script models the LD contents. Its formal specification, presented next, follows the one presented above for *script_lpo_cif*.

Definition 45. A scriptstate s of *script_lpo_ld* is a term of the form $\langle q, requestOrigin, context, email, key, uc, handledInputs, refXHRctx, refXHRLOauth, PIFindex \rangle$ with $q \in \mathbb{S}$, $requestOrigin \in Origins \cup \{\perp\}$, $context \in \mathcal{T}_{\mathcal{N}}$, $email \in ID \cup \{\perp\}$, $key \in \mathcal{K} \cup \{\perp\}$,

$uc \in \mathcal{T}_{\mathcal{N}}$, $handledInputs \subset \mathbb{N}$, $refXHRctx, refXHRLPOauth \in \mathcal{N} \cup \{\perp\}$, $PIFindex \in \mathbb{N} \cup \{\perp\}$. The initial scriptstate $initState_{ld}$ is the state $\langle \text{init}, \perp, \perp, \perp, \perp, \perp, \langle \rangle, \perp, \perp, \perp \rangle$.

Before we provide the formal specification of the relation that defines the behavior of $script_lpo_ld$, we present an informal description. The behavior mainly depends on the state q the script is in.

$q \equiv \text{init}$ This is the initial state. Its only transition takes no input and outputs a postMessage `ldready` to its parent window and transitions to `start`.

$q \equiv \text{start}$ In this state, the script expects a request postMessage. The sender's origin of this postMessage is recorded as the requesting origin in the scriptstate. An XHR is sent to LPO with a GET request to the path `/ctx` and then the script transitions to the state `receiveContext`.

$q \equiv \text{receiveContext}$ In this state, the script expects an XHR response as input containing the session context. This context is saved as the current context in the scriptstate. The script checks if an `idpNonce` is recorded in the `sessionStorage`. The presence of this nonce indicates that there was a run of `script_lpo_ld` in the same window previously. Indexed by this nonce, there can be an email address (identity) recorded in the `localStorage` which is then copied to the script's state. Otherwise an email address is non-deterministically chosen (and copied to the script's state) out of the email addresses owned by the browser.

The script now always issues the command to create an iframe, the PIF. The URL for the PIF is determined by the domain of the email address now recorded in the state. The script then transitions to the state `startPIF`.

$q = \text{startPIF}$ In this state, the script waits for a postMessage from the PIF containing a ping message. If such a message is received and the sender's window and origin match the PIF, the script sends a pong message back to the PIF and transitions to the state `runPIF`.

$q = \text{runPIF}$ This is the state in which `script_lpo_ld` interacts with the PIF. This state handles all postMessages the LD expects to receive from the latest PIF (as recorded in `PIFindex` in its state). If the postMessage received was sent from the PIF's window and the PIF's origin, it behaves as follows, depending on the first element of the received postMessage:

postMessage beginProvisioning The script responds with a postMessage to the PIF containing the email address of the identity which is to authenticate to the relying party (as recorded in the LD's state).

postMessage genKeyPair The script creates a fresh key pair (i.e. the LD chooses a fresh nonce) and sends the public key contained in an postMessage to the PIF.

postMessage registerCertificate The script stores the UC received in this postMessage in the LD's state. If the context contained in the script's state indicates that the browser is authenticated to LPO, the script transitions to the state `createCAPforRP`. Otherwise, the script transitions to the state `createCAPforLPO`.

postMessage raiseProvisioningFailure This message indicates that no one is logged in. The script now chooses a fresh nonce, the so-called *idpNonce*, which is stored in the `sessionStorage`. In the `localStorage`, this nonce is used as

a key under which the email address is stored, the LD is currently trying to get an UC for. The script navigates the window it is running to the authentication path at the identity provider responsible for the email address.

- $q = \text{createCAPforLD}$ In this state, the script creates an IA for LPO, combines it with the UC (stored in the script's state) to a CAP and sends the CAP to LPO in an XHR. The nonce identifying the XHR is stored as `refXHRLP0auth` in the script's state.
- $q = \text{receiveLP0authresponse}$ In this state, the script expects the response to the XHR identified by the nonce `refXHRLP0auth`. If the response indicates a successful authentication at LPO, the context recorded in the script's state is changed accordingly and the script transitions to the state `createCAPforRP`.
- $q = \text{createCAPforRP}$ In this state, the script creates an IA for the request origin (as recorded in the script's state), combines the IA with the UC to a CAP, and sends the CAP in a `postMessage` to its parent restricting the receiver to the request origin. The script records in the `localStorage` that the email address it is currently using is logged in at the request origin. The script then transitions to the state `null`.
- $q \equiv \text{null}$ In this state, the script does nothing.

We now formally specify the relation $\text{script_lpo_ld} \subseteq (\mathcal{T}_{\mathcal{N}} \times 2^{\mathcal{N}}) \times \mathcal{T}_{\mathcal{N}}$ of the LD's scripting process. Just like in Appendix C.2, we describe this relation by a non-deterministic algorithm. Like all scripts, the input term given to this script is determined by the browser and the browser expects a term of a specific form (see Algorithm 6)

Algorithm 14 Relation of script_lpo_ld

Input: $\langle \text{tree}, \text{docnonce}, \text{scriptstate}, \text{scriptinputs}, \text{cookies}, \text{localStorage}, \text{sessionStorage}, \text{ids}, \text{secret} \rangle, \text{nonces}$

- 1: **let** $s' := \text{scriptstate}$
- 2: **let** $\text{cookies}' := \text{cookies}$
- 3: **let** $\text{localStorage}' := \text{localStorage}$
- 4: **let** $\text{sessionStorage}' := \text{sessionStorage}$
- 5: **switch** $s'.q$ **do**
- 6: **case** `init`
- 7: **let** $\text{command} := \langle \text{POSTMESSAGE}, \text{OPENERWINDOW}(\text{tree}, \text{docnonce}), \hookrightarrow \langle \text{ldready}, \langle \rangle \rangle, \perp \rangle$
- 8: **let** $s'.q := \text{start}$
- 9: **stop** $\langle s', \text{cookies}', \text{localStorage}', \text{sessionStorage}', \text{command} \rangle$
- 10: **case** `start`
- 11: **let** $\text{input}, s' := \text{CHOOSEINPUT}(s', \text{scriptinputs})$
- 12: **if** $\pi_1(\text{input}) \equiv \text{POSTMESSAGE}$ **then**
- 13: **let** $\text{senderWindow} := \pi_2(\text{input})$
- 14: **let** $\text{senderOrigin} := \pi_3(\text{input})$
- 15: **let** $m := \pi_4(\text{input})$
- 16: **if** $m \equiv \langle \text{request}, \langle \rangle \rangle$ **then**
- 17: **let** $s'.\text{requestOrigin} := \text{senderOrigin}$
- 18: **let** $s'.\text{refXHRctx} \leftarrow \text{nonces}$
- 19: **let** $\text{command} := \langle \text{XMLHTTPREQUEST}, \text{URL}_{/\text{ctx}}^{\text{LPO}}, \text{GET}, \langle \rangle, s'.\text{refXHRctx} \rangle$
- 20: **let** $s'.q := \text{receiveContext}$
- 21: **stop** $\langle s', \text{cookies}', \text{localStorage}', \text{sessionStorage}', \text{command} \rangle$

```

22:         end if
23:     end if
24:     case receiveContext
25:         let input, s' := CHOOSEINPUT(s', scriptinputs)
26:         if ( $\pi_1(input) \equiv \text{XMLHTTPREQUEST}$ )  $\wedge$  ( $\pi_3(input) \equiv s'.\text{refXHRctx}$ ) then
27:             let s'.context :=  $\pi_2(input)$ 
28:             let s'.q := startPIF
29:             let idpnonce := sessionStorage[idpnonce]
30:             if idpnonce  $\equiv$   $\langle \rangle$   $\vee$  localStorage[idpnonce]  $\equiv$   $\langle \rangle$  then
31:                 let s'.email  $\leftarrow$  ids
32:             else
33:                 let s'.email := localStorage[idpnonce]
34:                 let sessionStorage[idpnonce] :=  $\langle \rangle$ 
35:             end if
36:             let url :=  $\langle \text{URL}, S, \pi_2(s'.\text{email}), /pif \rangle$ 
37:             let s'.PIFindex := |subwindows| + 1
                                      $\triangleright$  Index of the next subwindow to be created.
38:             let command :=  $\langle \text{IFRAME}, url, \_SELF \rangle$ 
39:             stop  $\langle s', \text{cookies}', \text{localStorage}', \text{sessionStorage}', \text{command} \rangle$ 
40:         end if
41:     case startPIF
42:         let idpOrigin :=  $\langle \pi_2(s'.\text{email}), S \rangle$ 
43:         let input, s' := CHOOSEINPUT(s', scriptinputs)
44:         let pifNonce :=  $\pi_{s'.\text{PIFindex}}(\text{subwindows}).\text{nonce}$ 
45:         if  $\pi_1(input) \equiv \text{POSTMESSAGE}$  then
46:             let senderWindow :=  $\pi_2(input)$ 
47:             let senderOrigin :=  $\pi_3(input)$ 
48:             let m :=  $\pi_4(input)$ 
49:             if  $m \equiv \text{ping} \wedge \text{senderWindow} \equiv \text{pifNonce}$ 
                $\hookrightarrow \wedge \text{senderOrigin} \equiv \text{idpOrigin}$  then
50:                 let command :=  $\langle \text{POSTMESSAGE}, \text{pifNonce}, \text{pong}, \text{idpOrigin} \rangle$ 
51:                 let s'.q := runPIF
52:                 stop  $\langle s', \text{cookies}', \text{localStorage}', \text{sessionStorage}', \text{command} \rangle$ 
53:             end if
54:         end if
55:     case runPIF
56:         let idpOrigin :=  $\langle \pi_2(s'.\text{email}), S \rangle$ 
57:         let input, s' := CHOOSEINPUT(s', scriptinputs)
58:         let pifNonce :=  $\pi_{s'.\text{PIFindex}}(\text{subwindows}).\text{nonce}$ 
59:         if  $\pi_1(input) \equiv \text{POSTMESSAGE}$  then
60:             let senderWindow :=  $\pi_2(input)$ 
61:             let senderOrigin :=  $\pi_3(input)$ 
62:             let m :=  $\pi_4(input)$ 
63:             if  $\text{senderWindow} \equiv \text{pifNonce} \wedge \text{senderOrigin} \equiv \text{idpOrigin}$  then
64:                 switch  $\pi_1(m)$  do
65:                     case beginProvisioning
66:                         let jschannel_nonce :=  $\pi_2(m)$ 
67:                         let command :=  $\langle \text{POSTMESSAGE}, \text{pifNonce},$ 
                $\hookrightarrow \langle \text{jschannel\_nonce}, s'.\text{email} \rangle, \text{idpOrigin} \rangle$ 

```

```

68:         stop ⟨s', cookies', localStorage', sessionStorage', command⟩
69:     case genKeyPair
70:         let jschannel_nonce := π2(m)
71:         let s'.key ← nonces
72:         let command := ⟨POSTMESSAGE, pifNonce,
73:             ↪ ⟨jschannel_nonce, pub(s'.key)⟩, idpOrigin⟩
74:         stop ⟨s', cookies', localStorage', sessionStorage', command⟩
75:     case registerCertificate
76:         if π1(extractmsg(π2(m))) ≡ s'.email ∧ s'.email ≠ ⟨⟩ then
77:             ▷ This check is our fix against identity injection.
78:             let s'.uc := π2(m)
79:             let loggedIn := π1(s'.context)
80:             if loggedIn ≡ ⊤ then
81:                 let s'.q := createCAPforRP
82:             end if
83:             let s'.q := createCAPforLPO
84:         end if
85:         stop ⟨s', cookies', localStorage', sessionStorage', command⟩
86:     case raiseProvisioningFailure
87:         let idpnonce ← nonces
88:         let localStorage'[idpnonce] := s'.email
89:         let sessionStorage'[idpnonce] := idpnonce
90:         let command := ⟨HREF, ⟨URL, S, π2(s'.email), ⟨⟩⟩, _SELF⟩
91:         stop ⟨s', cookies', localStorage', sessionStorage', command⟩
92:     end if
93:     case createCAPforLPO
94:         let ia := sig(⟨dom(LPO), S⟩, s'.key)
95:         let cap := ⟨s'.uc, ia⟩
96:         let body := ⟨cap, π2(s'.context)⟩
97:         let s'.refXHRLPOauth ← nonces
98:         let command := ⟨XMLHTTPREQUEST, URLLPO/auth, POST, body, s'.refXHRLPOauth⟩
99:         let s'.q := receiveLPOauthresponse
100:     stop ⟨s', cookies', localStorage', sessionStorage', command⟩
101:     case receiveLPOauthresponse
102:         let input, s' := CHOOSEINPUT(s', scriptinputs)
103:         if (π1(input) ≡ XMLHTTPREQUEST) ∧ (π3(input) ≡ s'.refXHRLPOauth)
104:             ↪ ∧ π2(input) ≡ ⊤ then
105:                 let π1(s'.context) := ⊤
106:                 let s'.q := createCAPforRP
107:                 stop ⟨s', cookies', localStorage', sessionStorage', command⟩
108:             end if
109:     case createCAPforRP
110:         let ia := sig(s'.requestOrigin, s'.key)
111:         let cap := ⟨s'.uc, ia⟩
112:         let command := ⟨POSTMESSAGE, OPENERWINDOW(tree, docnonce),
113:             ↪ ⟨response, cap⟩, s'.requestOrigin⟩
114:         let s'.q := null

```

112: **let** *localStorage'*[siteInfo][*s'.requestOrigin*] := *s'.email*
113: **stop** $\langle s', cookies', localStorage', sessionStorage', command \rangle$
114: **stop** $\langle scriptstate, cookies, localStorage, sessionStorage, \langle \rangle \rangle$

Relying Party Web Page Script (script_rp_index). This script models the default web page at a RP. The user usually triggers the login process on this page. Its formal specification, presented next, follows the one presented for the other scripts above.

Definition 46. A scriptstate *s* of *script_rp_index* is a term of the form $\langle q, CIFindex, LDindex, dialogRunning, cap, handledInputs, refXHRcap \rangle$ with $q \in \mathbb{S}$, $CIFindex \in \mathbb{N} \cup \{\perp\}$, $dialogRunning \in \{\top, \perp\}$, $cap \in \mathcal{T}_{\mathcal{N}}$, $handledInputs \subset \langle \rangle \mathbb{N}$, $refXHRcap \in \mathcal{N} \cup \{\perp\}$. We call *s* the initial scriptstate of *script_rp_index* iff $s \equiv \langle \text{init}, \perp, \perp, \perp, \langle \rangle, \langle \rangle, \perp \rangle$.

Before we provide the formal specification of the relation that defines the behavior of *script_rp_index*, we present an informal description. The behavior mainly depends on the state *q* the script is in.

$q \equiv \text{init}$ This is the initial state. The script creates the CIF iframe and then transitions to `receiveCIFReady`.

$q \equiv \text{receiveCIFReady}$ In this state, the script expects a `cifready` `postMessage` from the CIF iframe with the sender origin of LPO. The script chooses some ID, $\langle \rangle$, or \perp and sends this in a `loaded` `postMessage` to the CIF iframe with receiver's origin set to the origin of LPO.¹⁴ It then transitions to the state `default`.

$q \equiv \text{default}$ In this state, the script chooses non-deterministically between (1) opening the LD subwindow and then transitioning to the same state or (2) handling one of the following `postMessages` (identified by their first element):

postMessage login This message has to be sent from the CIF with origin of LPO. Handling this `postMessage` stores the CAP (contained in the `postMessage`) in the scriptstate and then transitions to the `sendCAP` state.

postMessage logout This message has to be sent from the CIF with origin of LPO. Handling this `postMessage` has no effect and results in the same state.

postMessage ldready This message can only be handled after the LD has been opened and before a `response` `postMessage` has been received. The `ldready` `postMessage` has to be sent from the origin of LPO. The script sends a `request` `postMessage` to the LD and stays in the `default` state.

postMessage response This message can only be handled after the LD has been opened and before another `response` `postMessage` has been received. The `ldready` `postMessage` has to be sent from the origin of LPO. Handling this `postMessage` stores the CAP (contained in the `postMessage`) in the scriptstate, closes the LD, and then transitions to the `dlgClosed` state.

$q \equiv \text{dlgClosed}$ In this state, the script sends a `loggedInUser` `postMessage` to the CIF and transitions to the `loggedInUser` state.

¹⁴From the point of view of the real scripts running at RP either some ID is considered to be logged in (e.g. from some former "session"), or no one is considered to be logged in ($\langle \rangle$), or the script *script_rp_index* does not know if it should consider anyone to be logged in (\perp). This is overapproximated here by allowing *script_rp_index* to choose non-deterministically between these cases.

$q \equiv \text{loggedInUser}$ In this state, the script sends a `dlgCmplIt postMessage` to the CIF and transitions to the `sendCAP` state.

$q \equiv \text{sendCAP}$ In this state, the script sends the CAP to RP as a POST XHR and then transitions to the `receiveServiceToken` state.

$q \equiv \text{receiveServiceToken}$ In this state, the script receives $\langle n, i \rangle$ from RP, but does not do anything with it. The script then transitions to the `default` state.

We now formally specify the relation $\text{script_rp_index} \subseteq (\mathcal{T}_{\mathcal{N}} \times 2^{\mathcal{N}}) \times \mathcal{T}_{\mathcal{N}}$ of the RP-Doc's scripting process. Just like in Appendix C.2, we describe this relation by a non-deterministic algorithm. Like all scripts, the input term given to this script is determined by the browser and the browser expects a term of a specific form (see Algorithm 6). Following Algorithm 15, we provide some more explanation.

Algorithm 15 Relation of script_rp_index

Input: $\langle \text{tree}, \text{docnonce}, \text{scriptstate}, \text{scriptinputs}, \text{cookies}, \text{localStorage}, \text{sessionStorage}, \hookrightarrow \text{ids}, \text{secret} \rangle, \text{nonces}$

```

1: let  $s' := \text{scriptstate}$ 
2: let  $\text{cookies}' := \text{cookies}$ 
3: let  $\text{localStorage}' := \text{localStorage}$ 
4: let  $\text{sessionStorage}' := \text{sessionStorage}$ 
5: switch  $s'.q$  do
6:   case init
7:     let  $\text{command} := \langle \text{IFRAME}, \text{URL}_{\text{cif}}^{\text{LPO}}, \text{GETWINDOW}(\text{tree}, \text{docnonce}) \rangle$ 
8:     let  $s'.q := \text{receiveCIFReady}$ 
9:     let  $\text{subwindows} := \text{SUBWINDOWS}(\text{tree}, \text{docnonce})$ 
10:    let  $s'.\text{CIFindex} := |\text{subwindows}| + 1$   $\triangleright$  Index of the next subwindow to be created.
11:    stop  $\langle s', \text{cookies}', \text{localStorage}', \text{sessionStorage}', \text{command} \rangle$ 
12:   case receiveCIFReady
13:     let  $\text{input}, s' := \text{CHOOSEINPUT}(s', \text{scriptinputs})$ 
14:     if  $\pi_1(\text{input}) \equiv \text{POSTMESSAGE}$  then
15:       let  $\text{senderWindow} := \pi_2(\text{input})$ 
16:       let  $\text{senderOrigin} := \pi_3(\text{input})$ 
17:       let  $m := \pi_4(\text{input})$ 
18:       let  $\text{subwindows} := \text{SUBWINDOWS}(\text{tree}, \text{docnonce})$ 
19:       if  $(m \equiv \langle \text{cifready}, \langle \rangle \rangle)$ 
20:          $\hookrightarrow \wedge(\text{senderOrigin} \equiv \text{origin}_{\text{LPO}})$ 
21:          $\hookrightarrow \wedge(\text{senderWindow} \equiv \pi_{s'.\text{CIFindex}}(\text{subwindows}).\text{nonce})$  then
22:           let  $\text{id} \leftarrow \{\perp, \langle \rangle\} \cup \text{ID}$ 
23:           let  $\text{command} := \langle \text{POSTMESSAGE}, \pi_{s'.\text{CIFindex}}(\text{subwindows}), \hookrightarrow \langle \text{loaded}, \text{id} \rangle, \text{origin}_{\text{LPO}} \rangle$ 
24:           let  $s'.q := \text{default}$ 
25:           stop  $\langle s', \text{cookies}', \text{localStorage}', \text{sessionStorage}', \text{command} \rangle$ 
26:         end if
27:       end if
28:     case default
29:       if  $s'.\text{dialogRunning} \equiv \perp$  then
30:         let  $\text{choice} \leftarrow \{\text{openLD}, \text{handlePM}\}$ 
31:       else

```

```

30:     let choice := handlePM
31:   end if
32:   if choice ≡ openLD then
33:     let s'.dialogRunning := ⊤
34:     let command := ⟨HREF, URLLPO/ld, _BLANK⟩
35:     let s'.q := default
36:     stop ⟨s', cookies', localStorage', sessionStorage', command⟩
37:   else
38:     let input, s' := CHOOSEINPUT(s', scriptinputs)
39:     if π1(input) ≡ POSTMESSAGE then
40:       let senderWindow := π2(input)
41:       let senderOrigin := π3(input)
42:       let m := π4(input)
43:       let subwindows := SUBWINDOWS(tree, docnonce)
44:       if senderOrigin ≡ originLPO then
45:         if senderWindow ≡ πs', CIFindex(subwindows).nonce then
46:           if π1(m) ≡ login then
47:             let s'.cap := π2(m)
48:             let s'.q := sendCAP
49:             stop ⟨s', cookies', localStorage', sessionStorage', ⟨⟩⟩
50:           else if π1(m) ≡ logout then
51:             let s'.q := default
52:             stop ⟨s', cookies', localStorage', sessionStorage', ⟨⟩⟩
53:           end if
54:         else if s'.dialogRunning ≡ ⊤ then
55:           if π1(m) ≡ ldready then
56:             let command := ⟨POSTMESSAGE,
57:               ↪ AUXWINDOW(tree, docnonce), ⟨request, ⟨⟩⟩, originLPO⟩
58:             let s'.q := default
59:             stop ⟨s', cookies', localStorage', sessionStorage', command⟩
60:           else if π1(m) ≡ response then
61:             let s'.dialogRunning := ⊥
62:             let s'.cap := π2(m)
63:             let command := ⟨CLOSE, AUXWINDOW(tree, docnonce)⟩
64:             let s'.q := dlgClosed
65:             stop ⟨s', cookies', localStorage', sessionStorage', command⟩
66:           end if
67:         end if
68:       end if
69:     end if
70:   case dlgClosed
71:     let subwindows := SUBWINDOWS(tree, docnonce)
72:     let id := π1(extractmsg(π1(s'.cap))) ▷ Extract ID from CAP.
73:     let command := ⟨POSTMESSAGE, πs', CIFindex(subwindows).nonce,
74:       ↪ ⟨loggedInUser, id⟩, originLPO⟩
75:     let s'.q := loggedInUser
76:     stop ⟨s', cookies', localStorage', sessionStorage', command⟩
77:   case loggedInUser

```



```

77:   let subwindows := SUBWINDOWS(tree, docnonce)
78:   let command :=
    ↪ ⟨POSTMESSAGE,  $\pi_{s'}$ .CIFindex(subwindows).nonce, ⟨dlgCmplt, ⟨⟩⟩, originLPO)
79:   let s'.q := sendCAP
80:   stop ⟨s', cookies', localStorage', sessionStorage', command⟩
81:   case sendCAP
82:     let s'.refXHRcap ← nonces
83:     let host, protocol such that
      ↪ ⟨host, protocol⟩ ≡ GETORIGIN(tree, docnonce)
      ↪ if possible; otherwise stop
      ↪ ⟨scriptstate, cookies, localStorage, sessionStorage, command⟩
84:     let command := ⟨XMLHTTPREQUEST, ⟨URL, protocol, host, /, ⟨⟩⟩, POST, s'.cap,
      ↪ s'.refXHRcap⟩ ▷ Relay received CAP to RP.
85:     let s'.q := receiveServiceToken
86:     stop ⟨s', cookies', localStorage', sessionStorage', command⟩
87:     case receiveServiceToken
88:       let input, s' := CHOOSEINPUT(s', scriptinputs)
89:       if ( $\pi_1$ (input) ≡ XMLHTTPREQUEST) ∧ ( $\pi_3$ (input) ≡ s'.refXHRcap) then
90:         let s'.q := default
91:         stop ⟨s', cookies', localStorage', sessionStorage', ⟨⟩⟩
92:       end if
93:   stop ⟨scriptstate, cookies, localStorage, sessionStorage, ⟨⟩⟩

```

In Lines 7–11 and 33–36 the script asks the browser to create iframes. To obtain the window reference for these iframes, the script first determines the current number of subwindows and stores it (incremented by 1) in the scriptstate (CIFindex and LDindex, respectively). When the script is invoked the next time, the iframe the script asked to be created will have been added to the sequence of subwindows by the browser directly following the previously existing subwindows. The script can therefore access the iframe by the indexes CIFindex and LDindex, respectively.

Identity Provider Authentication Dialog Script (script_idp_ad). This script runs in the LD after *script_lpo_ld* has navigated the LD window. The purpose of this script is to authenticate the browser to the identity provider.

The script non-deterministically chooses if it sends authentication data to the IdP (i.e. its origin) via an XHR, or if it navigates the window to an URL at LPO which servers *script_lpo_ld*. Note that *script_idp_ad* does not read or change its scriptstate. Hence, we omit the definition of the scriptstate for this script.

Algorithm 16 Relation of *script_idp_ad*

Input: ⟨*tree*, *docnonce*, *scriptstate*, *scriptinputs*, *cookies*, *localStorage*, *sessionStorage*,
 ↪ *ids*, *secret*⟩, *nonces*

- 1: **let** *action* ← {authenticate, navigate}
- 2: **if** *action* ≡ authenticate **then**
- 3: **let** *email* ← *ids*
- 4: **let** *body* := ⟨*email*, *secret*⟩
- 5: **let** *host*, *protocol* **such that**
 ↪ ⟨*host*, *protocol*⟩ ≡ GETORIGIN(*tree*, *docnonce*)

```

    ↪ if possible; otherwise
    ↪ stop  $\langle \text{scriptstate}, \text{cookies}, \text{localStorage}, \text{sessionStorage}, \langle \rangle \rangle$ 
6: let  $\text{command} := \langle \text{XMLHTTPREQUEST}, \langle \text{URL}, \text{protocol}, \text{host}, / \text{auth}, \langle \rangle \rangle, \text{POST}, \text{body}, \perp \rangle$ 
7: stop  $\langle \text{scriptstate}, \text{cookies}', \text{localStorage}', \text{sessionStorage}', \text{command} \rangle$ 
8: else
9: let  $\text{command} := \langle \text{HREF}, \langle \text{URL}, \text{S}, \text{dom}(\text{LPO}), / \text{ld}, \langle \rangle \rangle, \text{\_SELF} \rangle$ 
10: stop  $\langle \text{scriptstate}, \text{cookies}', \text{localStorage}', \text{sessionStorage}', \text{command} \rangle$ 
11: end if

```

Identity Provider Provisioning Iframe Script (script_idp_pif). This script acts as a proxy between the LD or CIF and the IdP server.

Definition 47. A scriptstate s of *script_idp_pif* is a term of the form $\langle q, \text{emails}, \text{pubkeys}, \text{ucs}, \text{provisioningnonces}, \text{genkeypairnonces}, \text{xhrnonces}, \text{handledInputs} \rangle$ with $q \in \mathbb{S}$, $\text{emails}, \text{pubkeys}, \text{ucs} \in \mathcal{T}_{\mathcal{N}}$, $\text{provisioningnonces}, \text{genkeypairnonces}, \text{xhrnonces} \in \mathcal{N} \cup \{\perp\}$, $\text{handledInputs} \subset \langle \rangle \mathbb{N}$. We call s the initial scriptstate of *script_idp_pif* iff $s \equiv \langle \text{init}, \langle \rangle, \langle \rangle, \langle \rangle, \perp, \perp, \perp \rangle$.

Before we provide the formal specification of the relation that defines the behavior of *script_idp_pif*, we present an informal description. The behavior mainly depends on the state q the script is in.

$q = \text{init}$ This is the initial state. Its only transition handles no input and outputs a `postMessage ping` to its parent window, which has to have the origin of LPO, and transitions to `waiting`.

$q = \text{waiting}$ In this state, the script expects a `postMessage` containing either `ping` or `pong`, which has to be sent by the parent window from the origin of LPO. If such a `postMessage` has been received, the script transitions to `default`.

$q = \text{default}$ In this state, the script chooses an action non-deterministically out of the following:

`beginprovisioning` The script sends a `postMessage` to the parent window, which has to have the origin of LPO, indicating that the provisioning process of a UC should start. A fresh nonce is chosen, stored in the script's state, and included in this `postMessage`. The `postMessage` requests the email address of the user from the receiver. The address is to be sent to the PIF in a `postMessage` which is identified by the nonce in the request.

`genkeypair` The script sends a `postMessage` to the parent window, which has to have the origin of LPO, indicating that a new key pair should be generated. This `postMessage` requests the public key of this fresh key pair. As above, a nonce is included to identify the response corresponding to the request.

`registercert` The script sends a `postMessage` containing a UC to the parent window, which has to have the origin of LPO. This `postMessage` is only sent if the script has received a UC before.

`raisefailure` The script sends a `postMessage` to the parent window, which has to have the origin of LPO, indicating that the browser is currently not authenticated to the identity provider.

`requestuc` The script sends an XHR to the origin of the current document if the scriptstate contains at least one email address and one public key. The message

contains a non-deterministically chosen email address and a public key (from the scriptstate). The nonce identifying this XHR is non-deterministically chosen and stored in the scriptstate.

handleresponse The script chooses non-deterministically a script input and distinguishes if this input is a postMessage or an XHR response.

If the chosen input is a postMessage, it is checked if the postMessage was sent by the parent window and if this window has the origin of LPO. If this check is successful, it is checked if the message contains a nonce, which was previously been recorded in the script's state. If this nonce indicates that this message is a response to a beginProvisioning postMessage, the second part is assumed to contain an email address. This address is then recorded in the script's state. If the nonce indicates that this message is a response to a genKeyPair postMessage, the second part is assumed to contain a public key. This public key is then recorded in the script's state.

If the chosen input is an XHR response, it is checked if the nonce identifying the XHR is recorded in the script's state. If this is the case, the message is assumed to contain an UC. The content of the message is stored in the script's state.

Algorithm 17 Relation of *script_idp_pif*

Input: $\langle tree, docnonce, scriptstate, scriptinputs, cookies, localStorage, sessionStorage, \hookrightarrow ids, secret \rangle, nonces$

```

1: let  $s' := scriptstate$ 
2: switch  $s'.q$  do
3:   case init
4:     let  $command := \langle POSTMESSAGE, PARENTWINDOW(tree, docnonce), \hookrightarrow \langle ping, \langle \rangle \rangle, \langle dom(LPO), S \rangle \rangle$ 
5:     let  $s'.q := waiting$ 
6:     stop  $\langle s', cookies, localStorage, sessionStorage, command \rangle$ 
7:   case waiting
8:     let  $input, s' := CHOOSEINPUT(s', scriptinputs)$ 
9:     let  $senderWindow := \pi_2(input)$ 
10:    let  $senderOrigin := \pi_3(input)$ 
11:    let  $m := \pi_4(input)$ 
12:    if  $\pi_1(input) \in \{ping, pong\}$ 
13:       $\hookrightarrow \wedge senderWindow \equiv PARENTWINDOW(tree, docnonce)$ 
14:       $\hookrightarrow \wedge senderOrigin \equiv \langle dom(LPO), S \rangle$  then
15:        let  $s'.q := default$ 
16:      end if
17:    stop  $\langle s', cookies, localStorage, sessionStorage, \langle \rangle \rangle$ 
18:   case default
19:     let  $action \leftarrow \{beginprovisioning, genkeypair, registercert, \hookrightarrow raisefailure, requestuc, handleresponse\}$ 
20:     switch  $action$  do
21:       case beginprovisioning
22:         let  $jschannel\_nonce \leftarrow nonces$ 

```

```

21:         let command := ⟨POSTMESSAGE, PARENTWINDOW(tree, docnonce),
    ↪ ⟨beginProvisioning, jschannel_nonce⟩, dom(LPO)⟩
22:         let s'.provisioningnonces :=
    ↪ s'.provisioningnonces + ∅ jschannel_nonce
23:         stop ⟨s', cookies, localStorage, sessionStorage, command⟩
24:     case genkeypair
25:         let jschannel_nonce ← nonces
26:         let command := ⟨POSTMESSAGE, PARENTWINDOW(tree, docnonce),
    ↪ ⟨genKeyPair, jschannel_nonce⟩, dom(LPO)⟩
27:         let s'.genkeypairnonces :=
    ↪ s'.genkeypairnonces + ∅ jschannel_nonce
28:         stop ⟨s', cookies, localStorage, sessionStorage, command⟩
29:     case registercert
30:         if s'.ucs ≠ ⟨⟩ then
31:             let uc ← s'.ucs
32:             let command := ⟨POSTMESSAGE, PARENTWINDOW(tree, docnonce),
    ↪ ⟨registerCertificate, uc⟩, dom(LPO)⟩
33:             stop ⟨s', cookies, localStorage, sessionStorage, command⟩
34:         end if
35:     case raisefailure
36:         let command := ⟨POSTMESSAGE, PARENTWINDOW(tree, docnonce),
    ↪ ⟨raiseProvisioningFailure, ⊥⟩, dom(LPO)⟩
37:         stop ⟨s', cookies, localStorage, sessionStorage, command⟩
38:     case requestuc
39:         if s'.emails ≠ ⟨⟩ ∧ s'.pubkeys ≠ ⟨⟩ then
40:             let email ← s'.emails
41:             let pubkey ← s'.pubkeys
42:             let body := ⟨email, pubkey⟩
43:             let xhrnonce ← nonces
44:             let s'.xhronces := s'.xhronces + ∅ xhrnonce
45:             let host.protocol such that
    ↪ ⟨host, protocol⟩ ≡ GETORIGIN(tree, docnonce)
    ↪ if possible; otherwise
    ↪ stop ⟨s', cookies, localStorage, sessionStorage, ⟨⟩⟩
46:             let command := ⟨XMLHTTPREQUEST,
    ↪ ⟨URL, protocol, host, /certreq, ⟨⟩⟩, POST, body, xhrnonce⟩
47:             stop ⟨s', cookies, localStorage, sessionStorage, command⟩
48:         end if
49:     case handleresponse
50:         let input, s' := CHOOSEINPUT(s', scriptinputs)
51:         if  $\pi_1(\textit{input}) \equiv \text{POSTMESSAGE}$  then
52:             let senderWindow :=  $\pi_2(\textit{input})$ 
53:             let senderOrigin :=  $\pi_3(\textit{input})$ 
54:             let m :=  $\pi_4(\textit{input})$ 
55:             if senderWindow ≡ PARENTWINDOW(tree, docnonce)
    ↪  $\wedge \textit{senderOrigin} \equiv \langle \text{dom}(\text{LPO}), \text{S} \rangle$  then
56:                 if  $\pi_1(\textit{m}) \in \textit{s'}.provisioningnonces$  then
57:                     let s'.emails := s'.emails + ∅  $\pi_2(\textit{m})$ 
58:                 else if  $\pi_1(\textit{m}) \in \textit{s'}.genkeypairnonces$  then

```

```

59:           let  $s'.pubkeys := s'.pubkeys + \langle \pi_2(m) \rangle$ 
60:         end if
61:         stop  $\langle s', cookies, localStorage, sessionStorage, \langle \rangle \rangle$ 
62:       end if
63:     else if  $\pi_1(input) \equiv \text{XMLHTTPREQUEST}$ 
64:        $\hookrightarrow \wedge \pi_3(input) \in s'.xhronces$  then
65:         let  $s'.ucs := s'.ucs + \langle \pi_2(input) \rangle$ 
66:         stop  $\langle s', cookies, localStorage, sessionStorage, \langle \rangle \rangle$ 
67:       end if
67: stop  $\langle scriptstate, cookies, localStorage, sessionStorage, \langle \rangle \rangle$ 

```

G Formal Security Properties

The security properties for BrowserID, informally introduced in Section 5.2, are formally defined as follows. First note that every RP service token $\langle n, i \rangle$ recorded in RP was created by RP as the result of a unique HTTPS POST request m with a valid CAP for ID i . We refer to m as the *request corresponding to* $\langle n, i \rangle$.

Definition 48. *Let \mathcal{BID} be a BrowserID web system. We say that \mathcal{BID} is secure if for every run ρ of \mathcal{BID} , every state (S_j, E_j) in ρ , every $r \in \text{RP}$ that is honest in S_j , every RP service token of the form $\langle n, i \rangle$ recorded in r in the state $S_j(r)$, the following two conditions are satisfied:*

(A) *If $\langle n, i \rangle$ is derivable from the attackers knowledge in S_j (i.e., $\langle n, i \rangle \in d_{N_{\text{attacker}}}(S_j(\text{attacker}))$), then it follows that the browser b owning i is fully corrupted in S_j (i.e., the value of `isCorrupted` is FULLCORRUPT) or `governor(i)` is not an honest IdP (in S_j).*

(B) *If the request corresponding to $\langle n, i \rangle$ was sent by some $b \in \mathcal{B}$ which is honest in S_j , then b owns i .*

H Proof of Theorem 1

In order to prove Theorem 1, we have to prove Conditions A and B of Definition 48. We prove these conditions separately. First, we provide an overview of the proofs.

H.1 Overview

For Condition (A), we analyze the request to an honest RP r upon which r returned a service token $\langle n, i \rangle$, where i is an ID and n a nonce. We show that it must contain a valid CAP (for the identity i). For this, it must in particular contain a valid UC and a matching IA. We show that the UC must have been created by the IdP that governs the identity i (which is honest by assumption). We can then show that only b can request a UC at the IdP for the identity i , and that b does not leak the private key that corresponds to the public key used for this UC, and that this key was chosen from b 's set of fresh nonces. Thus, only b can know the key that is used in the creation of the UC in the CAP. We show that neither the private key corresponding to the public key in the UC, nor the IA can leak to the attacker. Thus, the attacker cannot have sent the request corresponding to

$\langle n, i \rangle$ to the RP r . Also, $\langle n, i \rangle$ does not leak to the attacker. The attacker can therefore not know $\langle n, i \rangle$, which contradicts the assumption and proves that Condition **(A)** is satisfied.

For Condition **(B)**, we focus on the request corresponding to $\langle n, i \rangle$ as well. We observe that if the request was sent by b , the script that initiated the request was *script_rp_index*, which again got the CAP that is finally used in the request from either *script_lpo_cif* or *script_lpo_ld* (any other sources, including the attacker script, can be ruled out). In both of these scripts, the identity in the CAP is checked against the list of identities of the browser (here, the proposed patch comes into play). This ensures that the request corresponding to $\langle n, i \rangle$ contains a CAP for an identity of the browser, which contradicts the assumption that Condition **(B)** is not satisfied and thus proves the theorem.

H.2 Condition A

We assume that Condition A is not satisfied and prove that this leads to a contradiction. That is, we make the following assumption: There is a run $\rho = s_0, s_1, \dots$ of *BID*, a state $s_j = (S_j, E_j)$ in ρ , an $r \in \text{RP}$ that is honest in S_j , an RP service token of the form $\langle n, i \rangle$ recorded in r in the state $S_j(r)$ such that $\langle n, i \rangle \in d_{\text{attacker}}(S_j(\text{attacker}))$ and the browser b owning i is not fully corrupted in S_j and $\text{governor}(i)$ is an honest IdP in S_j .

By definition of RPs, for $\langle n, i \rangle$ there exists a corresponding HTTPS request received by r , which we call req_{cap} , and a corresponding response $resp_{\text{cap}}$. The request must contain a valid CAP c and must have been sent by some atomic process p to r . The response must contain $\langle n, i \rangle$ and it must be encrypted by some symmetric encryption key k sent in req_{cap} .

In particular, it follows that the request and the response must be of the following form, where $d_r \in \text{dom}(r)$ is the domain of r , $n_{\text{cap}}, k \in \mathcal{N}$ are some nonces, $path, params \in \mathcal{T}_{\mathcal{N}}$, c is some valid CAP, and sts is the Strict-Transport-Security header (as in the definition of RP's relation):

$$req_{\text{cap}} = \text{enc}_a(\langle \langle \text{HTTPReq}, n_{\text{cap}}, \text{POST}, d_r, path, params, [\text{Origin} : \langle d_r, \mathbb{S} \rangle], c \rangle, k \rangle, \text{pub}(\text{key}(d_r))) \quad (16)$$

$$resp_{\text{cap}} = \text{enc}_s(\langle \text{HTTPResp}, n_{\text{cap}}, 200, \langle sts \rangle, \langle n, i \rangle \rangle, k) \quad (17)$$

Moreover, there must exist a processing step of the following form, where $m \leq j$, $a_r \in \text{addr}(r)$, and x is some address:

$$s_{m-1} \xrightarrow[r \rightarrow \{(x:a_r:resp_{\text{cap}})\}]{(a_r:x:req_{\text{cap}}) \rightarrow r} s_m .$$

From the assumption and the definition of RPs it follows that c is of the following form:

$$\begin{aligned} c &= \langle uc, ia \rangle \\ &\equiv \langle \text{sig}(\langle i, \text{pub}(k_u) \rangle, k_{\text{sign}}), \text{sig}(\langle d_r, \mathbb{S} \rangle, k_u) \rangle \end{aligned}$$

where k_u and k_{sign} are some private keys. When we write $i = \langle i_{\text{name}}, i_{\text{domain}} \rangle$, we have that:

$$c \equiv \langle \text{sig}(\langle \langle i_{\text{name}}, i_{\text{domain}} \rangle, \text{pub}(k_u) \rangle, k_{\text{sign}}), \text{sig}(\langle d_r, \mathbb{S} \rangle, k_u) \rangle .$$

As r accepts the CAP c , we know that $\text{pub}(k_{\text{sign}}) \equiv S_j(r).\text{signkeys}[i_{\text{domain}}]$. As the subterm signkeys of r 's state is never changed, we have $S_j(r).\text{signkeys} = S_0(r).\text{signkeys}$. With the definition of the initial state of r (See Definition 42), we have that $\text{pub}(k_{\text{sign}}) \equiv S_j(r).\text{signkeys}[i_{\text{domain}}] \equiv \text{pub}(\text{signkey}(\text{dom}^{-1}(i_{\text{domain}})))$.

The private key $\text{signkey}(\text{dom}^{-1}(i_{\text{domain}}))$ is initially only known to the DY process $\text{idp} := \text{dom}^{-1}(i_{\text{domain}}) = \text{governor}(i)$. From the assumption we know that idp is an honest IdP (and not the attacker, a corrupted IdP, or some other DY process). As we can see in Algorithm 11 (that defines the behavior of IdPs), the signkey can only be used in Line 4 and in Line 32. We know that Line 4 cannot be invoked as long as idp is honest, which it is in s_j and ever since s_0 . For Line 32, we see that the key is not sent out to other processes. In s_j , the key can therefore not have been leaked to any other DY processes.

Knowing that in or before s_j , only idp can derive k_{sign} from its knowledge, it is easy to see that only idp can derive $\text{sig}(x, k_{\text{sign}})$ for any x , and in particular, uc .

Now we want to see exactly how idp creates uc and which data it uses in this process.

We have already seen that idp creates the uc in Line 32 of Algorithm 11. There may be more than one processing step in ρ where idp outputs uc .

Lemma 4. *For all processing steps of the form*

$$s_{\beta-1} \xrightarrow[\text{idp} \rightarrow \{(x:a_{\text{idp}}:\text{resp}_{uc})\}]{(a_{\text{idp}}:x:\text{req}_{uc}) \rightarrow \text{idp}} s_{\beta} \quad (18)$$

(for some addresses x , a_{idp} with $s_{\beta} < s_j$, where resp_{uc} is an encrypted HTTP response with the body (uc)) it holds that req_{uc} was emitted by b .

Proof. To reach Line 32 of Algorithm 11, several conditions have to be met for req_{uc} : It must be an encrypted HTTPS POST request with the path $/\text{certreq}$. The body of req_{uc} must be congruent to $\langle i, \text{pub}(k_u) \rangle$. The request must contain a cookie with the name sessionid and some value sessionid . This value must be a valid key for the dictionary $s'.\text{sessions}$ and

$$i \in {}^{\langle \rangle} s'.\text{sessions}[\text{sessionid}] . \quad (19)$$

Initially, $s'.\text{sessions}$ is empty. It is only populated in Line 22 of Algorithm 11. This line must have been executed in a previous processing step of the following form:

$$s_{\alpha-1} \xrightarrow[\text{idp} \rightarrow \{(x:a_{\text{idp}}:\text{resp}_{\text{auth}})\}]{(a_{\text{idp}}:x:\text{req}_{\text{auth}}) \rightarrow \text{idp}} s_{\alpha} \quad (20)$$

(for some addresses x , a_{idp} with $s_{\alpha} < s_{\beta}$). In this step, $s'.\text{sessions}$ was populated with a new entry for the session id sessionid .

From Algorithm 11 we can see that req_{auth} must meet the following conditions: It must be an HTTPS POST request, must contain a specific Origin header and its body must contain a pair $\langle i_{\text{in}}, \text{secret}_{\text{in}} \rangle$ such that the id/password combination matches a combination stored in $S_{\alpha-1}(\text{idp}).\text{users}$. As we have that $S_{\alpha-1}(\text{idp}).\text{users} = S_0(\text{idp}).\text{users}$ and with the initial definition

$$S_0(\text{idp}).\text{users} = \{ \{ \langle s, \langle \text{IDsofSecret}(s) \rangle \rangle | \text{Secrets}^i \} \} \quad (21)$$

we can see that $i_{\text{in}} \in \text{IDsofSecret}(secret_{\text{in}})$. As the list of authenticated ids in the session is then (in Line 22 of Algorithm 11) populated with $\text{IDsofSecret}(secret_{\text{in}})$ and with (19) we have that $i \in \text{IDsofSecret}(secret_{\text{in}})$. Now, IDsofSecret assigns the IDs to their secrets according to secretOfID , i.e., it must hold that

$$\text{secretOfID}(i) = secret_{\text{in}} . \quad (22)$$

This secret can be owned by at most one browser, and according to the definitions of the initial knowledge of the DY processes in **F**, it is initially only known to the owner of the secret $\text{ownerOfSecret}(secret_{\text{in}})$ (see Section F.8) and to one specific IdP (see Section F.11), in this case $i_{\text{domain}} \in \text{dom}(idp)$ (because otherwise, idp would not accept this ID).

From Algorithm 11 we can see that the IdP never uses this secret to create messages as long as it is honest, which it is by precondition.

With (22) we see that initially, only $\text{ownerOfSecret}(\text{secretOfID}(i)) = \text{ownerOfID}(i)$ knows the secret $secret_{\text{in}}$, which, by assumption, is not fully corrupted in s_j , and thus, with the request order given for (18) and (20) is not fully corrupted in s_α . (Once fully corrupted, browsers stay fully corrupted.)

(*): Honest browsers release secrets only to scripts that are loaded from a specific origin. In this case, according to the initial state given in Section F.8, the secret $\text{secretOfID}(i)$ is only released to scripts from the origin $\langle i_{\text{domain}}, \mathcal{S} \rangle$. For any such script (or document), with Lemma 2 and the definition of the browser's key mapping in Section F.8, we can see that any script that has access to the secret was sent by idp . This DY process is also the governor of i , which is, by assumption, not corrupted. Therefore, idp can only deliver either the script $script_idp_pif$ or the script $script_idp_ad$. We can now check, that both scripts, running in a browser, never send this secret to any other DY process than idp , and trigger only encrypted requests to do so.

In $script_idp_pif$ (Algorithm 17), the subterm $secret$ of the state is not used at all; therefore, the script triggers no outgoing message containing the secret at all.

In $script_idp_ad$ (Algorithm 16), $secret$ is only used as a part of a an HTTP request to the document's own origin (which therefore is the origin for which the secret is stored in the browser's list of secrets, which therefore must be $\langle i_{\text{domain}}, \mathcal{S} \rangle$). The request's data is not stored in the script's state.

We now know that all entities that have access to $secret$ (the browser b and the IdP idp) never leak it. As idp never creates any HTTP(S) requests, b must have created req_{auth} before the processing step $s_{\alpha-1} \rightarrow s_\alpha$.

In this processing step, idp creates a new session id ($sessionid$). This id is sent out only once (in Line 25 of Algorithm 11), which, in our case, is $resp_{\text{auth}}$. With Corollary 1 we can see that from this (encrypted) response $resp_{\text{auth}}$, only b can derive the contents, especially the contents of the Set-Cookie header. As in b , the cookie is stored as a *secure, HTTP only* cookie, b releases the contents of this cookie only as a Cookie header to the origin $\langle i_{\text{domain}}, \mathcal{S} \rangle$. Given the keymapping in b 's state, requests to this origin are handled by idp , and with Algorithm 11 it is easy to see that the Cookie header is only used for validating the UC request, but is not used anywhere else. All in all, b and idp do not leak the session id $sessionid$.

As $sessionid$ is an important part of req_{uc} , we can see that this request must have been emitted by b . \square

Lemma 5. *The secret key k_u was chosen by the browser b from its own nonces, i.e., $k_u \in N^b$.*

Proof. First of all, we know that for idp to generate uc , there must be a processing step in ρ of the form (described in Lemma 4):

$$s_{\beta-1} \xrightarrow[idp \rightarrow \{(x:a_{idp}:resp_{uc})\}]{(a_{idp}:x:req_{uc}) \rightarrow idp} s_{\beta} \quad (23)$$

(for some addresses x , a_{idp} with $s_{\beta} < s_j$, where $resp_{uc}$ is an encrypted HTTP response with the body $\langle uc \rangle$). For the request req_{uc} , the method must be POST and the path component must be `/certreq`.

With Lemma 4 we know that req_{uc} was emitted by b , which is honest at this point in the run. With the same arguments as in (*) we can see that either *script_idp_pif* or the script *script_idp_ad* initiated req_{uc} .

For *script_idp_ad* it is easy to see that this script never sends a POST request to *idp*.

The script *script_idp_pif* can only send a POST request to `/certreq` in Line 47 of Algorithm 17. In this case, the public key is chosen from the subterm `pubkeys` of the script's state. This subterm is only populated in Line 59 of Algorithm 17. It can only be populated by a `postMessage pm` from an immediate parent window and from the origin $\langle \text{dom}(\text{LPO}), \mathcal{S} \rangle$ (given how a browser checks and transmits `postMessages`, see Line 97f. of Algorithm 6). Further, the message in pm must be of the form $\langle n, \text{pub}(k_u) \rangle$ where n is a nonce that was freshly chosen for a `genKeyPair, n` `postMessage` in Line 28 of Algorithm 17.

Given that b 's keymapping assigns the private key of LPO to the domain of LPO and with Lemma 3 we see that the only scripts that can send such a `postMessage` are *script_lpo_cif* and *script_lpo_ld*.

In the script *script_lpo_cif* (Algorithm 13), `postMessages` of the form of pm can only be sent in Line 110 (the message sent in Line 105 would not carry the correct nonce for a response to a `genKeyPair` message).

The same holds true for the script *script_lpo_ld* (Algorithm 13).

Therefore, the key k_u is a nonce that was chosen from the browser's nonces. \square

Lemma 6. *k_u does not leak from b .*

Proof. As we have seen above, the key k_u was chosen either in the script *script_lpo_cif* or in the script *script_lpo_ld* running in the honest browser b .

In both scripts, any nonce that is chosen from the script's *nonces* will not be given to the script (as part of *nonces*) by the browser again, thus, the nonce was chosen freshly. Further, the nonce is stored in the subterm `key` of the script's state and (besides the derivation of the public key) is only used to sign IAs.

There are no other scripts running in the origin of $\langle \text{dom}(\text{LPO}), \mathcal{S} \rangle$. The (honest) browser b does not leak the script's state. Therefore, k_u does not leak from b . \square

With Lemma 4, 5, and 6, we can see that only b knows k_u and the attacker cannot know k_u . Therefore, only b can create the $ia = \text{sig}(\langle d_r, \mathcal{S} \rangle, k_u)$. As k_u is only accessible to scripts with the origin $\langle \text{dom}(\text{LPO}), \mathcal{S} \rangle$, only the script *script_lpo_cif* or the script

$script_lpo_ld$ can create $\text{sig}(\langle d_r, S \rangle, k_u)$. In both scripts, after creation, ia is sent in post-Message only to scripts that have the origin for which ia was created ($= \langle d_r, S \rangle$). With Lemma 3 and the definition of relying parties (see Algorithm 10) we see, that the only potential receiver is $script_rp_index$.

After receiving this response postMessage, $script_rp_index$ stores the UC and the IA in the subterm called cap of its scriptstate (see Algorithm 15, Line 47). After doing so, this subterm is read only in Line 72 (where only the identity is extracted) and in Line 84. There, the ia is sent to r (in the encrypted request req_{cap}).

The RP r , which is not corrupted, and the browser b do not leak ia . After receiving ia , r sends the newly created service token $\langle n, i \rangle$ to b , which ignores it (see Algorithm 15 Line 90f.). Therefore, b and r do not leak $\langle n, i \rangle$.

Therefore, the attacker cannot know $\langle n, i \rangle$ in S_j , i.e., $\langle n, i \rangle \notin d_{N^{\text{attacker}}}(S_j(\text{attacker}))$. This is a contradiction to our assumption. \square

H.3 Condition B

Similar to before, we assume that Condition B does not hold and lead this to a contradiction. We therefore make the following assumption: There is a run ρ of BID , some state $s_j = (S_j, E_j)$ in ρ , some $r \in RP$ that is honest in S_j , some RP service token of the form $\langle n, i \rangle$ recorded in r in the state $S_j(r)$, the request corresponding to $\langle n, i \rangle$ was sent by some $b \in B$ which is honest in S_j , and b does not own i .

By definition of RPs, for $\langle n, i \rangle$ there exists a corresponding HTTPS request received by r , which we call req_{cap} , and a corresponding response $resp_{cap}$. The request must contain a valid CAP c and must have been sent by some atomic process p to r . The response must contain $\langle n, i \rangle$ and it must be encrypted by some symmetric encryption key k sent in req_{cap} .

In particular, it follows that the request and the response must be of the following form, where $d_r \in \text{dom}(r)$ is the domain of r , $n_{cap}, k \in \mathcal{N}$ are some nonces, $path, params \in \mathcal{T}_{\mathcal{N}}$, c is some valid CAP, and sts is the Strict-Transport-Security header (as in the definition of RP's relation):

$$req_{cap} = \text{enc}_a(\langle \langle \text{HTTPReq}, n_{cap}, \text{POST}, d_r, path, params, [\text{Origin} : \langle d_r, S \rangle], c \rangle, k \rangle, \text{pub}(\text{key}(d_r))) \quad (24)$$

$$resp_{cap} = \text{enc}_s(\langle \text{HTTPResp}, n_{cap}, 200, \langle sts \rangle, \langle n, i \rangle \rangle, k) \quad (25)$$

Moreover, there must exist a processing step of the following form, where $m \leq j$, $a_r \in \text{addr}(r)$, and x is some address:

$$s_{m-1} \xrightarrow[r \rightarrow \{(x:a_r:resp_{cap})\}]{(a_r:x:req_{cap}) \rightarrow r} s_m \quad .$$

From the assumption and the definition of RPs it follows that c is of the following form:

$$\begin{aligned} c &= \langle uc, ia \rangle \\ &\equiv \langle \text{sig}(\langle i, \text{pub}(k_u) \rangle, k_{\text{sign}}), \text{sig}(\langle d_r, S \rangle, k_u) \rangle \end{aligned}$$

where k_u and k_{sign} are some private keys. When we write $i = \langle i_{\text{name}}, i_{\text{domain}} \rangle$, we have that:

$$c \equiv \langle \text{sig}(\langle \langle i_{\text{name}}, i_{\text{domain}} \rangle, \text{pub}(k_u) \rangle, k_{\text{sign}}), \text{sig}(\langle d_r, S \rangle, k_u) \rangle .$$

With Lemma 3 we see that this request was initiated by a script that b extracted from an HTTPS response by r . The only script that r sends in its responses is *script_rp_index*.

In this script (Algorithm 14), the only place where a request is initiated is in Line 47. We can see that the cap c is taken from the script's state, i.e., $s'.cap \equiv c$ before the execution of Line 47 must hold. Initially, this term is empty, therefore the value must have been set during the prior execution of the script. This happens in Line 47 and in Line 61 of the algorithm. For both lines to be executed, there must arrive a postMessage at *script_rp_index* (either a login or a response postMessage) from the origin of LPO.

With Lemma 2, Lemma 3, and the definition of the web browser, we can see that the message must indeed come from one of LPO's scripts, that is, either *script_lpo_ld* or *script_lpo_cif*. Before we proceed by showing that both scripts never send a UC for an identity that is not owned by browser b to the script *script_rp_index* (and later to r), we first proof the following lemma:

Lemma 7. *The value of $s'.email$ in *script_lpo_ld* is always either one of the browser's identities or empty.*

Proof. We show this by induction:

Base case: The value of $s'.email$ is initially empty (see initial scriptstate).

Induction step: The value is set only in Lines 31 and 33. In the first case, the identity is chosen non-deterministically from the browser's identities *ids*, which are the identities that the browser owns (see Section F.8).

In the second case, the value of $s'.email$ is taken from the localStorage, with the help of the key *idpnonce* that is taken from the sessionStorage. We can now show that what is retrieved from the localStorage is either empty or a previous value of $s'.email$:

First, we show that the value of *idpnonce*, taken from sessionStorage in Line 29, is always a nonce or empty: The browser's sessionStorage is separated by origins (and root windows), and therefore, only scripts under the origin of LPO have read or write access. Thus, the only two scripts that can possibly write the *idpnonce* value are *script_lpo_cif* and *script_lpo_ld*. The script *script_lpo_cif* does not write to sessionStorage. The script *script_lpo_ld* only writes to sessionStorage in Line 87. It only writes a fresh nonce (chosen in Line 85). Therefore, the value of *idpnonce* is always a nonce (or empty).

As we are already in the second case of the if-statement in Line 30 (we know that Line 33 was executed) *idpnonce* cannot be empty and must be a nonce.

Now, we can show that *localStorage[idpnonce]* is either empty or a previous value of $s'.email$: The browser's localStorage is separated by origins, and therefore, only scripts under the origin of LPO have read or write access. As above, the only two scripts that can write values to the localStorage are *script_lpo_cif* and *script_lpo_ld*. The script *script_lpo_cif* does not write to localStorage (it only removes subterms from localStorage in Line 47). We can thus focus on *script_lpo_ld*.

There are two lines where this script writes to the localStorage: Lines 112 and 86. We can safely ignore the first case, as it does not use a nonce as a key (but the fixed string *siteInfo* instead). In the latter case, it writes a value of $s'.email$.

This concludes the induction. □

We can now show (for both scripts), that they never send a UC for an identity that is not owned by the browser b :

(I) For *script_lpo_ld* (Algorithm 14), it is easy to see that the UC that is finally used to create a CAP for RP in Line 109 is set in Line 76. There, the identity in the UC is checked against the identity in $s'.email$ in the script's state (and it is checked that $s'.email$ is not empty).

With Lemma 7 and the observations above we can conclude that in *script_lpo_ld*, it is not possible that a UC for an identity that the browser does not own is accepted. Therefore, the UC that is sent to *script_rp_index* is issued for an identity of the browser b .

(II) For *script_lpo_cif* (Algorithm 13), it is easy to see that the UC that is finally used in Line 125 is set in Line 113. There, the identity in the UC is checked against the value of $s'.email$ (and, that $s'.email$ is not empty). Initially, $s'.email$ is empty. It is set only in Line 64. There, it is taken from the localStorage, using the key `siteInfo`. As we have seen above, the only place where values are stored using this key is in Line 112 of Algorithm 14. There, it is taken from the script's $s'.email$, which, according to Lemma 7, is either empty or one of the browser's identities. Note that the value of `siteInfo` is a dictionary. The keys which are used inside of this dictionary are not relevant here, but only the values.

Thus, in *script_lpo_cif*, it is not possible that a UC for an identity that the browser does not own is accepted. Therefore, the UC that is sent to *script_rp_index* is issued for an identity of the browser b .

With (I) and (II), we see that all UCs that are sent to *script_rp_index* (and later to r) are issued for identities of the browser b . This contradicts the assumption, which proves that Condition B holds true. □

I BrowserID Login Flow Overviews

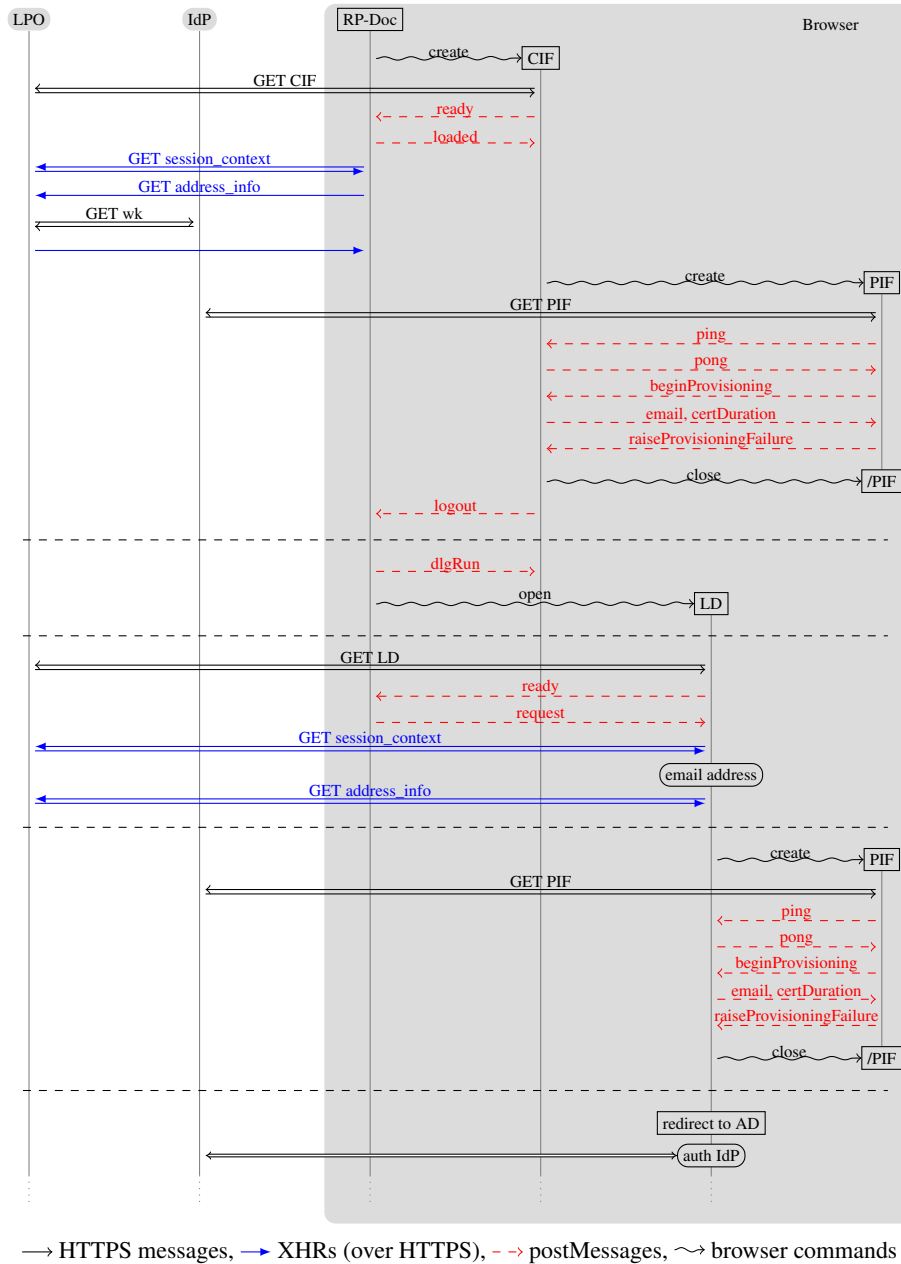



Fig. 8. BrowserID primary mode typical login flow overview (part 1 of 2).



→ HTTPS messages, → XHRs (over HTTPS), - - -> postMessages, ~> browser commands

Fig. 8. BrowserID primary mode typical login flow overview (part 2 of 2).

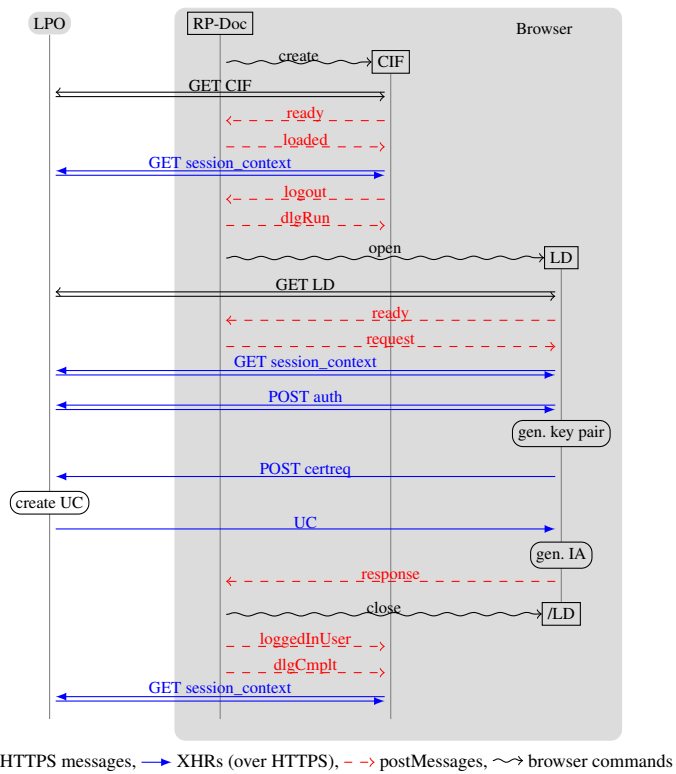


Fig. 9. BrowserID secondary mode typical login flow overview. Similar abstraction level as in Figure 8.