Provably Sound Browser-Based Enforcement of Web Session Integrity

Michele Bugliesi Stefano Calzavara Riccardo Focardi Wilayat Khan Mauro Tempesta
DAIS, Università Ca' Foscari Venezia
{bugliesi,calzavara,focardi,khan,mtempest}@dais.unive.it

***Abstract*—Enforcing protection at the browser side has recently become a popular approach for securing web authentication. Though interesting, existing attempts in the literature only address specific classes of attacks, and thus fall short of providing robust foundations to reason on web authentication security. In this paper we provide such foundations, by introducing a novel notion of web session integrity, which allows us to capture many existing attacks and spot some new ones. We then propose FF^+, a security-enhanced model of a web browser that provides a full-fledged and provably sound enforcement of web session integrity. We leverage our theory to develop SESSINT, a prototype extension for Google Chrome implementing the security mechanisms formalized in FF^+. SESSINT provides a level of security very close to FF^+, while keeping an eye at usability and user experience.**

I. INTRODUCTION

Despite the growing success of security-critical web applications, "today's Web authentication almost appears to be an exercise in demonstrating how an authentication process should *not* be realized" [28]. Besides its inherent weaknesses, password-based authentication is particularly vulnerable on the Web, since any password entered into a login form flows into the DOM of the page and is made available to any malicious script injected on it. Even when the password is not leaked during the login process, client authentication on the Web is still heavily at risk after the initial authentication step, since a large majority of web applications employ *cookies* to keep track of the authenticated sessions established upon password verification. The attack surface against cookie-based sessions is painfully large: authentication cookies can inadvertently be sent in clear over the wire [26], leaked to malicious websites through XSS flaws [21], or fixated by an attacker [27]. Moreover, untrusted parties may force the browser into creating arbitrary authenticated requests to trusted websites [10]. While current web application frameworks do allow to deploy web authentication safely at the server side, developers often misuse them, and/or are reluctant to adopt recommended security practices [36]. Enforcing protection at the browser side has thus become a popular approach for securing web authentication [15], [18]–[20], [29]–[31], [33]–[35]. Unfortunately, all the existing proposals in the literature only address very specific classes of known vulnerabilities, often lack rigorous security definitions and proofs, and eventually fall short of providing robust foundations for understanding the real effectiveness of client-side defenses for web authentication.

Contributions: In this paper we advocate the study of web authentication security through the introduction of a novel notion of (web) *session integrity*. Our theory draws on *reactive systems*, a formalism which has been previously proposed as an appropriate model of the browser behaviour [13]. Our definition of session integrity is particularly appealing, since it is browser-centric and thus naturally amenable for effective enforcement at the client side, without any background knowledge of the server behaviour. We show how our definition captures many existing attacks and spots some new ones.

We then introduce Flyweight Firefox (FF), a core model of a web browser distilled from the Featherweight Firefox model developed with the Coq proof assistant [11], [12], and we discuss FF^+, a security-enhanced extension of FF that provides a full-fledged enforcement of web session integrity. The runtime mechanisms underlying FF^+ are robust against both web threats and network attacks, and the resulting model is concrete enough to be amenable for an almost direct implementation, while at the same time being fit for a rigorous formal treatment and a security proof.

We leverage our theory to develop SESSINT, a prototype extension for Google Chrome enforcing the security policy formalized in FF^+. SESSINT is a proof of concept that the mechanisms proposed in FF^+ can be implemented in real browsers without affecting too much the user experience of many web applications. In our experiments we identify web scenarios where the security mechanisms of FF^+ need to be relaxed in order to regain usability or functionality of websites: in these cases, SESSINT warns users of the security risk, affecting as less as possible their navigation experience.

Related work: There exists a huge literature on attacks against web authentication, we refer to [28] for a good overview. The research community has proposed several solutions against these attacks in the last few years, based on server-side countermeasures [10], [21], [27], stronger web authentication schemes [4], [17], [23], [25], [28], or purely client-side solutions [15], [18]–[20], [29]–[31], [33]–[35]. In this paper we are particularly interested in the last research line, as browser-side defenses have a very wide scope and applicability: if a website does not comply with recommended security practices and/or is affected by a vulnerability, web authentication can often be protected by working solely at the browser's. Server-side defensive mechanisms or better web authentication protocols are clearly important and worth of study, since they can precisely fix the root cause of the vulnerabilities and prevent usability issues, but we consider these approaches orthogonal to our present endeavours.

We find existing client-side defenses very inspiring and we borrowed (and refined) a number of ideas from them in our work. Still, we observe that different solutions are designed around different threat models, hence it is not obvious how to soundly combine them in practice. We also notice that the lack of formal foundations in previous studies led to the development of sub-optimal solutions: we refer to Section II-A for a subtle attack which can be prevented at the browser side, but escapes state-of-the-art proposals against CSRF.

SessionShield [33] is a client-side proxy aimed at protecting authentication cookies from XSS attacks, by isolating them from JavaScript accesses. The solution protects the confidentiality of authentication cookies against web attacks, but does not enforce protection against network attacks. The same limitation applies to the competitor tool Noxes [30] and to Zan [35].

Several client-side solutions have been proposed against CSRF vulnerabilities [18], [19], [29], [31]. All these tools share the same idea of stripping authentication cookies from (selected classes of) cross-site requests, thus making CSRF attacks largely ineffective. Only the design of [19] has been formally validated, through bounded model-checking. However, the verification excludes from the threat model both XSS flaws and network attackers, which instead are two important aspects we consider in the present work.

Serene [20] is a browser-side solution against session fixation attacks. The core idea is to instruct the browser to attach to outgoing HTTP(S) requests only those authentication cookies which have been set via HTTP(S) headers, thus preventing cookies set by a malicious script from being used for authentication. Serene does not protect against network attacks, since network attackers can overwrite any cookie in the browser just by forging HTTP responses from the registering domain [1], [14]. The design of Serene has not been formally validated.

CookiExt [15] is a recent browser extension aimed at protecting the confidentiality of authentication cookies against both web and network attacks, by marking any authentication cookie received by the browser as both `HttpOnly` and `Secure`, and forcing a redirection from HTTP to HTTPS for supporting websites. The approach has been proved sound through a mechanized non-interference proof, but it does not ensure the *integrity* of authentication cookies and authenticated requests, thus leaving room for attacks like session fixation and CSRF.

A different approach to secure web authentication at the client side would be to extend the browser with a full-fledged information flow control policy, as in FlowFox [24]. At the time of writing, FlowFox does not support integrity policies, which would be central to enforcing our security notion.

Origin cookies have been proposed as a lightweight solution for protecting web sessions, by providing stronger integrity guarantees than standard cookies [14]. Origin isolation is a sound security principle and we leverage it in FF$^+$ / SESSINT. However, origin cookies do not solve the problem of protecting the first authentication step, i.e., when the password is sent from the browser to the server. Moreover, origin cookies do not directly support mixed HTTP/HTTPS websites, which instead are largely present on the Web and are supported by our solution (cf. Section IV). We also notice that the `Origin` attribute does not solve all the potential problems affecting cookie-based authentication: for instance, non-`HttpOnly` origin cookies can still be leaked via XSS, so it is not obvious what security guarantees are supported by the `Origin` attribute. On the other hand, origin cookies ensure protection against *related-domain* attackers, which is something we do not consider in our formal model for the sake of simplicity.

A seminal paper by Akhawe et al. [5] proposes a formal definition of web session integrity formulated in Alloy. Roughly speaking, the definition requires that the attacker is not involved in the "causal chain" of the events which lead to an authenticated HTTP(S) request being fired by the browser. The property is very syntactic, so it is hard to generalize it to new settings and carry out a precise comparison with our proposal. What we observe though is that the definition in [5] is only concerned about *web* attackers entering the causal chain: indeed, we argue that it would be difficult to extend the notion to deal with network attackers, since

the latter can enter the causal chain of any transaction which includes at least a communication over HTTP and trivially violate session integrity. Besides the differences in the definitions, we notice that the focus of our work is rather different with respect to [5]: here we target a security property which can be provably enforced at the browser side for *any* authenticated session, with no background knowledge about the intended server behaviour. The authors of [5], instead, use their property to verify some specific browser-server interactions (e.g., the WebAuth protocol) by bounded model-checking.

Similar considerations apply to WebSpi, a ProVerif library for modelling browsers and web applications [9]. While we find the WebSpi approach interesting and general, e.g., it has been applied also to verify cloud-storage services [8], we notice that authenticity properties in ProVerif are modelled through *correspondence assertions*: if we wanted to define web session integrity in these terms, we would need to explicitly model all the pages of the web server and its authentication goals, but this would make it difficult or even impossible to provide integrity guarantees for any authenticated session.

Armando et al. [6], [7] employ formal methods to analyse the security of existing Single Sign-On protocols, exposing real and dangerous attacks against web authentication. The approach is based on bounded model-checking, using SATMC. These papers, however, bear only limited similarities with the present work: their goal is protocol verification and the attacks they report are flaws in the protocol logic, rather than web application vulnerabilities. Their analysis abstracts from many browser-specific and web-specific aspects, which instead are central to the present paper.

Finally, we observe that our focus on client-side defenses has an important impact on the threat model we consider, which is significantly stronger than usual, since we assume that each web page may suffer of both XSS and CSRF. Given that these vulnerabilities are dangerous and widespread in practice, we argue that the design of browser-based defenses like FF^+/ SESSINT should be robust even in the presence of these server-side flaws.

Structure of the paper: Section II introduces our notion of session integrity and shows how it captures different attacks. Section III describes the browser-based enforcement of session integrity in FF^+. Section IV presents our SESSINT implementation. Section V concludes, while appendixes provide additional material and proofs.

II. SESSION INTEGRITY

Following [12], we define web browsers in terms of a very general notion of *reactive systems*, based on which we then define session integrity.

Definition 1 (Reactive System)**.** *A* reactive system *is a tuple $(\mathcal{C}, \mathcal{P}, \mathcal{I}, \mathcal{O}, \longrightarrow)$, where $\mathcal{C}$ and $\mathcal{P}$ are disjoint sets of consumer and producer states respectively, $\mathcal{I}$ and $\mathcal{O}$ are disjoint sets of input and output events respectively. The last component, $\longrightarrow$, is a labelled transition relation over the set of states $\mathcal{S} \triangleq \mathcal{C} \cup \mathcal{P}$ and the set of labels $\mathcal{A} \triangleq \mathcal{I} \cup \mathcal{O}$, defined by the following clauses:*

1) *$C \in \mathcal{C}$ and $C \xrightarrow{\alpha} Q$ imply $\alpha \in \mathcal{I}$ and $Q \in \mathcal{P}$;*
2) *$P \in \mathcal{P}$, $Q \in \mathcal{S}$ and $P \xrightarrow{\alpha} Q$ imply $\alpha \in \mathcal{O}$;*
3) *$C \in \mathcal{C}$ and $i \in \mathcal{I}$ imply $\exists P \in \mathcal{P} : C \xrightarrow{i} P$;*
4) *$P \in \mathcal{P}$ implies $\exists o \in \mathcal{O}, \exists Q \in \mathcal{S} : P \xrightarrow{o} Q$.*

A reactive system is an event-driven state machine that waits for an input, produces a sequence of outputs in response, and repeats the process indefinitely without ever getting stuck. We presuppose a lattice of security labels $(\mathcal{L}, \sqsubseteq)$, with bottom and top elements $\bot$ and $\top$. With each output event of a reactive system, we associate a label in $\mathcal{L}$ by way of a *trust* mapping $\tau : \mathcal{O} \to \mathcal{L}$. The intuition is that each label in the lattice corresponds to an interaction point for the reactive system (an *origin*, in the context of web systems), and $\tau(o) = l$ indicates that o is a message output by the reactive system (the browser) in an authenticated session with l's endpoint. We further stipulate that $\tau(o) = \bot$ whenever o does not belong to any authenticated session, and let $\tau_\bot$ stand for the trust mapping such that $\tau_\bot(o) = \bot$ for all $o \in \mathcal{O}$. Finally, we let trust change dynamically, noted $\tau \xrightarrow{o} \tau'$, upon certain output (authentication) events.

Definition 2 (Traces)**.** *Given a trust mapping τ and an input stream I, a reactive system in a state Q generates the output stream O iff the judgement $\tau \vdash Q(I) \rightsquigarrow O$ can be derived by the following inference rules:*

$$\text{(T-NIL)}\ \frac{}{\tau \vdash C([]) \rightsquigarrow []} \qquad \text{(T-IN)}\ \frac{C \xrightarrow{i} P \qquad \tau \vdash P(I) \rightsquigarrow O}{\tau \vdash C(i :: I) \rightsquigarrow O}$$

$$\text{(T-OUT)}\ \frac{P \xrightarrow{o} Q \qquad \tau \xrightarrow{o} \tau' \qquad \tau' \vdash Q(I) \rightsquigarrow O}{\tau \vdash P(I) \rightsquigarrow (o, \tau(o)) :: O}$$

A reactive system generates the trace (I, O) *if and only if $\tau_\bot \vdash C_0(I) \rightsquigarrow O$, where C_0 is the initial state of the reactive system.*

Most existing frameworks formalize integrity as a non-interference property predicating that the sensitive (high-level) outputs generated by a system should not depend on the tainted (low-level) information the system receives as an input. This simple idea becomes more complicated in the presence of active attackers, like

the network attackers we consider in this paper. Our proposal is thus reminiscent of *robustness* [22], [32], which intuitively ensures that an active attacker does not have more power than a passive attacker.

We characterize the attacker as a security label $l \in \mathcal{L}$, and define the behaviour of an attacked system in terms of a new output-generation relation $\tau, l, M \vdash Q(I) \rightsquigarrow O$, where M represents the messages the attacker was able to intercept or eavesdrop. The definition is parametric with respect to the relations of interception ($\dagger$), eavesdropping ($?$) and synthesis ($\Vdash$).

Definition 3 (Attacked Traces)**.** *Let l be an attacker. Given an input stream I and a trust mapping τ, an attacked reactive system in a given state Q generates an output stream O (written $\tau, l \vdash Q(I) \rightsquigarrow O$) if and only if the judgement $\tau, l, \emptyset \vdash Q(I) \rightsquigarrow O$ can be derived by the inference rules below:*

$$\text{(AT-NIL)}\ \frac{}{\tau, l, M \vdash C([]) \rightsquigarrow []} \qquad \text{(AT-IN)}\ \frac{C \xrightarrow{i} P \qquad \tau, l, M \vdash P(I) \rightsquigarrow O}{\tau, l, M \vdash C(i :: I) \rightsquigarrow O}$$

$$\text{(AT-OUT)}\ \frac{P \xrightarrow{o} Q \qquad \tau \xrightarrow{o} \tau' \qquad \tau', l, M \vdash Q(I) \rightsquigarrow O}{\tau, l, M \vdash P(I) \rightsquigarrow (o, \tau(o)) :: O}$$

$$\text{(AT-GETIN)}\ \frac{\tau, l \dagger i \qquad \tau, l, M \cup \{i\} \vdash Q(I) \rightsquigarrow O}{\tau, l, M \vdash Q(i :: I) \rightsquigarrow O} \qquad \text{(AT-GETOUT)}\ \frac{P \xrightarrow{o} Q \qquad \tau, l \dagger o \qquad \tau, l, M \cup \{o\} \vdash Q(I) \rightsquigarrow O}{\tau, l, M \vdash P(I) \rightsquigarrow O}$$

$$\text{(AT-HEARIN)}\ \frac{\tau, l ? i \qquad \tau, l, M \cup \{i\} \vdash Q(i :: I) \rightsquigarrow O}{\tau, l, M \vdash Q(i :: I) \rightsquigarrow O}$$

$$\text{(AT-HEAROUT)}\ \frac{P \xrightarrow{o} Q \qquad \tau \xrightarrow{o} \tau' \qquad \tau, l ? o \qquad \tau', l, M \cup \{o\} \vdash Q(I) \rightsquigarrow O}{\tau, l, M \vdash P(I) \rightsquigarrow (o, \tau(o)) :: O} \qquad \text{(AT-SYNIN)}\ \frac{C \xrightarrow{i} P \qquad \tau, l, M \Vdash i \qquad \tau, l, M \vdash P(I) \rightsquigarrow O}{\tau, l, M \vdash C(I) \rightsquigarrow O}$$

$$\text{(AT-SYNOUT)}\ \frac{\tau, l, M \Vdash o \qquad \tau \xrightarrow{o} \tau' \qquad \tau', l, M \vdash Q(I) \rightsquigarrow O}{\tau, l, M \vdash Q(I) \rightsquigarrow (o, \tau(o)) :: O}$$

A reactive system generates the attacked trace (l, I, O) *if and only if $\tau_\perp, l \vdash C_0(I) \rightsquigarrow O$, where C_0 is the initial state of the reactive system.*

Our definition of session integrity arises from contrasting the behaviour (i.e., the traces) of a reactive system in the presence, or absence, of an attacker. Given an output stream O, let $O \downarrow l$ denote the stream that results from O by considering only the events at trust level l.

Definition 4 (Session Integrity)**.** *A reactive system preserves* session integrity *for its trace (I, O) iff for all $l \in \mathcal{L}$, and all its attacked traces (l, I, O') one has:*

$$\forall l' \not\sqsubseteq l : O' \downarrow l' \text{ is a prefix of } O \downarrow l'.$$

A reactive system preserves session integrity if and only if it preserves session integrity for all its traces.

Session integrity ensures that the attacker has no effective way to interfere with any authenticated session within the set of traces. In particular, if the trust mapping remains constant at $\tau_\perp$ along the trace, no authentication event occurs in O and the attacker may only initiate its own authenticated sessions, at level l or lower. If instead the trust mapping does change, to include authenticated output events at level $l' \not\sqsubseteq l$, then the requirement that $O' \downarrow l'$ be a prefix of $O \downarrow l'$ ensures that the attacker will at best be able to interrupt the on-going sessions, but not otherwise intrude into them.

A. Web vulnerabilities as session integrity violations

We illustrate a series of attack scenarios, showing how they can be characterized as violations of our session integrity property. We refer to Appendix A for additional attacks captured by our model, i.e., password theft, login CSRF [10], and session fixation [27].

We picture the attack scenarios as diagrams in which the browser is the reactive system whose input/output events are represented by incoming/outgoing edges respectively. The inputs are generated by the user or correspond to responses from the servers (origins) the browser contacts. The outputs, in turn, are the requests made by the browser or by other origins. Each output is marked by its associated trust level. The diagrams also mark the dynamic changes to the trust mapping along the trace: these arise as a result of authentication events, whose effect is to upgrade the trust level of the cookies set upon authentication to the level of the authentication credentials. The trust level for the credentials is predefined, and given as assumptions *credential* : Origin, where each Origin corresponds to a label in the security lattice. All attack scenarios involve two origins, S and E, placed at incomparable levels in the security lattice: S is the browser's intended partner in the session, while E plays the role of the attacker (or compromised server). The diagrams provide a graphical representation of the attacked traces (cf. Definition 4). The formal encoding of the attacks in the FF model is given in Appendix A.

Cross-Site Request Forgery (Figure 1 (a)): Requested by the user, the browser establishes an authenticated session with S that the server associates with the cookie c: the cookie (the session) assumes a trust label S, based

on the assumption *pwd* : S. Later, the user opens a new page on site E in another browser tab, concluding the unattacked trace. The attacker, sitting at E, provides a response page which automatically triggers a further request to S (via XHR). Being directed to S, for which the browser has registered the cookie c, the new request includes c, thus effectively becoming part of the existing authenticated session with S in the attacked trace. Given that S $\not\sqsubseteq$ E, this violates the prefix condition in our integrity definition.

Reflected XSS (Figure 1 (b)): Like in the previous scenario, the browser establishes an authenticated session with S and associated with a cookie c : S, and later the user requests a new page on site E in another browser tab, concluding the unattacked trace. The response, provided by the attacker at E, redirects the browser to a new page $\hat{u}$ at S, passing a script as a parameter to the page. Assuming S is vulnerable to injection attacks, the script gets included in the response page at $\hat{u}$, which, when rendered, executes the script, thus leaking c to E. At this stage E may generate an output event at level S, which violates the integrity condition for the trace. In the diagram, we tacitly assume that the unattacked part of the trace is over HTTPS, the redirection forced by the attacker is over HTTP, and the cookie c is flagged as `Secure`. If the cookie was not flagged as `Secure`, the attack would resurface as a forgery, like in Figure 1 (a), since c would be attached to the request to $\hat{u}$.

Local CSRF (Figure 1 (c)): This scenario has the same structure as the reflected XSS attack represented in Figure 1 (b). The difference is that the attacker exploits the XSS vulnerability to mount a "same-site" request forgery via the injected script. As a result, unlike the XSS scenario of Figure 1 (b), this attack is effective even when the cookie is flagged as `HttpOnly`. Interestingly, this attack is not prevented by the standard browser-based protection mechanisms against CSRF [18], [19], [29], [31] that strip the cookies from cross-site requests, since the last request is not cross-site.

To the best of our knowledge, this last attack is not covered by literature on the subject. Having identified it and devising a technique to guarantee client-side protection against it represent a novel contribution.

III. Enforcing Session Integrity in FF$^+$

Here we introduce FF, a core model of a standard web browser. We then move from FF to FF$^+$, a security-enhanced variant of FF which enforces session integrity.

A. FF: *syntax and informal semantics*

We fix disjoint sets of names $\mathcal{N}$ (a, b, c, d, k, m, n, p) and variables $\mathcal{V}$ (w, x, y, z). A map M is a partial

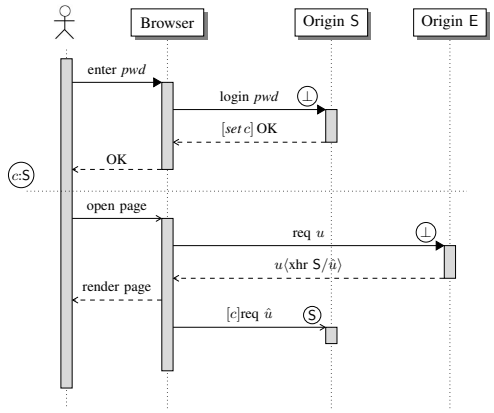

(a) Classic CSRF

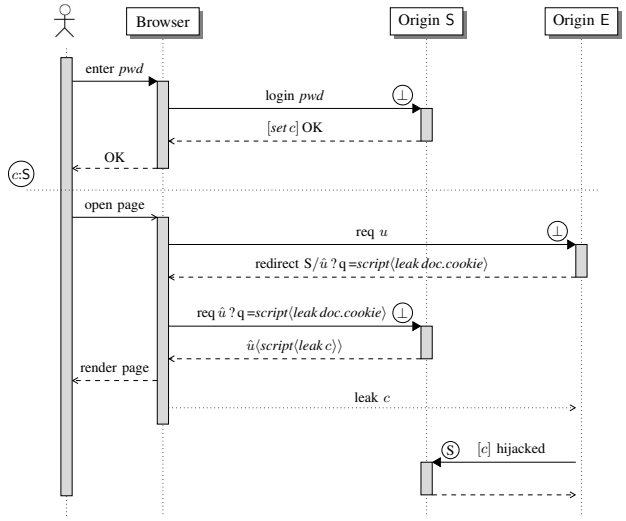

(b) Reflected XSS

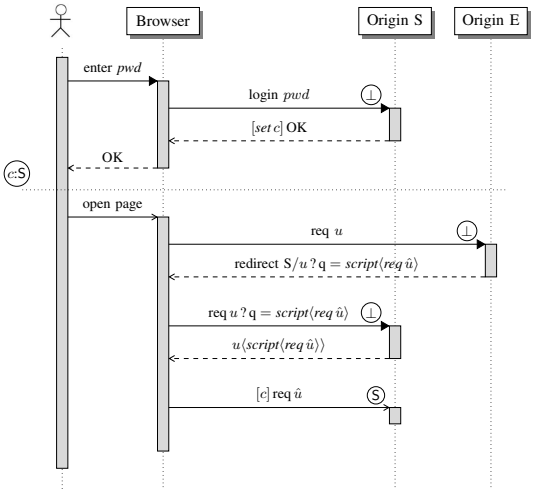

(c) Local CSRF

Figure 1: Violations of session integrity (*pwd* : S)

function from keys to values and we write $M(k) = v$ or $\{k \mapsto v\} \in M$ when the key k is bound to the value v in M; *dom*(M) denotes the domain of M and $\{\}$ is the empty map. Given two maps M_1 and M_2, $M_1 \triangleleft M_2$ is the map M such that $M(k) = v$ iff either $M_2(k) = v$ or $M_1(k) = v$ and $k \notin$ *dom*(M_2), while $M_1 \uplus M_2$ is the map $M_1 \triangleleft M_2$ whenever *dom*$(M_1) \cap$ *dom*$(M_2) = \emptyset$.

1) URLs: We let $\pi \in \{\mathsf{http}, \mathsf{https}\}$ note a protocol identifier. A URL $u \in \mathcal{U}$ is either the constant blank or a triple (π, d, v), where d is a domain name and v is a value encoding additional information, like the full path of the accessed resource or a query string. For $u = (\pi, d, v)$ we let $domain(u) = d$ and $path(u) = v$.

2) Cookies: Cookies are collected in maps ck such that $ck(k) = (n, f)$ whenever the cookie named k is bound to the value n and marked with $f \in \{\mathtt{H}, \mathtt{S}, \top, \bot\}$. Flag `H` models `HttpOnly` cookies, which must not be accessed by JavaScript and only be included in HTTP(S) requests to the registering domain. Flag `S`, in turn, models `Secure` cookies, which must only be sent over encrypted connections. Finally, flag $\bot$ is for cookies with no special security requirements, while $\top$ marks cookies which are both `HttpOnly` and `Secure`. We let $ck\_vals(ck) = \{n \mid \exists k, f : ck(k) = (n, f)\}$.

3) Values and expressions: We let v range over values, i.e., unit, URLs, names, variables and functions:

$$v ::= () \mid u \mid n \mid x \mid \lambda x.e.$$

We let e range over expressions of a simple scripting language which includes first-class functions, basic operations on cookies and the creation of AJAX requests:

$$\begin{array}{rcl} e & ::= & v\ v' \mid \mathsf{let}\ x = e\ \mathsf{in}\ e' \mid v? \mid v!\langle v', f\rangle \\ & \mid & \mathsf{xhr}(v, v') \mid \mathsf{auth}(v, v') \mid v. \end{array}$$

$(\lambda x.e)\ v$ evaluates to $e\{v/x\}$; $\mathsf{let}\ x = e\ \mathsf{in}\ e'$ first evaluates e to a value v and then behaves as $e'\{v/x\}$; $k?$ returns the value of cookie k, provided that k is not flagged `HttpOnly`; $k!\langle n, f\rangle$ with $f \in \{\bot, \mathtt{S}\}$ stores the cookie $\{k \mapsto (n, f)\}$ in the cookie jar, ensuring that no existing `HttpOnly` cookie is overwritten. The expression $\mathsf{xhr}(u, \lambda x.e)$ sends an AJAX request to u and, whenever a value v is available as a response, it behaves as $e\{v/x\}$. Finally, $\mathsf{auth}(u, p)$ sends to the URL u the password p.

4) Event handlers: We let h range over (sets of) event handlers, i.e., maps from names to functions. If $h(k) = \lambda x.e$, a handler registered on k is ready to run e, with x bound to the value received along with the firing event. FF handlers model two different aspects of web browsing: first, we use them to encode event-driven JavaScript programming (indeed, we represent the DOM with a set of event handlers). Second, a new handler is instantiated when an AJAX request is sent to a server and it is triggered only when a response is sent back.

5) Pages: Pages are triples $page ::= (u, h, h')$, where u keeps track of the *origin* of the page, h is a set of event handlers registered on the DOM, and h' is a dynamic set of handlers, which grows/shrinks when new AJAX requests/responses are sent/received by the page.

6) Events: Input events i are defined as follows:

$$\begin{array}{rcl} i & ::= & \mathsf{load}(u) \mid \mathsf{text}(p, k, n) \\ & \mid & \mathsf{doc\_resp}(n, ck, u, u', h, e) \\ & \mid & \mathsf{xhr\_resp}(n, ck, u, u', v). \end{array}$$

Event $\mathsf{load}(u)$ models the user navigating the web browser to u: the browser reacts to the event by opening a new network connection to u and sending a request for the document located there. Event $\mathsf{text}(p, k, n)$ corresponds to the user inserting a value n in the text field k of page p: if p contains a set of handlers h such that $h(k) = \lambda x.e$, the event triggers the expression $e\{n/x\}$. Event $\mathsf{doc\_resp}(n, ck, u, \mathsf{blank}, h, e)$ models the receipt of a response from u over the network connection n: the browser will store the cookies ck in its cookie jar, render the document structure (modelled as the set of handlers h) and then run the expression e. Event $\mathsf{doc\_resp}(n, ck, u, u', h, e)$ with $u' \neq \mathsf{blank}$ represents a redirect from u to u': in this case, the cookies ck are stored by the browser, but both h and e are ignored. Event $\mathsf{xhr\_resp}(n, ck, u, \mathsf{blank}, v)$ corresponds to the receipt of an AJAX response from u over the network connection n: the browser will store the cookies ck, then it will retrieve the continuation $\lambda x.e$ which must be triggered by the response, and it will run the expression $e\{v/x\}$. Again, event $\mathsf{xhr\_resp}(n, ck, u, u', v)$ with $u' \neq \mathsf{blank}$ models a redirect from u to u' triggered by an AJAX response (where ck is stored, v is ignored).

Output events o are defined as follows:

$$\begin{array}{rcl} o & ::= & \bullet \mid \mathsf{doc\_req}(ck, u) \mid \mathsf{xhr\_req}(ck, u) \\ & \mid & \mathsf{login}(ck, u, p). \end{array}$$

The dummy event $\bullet$ represents a silent reaction to an input event with no observable side-effect. Event $\mathsf{doc\_req}(ck, u)$ models a document request to u, attaching the cookies ck: it is triggered either by a $\mathsf{load}(u)$ event, or when the browser follows a redirect targeted at u after a document response. Event $\mathsf{xhr\_req}(ck, u)$ models an AJAX request to u, attaching the cookies ck: it is triggered either by the expression $\mathsf{xhr}(u, \lambda x.e)$, or when the browser is redirected to u after an AJAX response. Finally, $\mathsf{login}(ck, u, p)$ represents a request to u which includes the password p, corresponding to the submission of a login form: the occurrence of this event may signal the establishment of a new session. The event

is triggered by the expression $\mathsf{auth}(u, p)$ and it includes the cookies ck which must be sent to u.

We let $\alpha ::= i \mid o$ range uniformly over input and output events. We refer to requests, responses and logins as *network* events.

7) Browser states: Browser states are 5-tuples $Q = \langle W, K, N, T, O\rangle$ where:

$$\begin{aligned} \textit{Windows } W &::= \{\} \mid \{p \mapsto \textit{page}\} \mid W \uplus W, \\ \textit{Cookies } K &::= \{\} \mid \{d \mapsto ck\} \mid K \uplus K, \\ \textit{Networks } N &::= \{\} \mid \{n \mapsto (u, v)\} \mid N \uplus N, \\ \textit{Tasks } T &::= \{\} \mid \{p \mapsto e\}, \\ \textit{Outputs } O &::= [\,] \mid o. \end{aligned}$$

The window store W maps fresh page identifiers to pages, while the cookie jar K maps domain names to the cookies they registered in the browser. The network connection store N keeps track of the open network connections: if $\{n \mapsto (u, v)\} \in N$, then the browser is waiting for a document/AJAX response from u (the role of v will be apparent in the formal semantics). We use T to represent *tasks*: if $\{p \mapsto e\} \in T$, then the expression e is running in the page p. Finally, O is a size-1 buffer of output events, which is convenient to interpret our model as a reactive system.

We say that $Q = \langle W, K, N, T, O\rangle$ is a *consumer* state when both T and O are empty and we denote it with C, otherwise we say that Q is a *producer* state and we denote it with P. The formal semantics of FF is given in Appendix B.

B. Session establishment

Our definition of session integrity relies on a lattice of security labels, which we instantiate next.

Definition 5 (Security Labels)**.** *The set of* security labels *$\mathcal{L}$, ranged over by l, is the smallest set generated by the following grammar:*

$$l := \bot \mid \top \mid \mathsf{evil} \mid \mathsf{net} \mid \pi(d) \textit{ with } \pi \in \{\mathsf{http}, \mathsf{https}\}.$$

We define $\sqsubseteq$ as the least pre-order over $\mathcal{L}$ with $\bot$ as a bottom element, $\top$ as a top element, induced by the axioms: $\{\mathsf{evil} \sqsubseteq \mathsf{http}(d), \mathsf{http}(d) \sqsubseteq \mathsf{net}, \mathsf{net} \sqsubseteq \mathsf{https}(d)\}$.

We assume a partial function $\mathit{url\_label} : \mathcal{U} \to \mathcal{L}$ such that $\mathit{url\_label}(u) = \pi(d)$ whenever $u = (\pi, d, v)$. We also stipulate that the set of names $\mathcal{N}$ is partitioned into the indexed family $\{\mathcal{N}_l\}_{l \in \mathcal{L}}$: this is needed to capture in the model the inability of the attacker to guess random secrets, like passwords or authentication cookie values.

We adopt password-based authentication to establish new sessions with remote web servers. Simply put, when a valid password is submitted to a website supporting authenticated access, a cookie is endorsed to identify the password's owner for the session. Formally, this amounts to instantiating the relation $\tau \xrightarrow{o} \tau'$ underlying the semantics of reactive systems (cf. Definition 2). For this purpose, we presuppose a function $\rho : \mathcal{N} \to \mathcal{L}$ with the following understanding: if $\rho(n) = \pi(d)$, then n is the user's password for the website at d and can be exchanged on the protocol π. We let $\rho(n) = \mathsf{evil}$ whenever n is a password identifying the attacker's account: for simplicity, we assume that this password can be used to establish authenticated sessions on any website. We assume ρ to be *consistent* with respect to the partitioning of names, i.e., we stipulate $\rho(n) \sqsubseteq l$ whenever $n \in \mathcal{N}_l$.

Let now $\mathcal{U}_{\mathit{auth}} \subseteq \mathcal{U}$ be the set of the URLs containing a login form for password-based authentication. We assume that $\mathcal{U}_{\mathit{auth}}$ is partitioned into two subsets $\mathcal{U}_{\mathit{ok}}$ and $\mathcal{U}_{\mathit{fix}}$. If a valid password c is sent to $u \in \mathcal{U}_{\mathit{ok}}$, a fresh authentication cookie is created by the server and employed to identify the password's owner; if $u \in \mathcal{U}_{\mathit{fix}}$, instead, the server may be subject to session fixation, hence it endorses for authentication a cookie already included in the login request. In both cases the (only) authentication cookie is chosen by a function $\kappa : \mathcal{U}_{\mathit{auth}} \to \mathcal{N}$ identifying its name, and the trust mapping is updated to reflect that any output event o including that cookie will have the trust level $\rho(c)$ bound to the password, much like in the examples of Section II-A. The formal details are in Table 1 and commented below.

Rule (A-SRV) models a login on $u \in \mathcal{U}_{\mathit{ok}}$. If c is a valid password, a fresh value n is picked from the name partition $\mathcal{N}_{\rho(c)}$ based on an underlying total order ($n \leftarrow \mathcal{N}_{\rho(c)}$). The value n will be used to identify the password's owner: specifically, we perform a pointwise join between the original trust function τ and the auxiliary trust function $\tau_{u,n,c}$ in the table, which raises to $\rho(c)$ the trust of the output events sent to $\mathit{domain}(u)$ which include the cookie $\{\kappa(u) \mapsto (n, f)\}$ for some f.

Rule (A-FIX), instead, models a login on $u \in \mathcal{U}_{\mathit{fix}}$. In this case, the value n bound to the key $k = \kappa(u)$ among the cookies ck sent to the server will be used to identify the password's owner. If $k \notin \mathit{dom}(ck)$, the authentication fails and rule (A-NIL) must be applied.

C. Threat model

We assume that all HTTPS traffic is signed using trusted certificates (unsigned HTTPS traffic is represented using HTTP). The attacker's power is characterized by a security label l, with the understanding that higher labels provide additional capabilities. A novel aspect of our threat model is that we assume the attacker has full control over *compromised* sessions, i.e.,

TABLE 1 Rules for password-based authentication

$$\text{(A-SRV)}\ \frac{u \in \mathcal{U}_{ok} \qquad n \leftarrow \mathcal{N}_{\rho(c)} \qquad \rho(c) \in \{url\_label(u), \mathsf{evil}\}}{\tau \xrightarrow{\mathsf{login}(ck,u,c)} \tau \sqcup \tau_{u,n,c}}$$

$$\text{(A-FIX)}\ \frac{u \in \mathcal{U}_{fix} \qquad \kappa(u) = k \qquad ck(k) = (n, f) \qquad \rho(c) \in \{url\_label(u), \mathsf{evil}\}}{\tau \xrightarrow{\mathsf{login}(ck,u,c)} \tau \sqcup \tau_{u,n,c}}$$

$$\text{(A-NIL)}\ \frac{\alpha \text{ has a different form}}{\tau \xrightarrow{\alpha} \tau}$$

where $\tau_{u,n,c}(o) = \begin{cases} \rho(c) & \text{if } o \in \{\{\mathsf{doc}, \mathsf{xhr}\}\_\mathsf{req}(ck', u') \mid domain(u) = domain(u') \wedge ck'(\kappa(u)) = n \wedge \tau(o) \sqsubseteq \rho(c)\} \\ \bot & \text{otherwise} \end{cases}$

authenticated sessions established using the attacker's credentials. If a network request belongs to a compromised session, we pessimistically assume that all the data included in the request are stored by the server in the attacker's account and later made available to him: this is useful to capture login CSRF attacks [10].

Formally, the threat model results from instantiating the definitions of interception (†), eavesdropping (?) and synthesis ($\Vdash$) in Definition 2. Let $ev\_label : \mathcal{A} \to \mathcal{L}$ be the function such that $ev\_label(\alpha) = url\_label(u)$ whenever α is a network event sent to/received from u; we assume $ev\_label(\alpha) = \top$ whenever α is not a network event.

The relations † and ? are defined as follows:

$$\text{(II-NET)}\ \frac{ev\_label(\alpha) \sqsubseteq l}{\tau, l \dagger \alpha} \qquad \text{(IH-NET)}\ \frac{ev\_label(\alpha) \sqcap \mathsf{net} \sqsubseteq l}{\tau, l \,?\, \alpha} \qquad \text{(IH-EVIL)}\ \frac{\tau(o) = \mathsf{evil}}{\tau, l \,?\, o}$$

According to rule (II-NET), a web attacker at level, say, $\mathsf{http}(d)$ can intercept only the network traffic sent to d either in clear or with no trusted certificates, while a network attacker can intercept all the HTTP traffic (and any HTTPS message directed to him). We remark that a net-level attacker cannot intercept arbitrary HTTPS traffic: indeed, since signed HTTPS communication ensures both freshness and integrity [3], the attacker cannot replay encrypted messages or otherwise tamper with HTTPS exchanges without breaking the communication session. Hence, preventing the interception of arbitrary HTTPS traffic ultimately amounts just to discarding denial of service attacks, which we are not interested to deal with in the present paper. Notice, however, that an HTTPS exchange can still be overheard by a net-level attacker using rule (IH-NET): network attackers are thus aware of all the network traffic, even though they may be unable to access its payload. Finally, rule (IH-EVIL) makes any request sent over compromised sessions available to the attacker, as we discussed above.

Defining the relation $\tau, l, M \Vdash \alpha$ is slightly more complex. We start by defining an auxiliary relation $\tau, l, M \Vdash n$, which identifies the names that can be generated by the attacker:

$$\text{(NS-BASE)}\ \frac{n \in \mathcal{N}_{l'} \qquad l' \sqsubseteq l}{\tau, l, M \Vdash n} \qquad \text{(NS-LOOK)}\ \frac{ev\_label(\alpha) \sqsubseteq l \qquad n \in fn(\alpha)}{\tau, l, M \cup \{\alpha\} \Vdash n} \qquad \text{(NS-EVIL)}\ \frac{\tau(o) = \mathsf{evil} \qquad n \in fn(o)}{\tau, l, M \cup \{o\} \Vdash n}$$

According to (NS-BASE), an l-attacker can generate any name in a name partition indexed by a label bounded above by l. By rule (NS-LOOK) the attacker may generate the free names of any network event α previously intercepted or overheard, provided that the attacker can inspect its payload. Finally, rule (NS-EVIL) grants the attacker the capability to generate any name communicated over compromised sessions.

Now, we can define the relation $\tau, l, M \Vdash \alpha$:

$$\text{(IS-GEN)}\ \frac{\alpha = i \Rightarrow ev\_label(\alpha) \sqsubseteq l \qquad \forall n \in fn(\alpha) : \tau, l, M \Vdash n}{\tau, l, M \Vdash \alpha} \qquad \text{(IS-REP)}\ \frac{\alpha \in M \qquad ev\_label(\alpha) \sqsubseteq \mathsf{net} \sqsubseteq l}{\tau, l, M \Vdash \alpha}$$

By rule (IS-GEN), an l-attacker can forge an input event i, provided that he can generate all the free names in i and the event label of i is bounded above by l: the latter condition ensures, for instance, that a net-level attacker cannot forge signed HTTPS traffic and that a web attacker $\mathsf{http}(d)$ cannot provide responses for another web server at d'. Rule (IS-GEN) also allows the attacker to send arbitrary output events to any server, provided that he is able to compose the request contents. Finally, rule (IS-REP) allows an attacker with network capabilities (side-condition $\mathsf{net} \sqsubseteq l$) to replay previously intercepted/ overheard traffic. Since HTTPS ensures freshness, the side-condition $ev\_label(\alpha) \sqsubseteq \mathsf{net}$ similarly guarantees that encrypted traffic cannot be replayed.

We conclude this section with a note on XSS attacks: we implicitly include them in our model, since our session integrity property quantifies over all the possible inputs made available to the browser. This universal quantification grants any attacker the capability to mount reflected XSS attacks on any website.

D. FF$^+$: a secure extension of FF

FF provides a faithful abstraction of current web browsers and, just like them, it is vulnerable to a variety of attacks. In this section, we discuss the design of FF$^+$, a security-enhanced extension of FF aimed at enforcing web session integrity.

1) Qualifiers: FF lacks the contextual information needed to apply a sound security policy for session integrity, since it does not track origin changes across network requests. We fix this by extending the structure of network connections and pages with *qualifiers*, by having $N ::= \{\} \mid \{n \mapsto (u, v, q)\} \mid N \uplus N$ and *page* $::= (u, h, h', q)$. A qualifier $q \in \{\checkmark, \times\}$ is just a boolean mark used to *taint track* the open network connections. Pages downloaded from a given connection inherit the qualifier assigned to the connection, and connections become tainted when a cross-origin redirect is performed over them. FF$^+$ enforces different security policies on a page based on the value of its qualifier.

2) Security contexts: FF must also be enhanced to prevent the risk of password theft. When the user enters a password into a login form, an event handler registered on the page can steal the password and leak it to the attacker. We address this issue by running each expression e inside a *security context*, i.e., a sandbox represented by a pair (e, l). If $l = \pi(d)$, then the expression e is allowed to communicate only with d on the protocol π. When a password n is disclosed to an expression e, we instantiate a new security context $(e, \rho(n))$, which provides FF$^+$ with the information needed to protect n. Clearly, this assumes that FF$^+$ keeps track of $\rho(n)$ for any password n input by the user, for instance by using an internal password manager[1]. Formally, we enrich the syntax of tasks by having $T ::= \{\} \mid \{p \mapsto (e, l)\}$.

3) Secure cookie operations: Updates to the cookie jar in FF$^+$ adopt a strong security policy, whereby authentication cookies received over HTTP are marked `HttpOnly`, while authentication cookies received over HTTPS are flagged both `HttpOnly` and `Secure`. If a `Secure` cookie is sent from the server to the browser over HTTP, which is one of the many quirks allowed on the Web, it is discarded by FF$^+$. Moreover, FF$^+$ strengthens the integrity of cookies set over HTTPS against network attacks, by ensuring that cookies which are marked as both `HttpOnly` and `Secure` are never overwritten by cookies set through HTTP responses. This is not ensured by standard web browsers [1] and previous proposals already highlighted the dangers connected to this practice [14]. The formal details correspond to the secure cookie update function *sec_upd_ck* in Table 2.

We also introduce a secure counterpart of the standard procedure employed by web browsers to select the cookies to be attached to a given network request. Specifically, FF$^+$ ensures that no outgoing cookie can have been fixated by an attacker: for HTTP requests we enforce protection against web attacks, by requiring that only `HttpOnly` cookies are sent to the web server. Since these cookies cannot be set by a script, they can only be fixated by network attacks. For HTTPS requests, instead, we target a higher level of protection and we ensure that any cookie attached to them cannot have been fixated, even by a network attacker. Accordingly with the previous discussion, we thus impose that only cookies which are marked as both `Secure` and `HttpOnly` are attached to HTTPS requests. The formal details amount to the definition of the function *get_http_ck* in Table 2.

4) Inputs: The transitions $C \xrightarrow{i} P$ in Table 3 describe how the consumer state C reacts to the input i by evolving into a producer state P. The definition of $C \xrightarrow{i} P$ consists only of two rules, i.e., (I-MIRROR) and (I-COMPLETE). The definition relies on the auxiliary relation $C \overset{i}{\mapsto} P$, which is the bulk of the semantics: this is convenient to interpret FF$^+$ as a reactive system.

We start with the behaviour of document requests and responses. When the user navigates the browser to a URL u, a new network connection n is created and it is assigned the qualifier ✓ by rule (I-LOAD). Moreover, a new document request event is generated and put in the output buffer. If a cross-origin redirect is received over n, the connection is given the qualifier ✕ and becomes tainted, and it will never be restored to an untainted state by rule (I-DOCREDIR). Further requests sent over a tainted connection n will never include cookies, to thwart CSRF attacks performed through a redirect: this policy is applied also to same-origin requests, to prevent local CSRF attacks similar to the one described in Section II-A. When a document response is eventually received over the network connection n, the connection is closed and a new page is stored in the browser by rule (I-DOCRESP). The page inherits the qualifier assigned to n and the cookie jar is updated only if n was marked as untainted: this is needed to prevent the attacker from corrupting the cookies stored in the browser through malicious redirects.

Text input events are handled by rule (I-TEXT) as anticipated, by letting the disclosed expression $e\{n/x\}$ run in the security context $\rho(n)$, which will ensure that the confidentiality of passwords is protected. Finally, AJAX responses are processed much like document re-

[1]For simplicity, in the formal model each password is associated to a single origin. Our implementation allows to reuse the same password on different websites (cf. Section IV).

TABLE 2 Secure management of the cookie jar

$$\frac{}{\{\} \nearrow \pi = \{\}} \qquad \frac{f \in \{\bot, \mathtt{H}\}}{\{k \mapsto (n,f)\} \nearrow \mathsf{http} = \{k \mapsto (n, \mathtt{H})\}} \qquad \frac{f \in \{\mathtt{S}, \top\}}{\{k \mapsto (n,f)\} \nearrow \mathsf{http} = \{\}}$$

$$\frac{}{\{k \mapsto (n,f)\} \nearrow \mathsf{https} = \{k \mapsto (n, \top)\}} \qquad \frac{ck_1 \nearrow \pi = ck_1' \qquad ck_2 \nearrow \pi = ck_2'}{(ck_1 \uplus ck_2) \nearrow \pi = ck_1' \uplus ck_2'}$$

For $u = (\pi, d, v)$, we let:

$$\mathit{sec\_upd\_ck}(K, u, ck) = \begin{cases} K \uplus \{d \mapsto (ck \nearrow \pi)\} & \text{if } d \notin \mathit{dom}(K) \\ K' \uplus \{d \mapsto (ck' \triangleleft (ck \nearrow \pi))\} & \text{if } K = K' \uplus \{d \mapsto ck'\} \wedge \pi = \mathsf{https} \\ K' \uplus \{d \mapsto (ck_h \triangleleft (ck \nearrow \pi \triangleleft ck_s))\} & \text{if } K = K' \uplus \{d \mapsto ck_h \uplus ck_s\} \wedge \pi = \mathsf{http} \end{cases}$$

where:

$$\forall k \in \mathit{dom}(ck_h) : ck_h(k) = (n,f) \Rightarrow f \in \{\bot, \mathtt{H}, \mathtt{S}\}$$
$$\forall k \in \mathit{dom}(ck_s) : ck_s(k) = (n,f) \Rightarrow f = \top.$$

Finally, we let $\mathit{get\_http\_ck}(K, u)$ be defined as the least map M such that:

$$M(k) = \begin{cases} (n, \top) & \text{if } u = (\mathsf{https}, d, v) \wedge \exists ck : K(d) = ck \wedge ck(k) = (n, \top) \\ (n, \mathtt{H}) & \text{if } u = (\mathsf{http}, d, v) \wedge \exists ck : K(d) = ck \wedge ck(k) = (n, \mathtt{H}) \end{cases}$$

sponses. The only interesting differences from a security perspective are in rule (I-XHRRESP), where we must additionally instantiate the label of the new security context to the url_label of the page which sent the AJAX request: this is needed to protect the confidentiality of passwords when the continuation of an AJAX request is executed. It is also worth noticing that we require the qualifier q of the network connection to match the qualifier of the page where the response is received: loading tainted scripts inside an untainted page would be unsound, since these scripts would be allowed to send authenticated requests (see below).

5) Outputs: Table 4 collects the transitions $P \xrightarrow{o} Q$, describing how a producer state P can generate an output o and evolve into another state Q. Several rules are standard, so we just comment the most interesting points.

Rule (O-SET) models the setting of a cookie via JavaScript. The upd_ck function stands for the standard cookie update operation available in web browsers (cf. Appendix B). The only point worth mentioning here is the security label $\bot$ required on the security context, which is needed to prevent confidentiality leaks resulting by setting a cookie containing password information.

Rule (O-XHR) is the most complex and models the sending of an AJAX request by an expression running on a page. Different security policies are applied, based on the qualifier of the page and the label of the security context where the expression is run. Let u' be the URL of the page, q the qualifier of the page, l the label of the security context, and u the destination of the AJAX request. When $l = \bot$, no password was previously typed by the user and no confidentiality policy is enforced; otherwise, FF^+ allows the sending of the request only if $l = url\_label(u)$. We also require $l = url\_label(u')$ to prevent a password leakage when the asynchronous continuation of an AJAX request is disclosed, as we anticipated in rule (I-XHRRESP). Moreover, FF^+ strips the cookies from outgoing requests when $url\_label(u) \neq url\_label(u')$ or $q = \mathsf{X}$: this prevents both classic and local CSRF attacks. Notice that only untainted pages are allowed to open untainted network connections via XHR.

We conclude with rule (O-LOGIN). The condition $\rho(c) = url\_label(u)$ prevents login CSRF attacks, where the user is authenticated as the attacker, while the requirement $l = url\_label(u)$ ensures the confidentiality of the password. We also require that any login form is submitted to a URL within the same origin of the page: this prevents the attacker from fooling the user into establishing new authenticated sessions with trusted websites, which would violate session integrity.

E. Formal results

We can prove that FF^+ enforces session integrity for any *well-formed* trace. Intuitively, well-formedness ensures a basic set of constraints on incoming input events, which are needed for our formal result, but have a limited practical impact. Clearly, we do not assume that the intruder is forced to produce well-formed inputs.

TABLE 3 Reactive semantics of FF^{+}: inputs

(I-LOAD)
$$\frac{ck = \mathit{get\_http\_ck}(K, u)}{\langle W, K, N, \{\}, [\,]\rangle \xrightarrow{\mathsf{load}(u)} \langle W, K, N \uplus \{n \mapsto (u, (), \checkmark)\}, \{\}, \mathsf{doc\_req}(ck, u)\rangle}$$

(I-TEXT)
$$\frac{W(p) = (u, h, h', q) \qquad h(k) = \lambda x.e}{\langle W, K, N, \{\}, [\,]\rangle \xrightarrow{\mathsf{text}(p,k,n)} \langle W, K, N, \{p \mapsto (e\{n/x\}, \rho(n))\}, [\,]\rangle}$$

(I-DOCRESP)
$$\frac{q = \checkmark \Rightarrow K' = \mathit{sec\_upd\_ck}(K, u, ck) \qquad q = \times \Rightarrow K' = K}{\langle W, K, N \uplus \{n \mapsto (u, (), q)\}, \{\}, [\,]\rangle \xrightarrow{\mathsf{doc\_resp}(n,ck,u,\mathsf{blank},h,e)} \langle W \uplus \{p \mapsto (u, h, \{\}, q)\}, K', N, \{p \mapsto (e, \bot)\}, [\,]\rangle}$$

(I-DOCREDIR)
$$\frac{\begin{array}{c} q = \checkmark \Rightarrow K' = \mathit{sec\_upd\_ck}(K, u, ck) \qquad q = \times \Rightarrow K' = K \\ q = \checkmark \wedge \mathit{url\_label}(u) = \mathit{url\_label}(u') \Rightarrow ck' = \mathit{get\_http\_ck}(K', u') \wedge q' = \checkmark \\ q = \times \vee \mathit{url\_label}(u) \neq \mathit{url\_label}(u') \Rightarrow ck' = \{\} \wedge q' = \times \end{array}}{\langle W, K, N \uplus \{n \mapsto (u, (), q)\}, \{\}, [\,]\rangle \xrightarrow{\mathsf{doc\_resp}(n,ck,u,u',h,e)} \langle W, K', N \uplus \{n \mapsto (u', (), q')\}, \{\}, \mathsf{doc\_req}(ck', u')\rangle}$$

(I-XHRRESP)
$$\frac{\begin{array}{c} q = \checkmark \Rightarrow K' = \mathit{sec\_upd\_ck}(K, u, ck) \qquad q = \times \Rightarrow K' = K \\ h' = h'' \uplus \{n \mapsto \lambda x.e\} \qquad W' = W \uplus \{p \mapsto (u', h, h'', q)\} \qquad l = \mathit{url\_label}(u') \end{array}}{\langle W \uplus \{p \mapsto (u', h, h', q)\}, K, N \uplus \{n \mapsto (u, p, q)\}, \{\}, [\,]\rangle \xrightarrow{\mathsf{xhr\_resp}(n,ck,u,\mathsf{blank},v)} \langle W', K', N, \{p \mapsto (e\{v/x\}, l)\}, [\,]\rangle}$$

(I-XHRREDIR)
$$\frac{\begin{array}{c} q = \checkmark \Rightarrow K' = \mathit{sec\_upd\_ck}(K, u, ck) \qquad q = \times \Rightarrow K' = K \\ q = \checkmark \wedge \mathit{url\_label}(u) = \mathit{url\_label}(u') \Rightarrow ck' = \mathit{get\_http\_ck}(K', u') \wedge q' = \checkmark \\ q = \times \vee \mathit{url\_label}(u) \neq \mathit{url\_label}(u') \Rightarrow ck' = \{\} \wedge q' = \times \end{array}}{\langle W, K, N \uplus \{n \mapsto (u, p, q)\}, \{\}, [\,]\rangle \xrightarrow{\mathsf{xhr\_resp}(n,ck,u,u',v)} \langle W, K', N \uplus \{n \mapsto (u', p, q')\}, \{\}, \mathsf{xhr\_req}(ck', u')\rangle}$$

(I-MIRROR)
$$\frac{C \overset{i}{\mapsto} P}{C \xrightarrow{i} P}$$

(I-COMPLETE)
$$\frac{\langle W, K, N, \{\}, [\,]\rangle \overset{i}{\not\mapsto}}{\langle W, K, N, \{\}, [\,]\rangle \xrightarrow{i} \langle W, K, N, \{\}, \bullet\rangle}$$

Notation: we write $C \overset{i}{\not\mapsto}$ whenever there does not exist P such that $C \overset{i}{\mapsto} P$.

We say that a URL u is well-formed (written $\vdash_\diamond u$) iff $\mathit{domain}(u) \in \mathcal{N}_\perp$ and there exists $l \sqsubseteq \mathit{url\_label}(u)$ such that $\mathit{path}(u) \in \mathcal{N}_l$.

Definition 6 (Well-formed Trace)**.** *An input event i is* well-formed *if and only if the judgement $\vdash_\diamond i$ can be proved through the following inference rules:*

(WF-LOAD)
$$\frac{\vdash_\diamond u}{\vdash_\diamond \mathsf{load}(u)}$$

(WF-TEXT)
$$\frac{n \in \mathcal{N}_{\rho(n)}}{\vdash_\diamond \mathsf{text}(p, k, n)}$$

(WF-XHR)
$$\frac{\begin{array}{c} l = \mathit{url\_label}(u) \qquad \mathit{ck\_vals}(ck) \subseteq \mathcal{N}_l \qquad \vdash_\diamond u \qquad \vdash_\diamond u' \\ \exists l' \sqsubseteq l : \mathit{path}(u') \in \mathcal{N}_{l'} \qquad \mathit{dom}(ck) \cup \mathit{fn}(v) \cup \{n\} \subseteq \mathcal{N}_\perp \end{array}}{\vdash_\diamond \mathsf{xhr\_resp}(n, ck, u, u', v)}$$

(WF-DOC)
$$\frac{\begin{array}{c} l = \mathit{url\_label}(u) \qquad \mathit{ck\_vals}(ck) \subseteq \mathcal{N}_l \\ \vdash_\diamond u \qquad \vdash_\diamond u' \qquad \exists l' \sqsubseteq l : \mathit{path}(u') \in \mathcal{N}_{l'} \\ \mathit{dom}(ck) \cup \mathit{fn}(h) \cup \mathit{fn}(e) \cup \{n\} \subseteq \mathcal{N}_\perp \end{array}}{\vdash_\diamond \mathsf{doc\_resp}(n, ck, u, u', h, e)}$$

We say that a trace (I, O) is well-formed iff so is every $i \in I$.

An explanation of the rules follows: rule (WF-LOAD) ensures that the user never types in the address bar a URL containing a password (or an authentication cookie value) which should not be disclosed to the remote server. Rule (WF-TEXT) rules out text inputs containing names corresponding to authentication cookie values: in

TABLE 4 Reactive semantics of FF^{+}: outputs

(O-APP)
$$\frac{}{\langle W, K, N, \{p \mapsto ((\lambda x.e)\ v, l)\}, [\,]\rangle \overset{\bullet}{\mapsto} \langle W, K, N, \{p \mapsto (e\{v/x\}, l)\}, [\,]\rangle}$$

(O-LETCTX)
$$\frac{\langle W, K, N, \{p \mapsto (e', l)\}, [\,]\rangle \overset{o}{\mapsto} \langle W', K', N', \{p \mapsto (e'', l)\}, [\,]\rangle}{\langle W, K, N, \{p \mapsto (\mathsf{let}\ x = e'\ \mathsf{in}\ e, l)\}, [\,]\rangle \overset{o}{\mapsto} \langle W', K', N', \{p \mapsto (\mathsf{let}\ x = e''\ \mathsf{in}\ e, l)\}, [\,]\rangle}$$

(O-LET)
$$\frac{}{\langle W, K, N, \{p \mapsto (\mathsf{let}\ x = v\ \mathsf{in}\ e, l)\}, [\,]\rangle \overset{\bullet}{\mapsto} \langle W, K, N, \{p \mapsto (e\{v/x\}, l)\}, [\,]\rangle}$$

(O-GET)
$$\frac{W(p) = (u, h, h', q) \qquad d = \mathit{domain}(u) \qquad \exists ck : K(d) = ck \wedge ck(k) = (n, f) \wedge f \in \{\bot, \mathtt{S}\}}{\langle W, K, N, \{p \mapsto (k?, l)\}, [\,]\rangle \overset{\bullet}{\mapsto} \langle W, K, N, \{p \mapsto (n, l)\}, [\,]\rangle}$$

(O-GETFAIL)
$$\frac{W(p) = (u, h, h', q) \qquad d = \mathit{domain}(u) \qquad \neg\exists ck : K(d) = ck \wedge ck(k) = (n, f) \wedge f \in \{\bot, \mathtt{S}\}}{\langle W, K, N, \{p \mapsto (k?, l)\}, [\,]\rangle \overset{\bullet}{\mapsto} \langle W, K, N, \{p \mapsto ((), l)\}, [\,]\rangle}$$

(O-SET)
$$\frac{\begin{array}{c} W(p) = (u, h, h', q) \qquad d = \mathit{domain}(u) \\ \neg\exists ck : K(d) = ck \wedge ck(k) = (m, f') \wedge f' \in \{\mathtt{H}, \top\} \qquad K' = \mathit{upd\_ck}(K, d, \{k \mapsto (n, f)\}) \end{array}}{\langle W, K, N, \{p \mapsto (k!\langle n, f\rangle, \bot)\}, [\,]\rangle \overset{\bullet}{\mapsto} \langle W, K', N, \{p \mapsto ((), \bot)\}, [\,]\rangle}$$

(O-SETFAIL)
$$\frac{W(p) = (u, h, h', q) \qquad d = \mathit{domain}(u) \qquad l \neq \bot \vee (\exists ck : K(d) = ck \wedge ck(k) = (m, f') \wedge f' \in \{\mathtt{H}, \top\})}{\langle W, K, N, \{p \mapsto (k!\langle n, f\rangle, l)\}, [\,]\rangle \overset{\bullet}{\mapsto} \langle W, K, N, \{p \mapsto ((), l)\}, [\,]\rangle}$$

(O-XHR)
$$\frac{\begin{array}{c} W' = W \uplus \{p \mapsto (u', h, h' \uplus \{n \mapsto \lambda x.e\}, q)\} \\ l \neq \bot \Rightarrow l = \mathit{url\_label}(u) = \mathit{url\_label}(u') \wedge q = \checkmark \\ q = \checkmark \wedge \mathit{url\_label}(u) = \mathit{url\_label}(u') \Rightarrow ck = \mathit{get\_http\_ck}(K, u) \wedge q' = \checkmark \\ q = \times \vee \mathit{url\_label}(u) \neq \mathit{url\_label}(u') \Rightarrow ck = \{\} \wedge q' = \times \end{array}}{\langle W \uplus \{p \mapsto (u', h, h', q)\}, K, N, \{p \mapsto (\mathsf{xhr}(u, \lambda x.e), l)\}, [\,]\rangle \xmapsto{\mathsf{xhr\_req}(ck, u)} \langle W', K, N \uplus \{n \mapsto (u, p, q')\}, \{p \mapsto ((), l)\}, [\,]\rangle}$$

(O-LOGIN)
$$\frac{W(p) = (u', h, h', \checkmark) \qquad \rho(c) = \mathit{url\_label}(u) \qquad l = \mathit{url\_label}(u) = \mathit{url\_label}(u') \qquad ck = \mathit{get\_http\_ck}(K, u)}{\langle W, K, N, \{p \mapsto (\mathsf{auth}(u, c), l)\}, [\,]\rangle \xmapsto{\mathsf{login}(ck, u, c)} \langle W, K, N \uplus \{n \mapsto (u, (), \checkmark)\}, \{p \mapsto ((), l)\}, [\,]\rangle}$$

(O-FLUSH)
$$\frac{}{\langle W, K, N, T, o\rangle \overset{o}{\mapsto} \langle W, K, N, T, [\,]\rangle}$$

(O-MIRROR)
$$\frac{P \overset{o}{\mapsto} Q}{P \overset{o}{\to} Q}$$

(O-COMPLETE)
$$\frac{\langle W, K, N, \{p \mapsto (e, l)\}, [\,]\rangle \not\mapsto}{\langle W, K, N, \{p \mapsto (e, l)\}, [\,]\rangle \overset{\bullet}{\to} \langle W, K, N, \{\}, [\,]\rangle}$$

Notation: we write $P \not\mapsto$ whenever there do not exist o and Q such that $P \overset{o}{\mapsto} Q$.

other words, we assume that the user is always entering either a password or some public data. Rules (WF-DOC) and (WF-XHR) ensure that cookies set by a honest server are picked from the correct name partition and only occur in the standard HTTP header: furthermore, we require that confidential data (e.g., passwords) never appear in the body of a response or in the cookie names.

Theorem 1 (Session Integrity)**.** FF^+ *enforces session integrity for any well-formed trace (with respect to the threat model in Section III-C).*

The proof draws on a label-indexed family of *simulation* relations, which connect the attacked trace with the original one. The proof is challenging, due to the significant differences which may arise between the two traces: full details are in Appendix C.

IV. Enforcing Session Integrity in SessInt

In this section, we discuss how to transfer FF^+ provable security into real browser security. To accomplish that, the following aspects must be taken in due account. First, the implementation of the required protection mechanisms should be designed so as to minimize their impact on the user experience: this is a difficult task, which requires careful design based on the search of the best possible trade-off between security and usability. Second, the design should lend itself to an implementation as a browser extension, to ease its deployment. When that is not possible, it is important to fine-tune the proposed security mechanisms so that they are still consistent with the theoretical model.

Below, we report on our development of SessInt, a proof-of-concept implementation of the integrity mechanisms of FF^+ as a browser extension for Google Chrome. While the current design is targeted at Chrome, and depends on its extension API, the same development appears possible on other major web browsers.

A. Implementing FF^+ security

We start by discussing how the different browsing events correspond to FF^+ events, and are handled by SessInt accordingly.

1) Address bar: Typing a URL in the address bar and loading a page corresponds to the load event in FF^+, i.e., we trust what the user types in the address bar. Here we have a first problem due to Chrome's API, which does not provide enough information to distinguish between a request triggered by the user typing in the address bar and a redirect, possibly caused by Javascript. For the moment, we have solved the issue by letting the user add a special character 'g' before inserting the URL, so that we can capture this input via the Chrome omnibox API and detect, accordingly, that the URL has been typed in the address bar.

2) User clicks: Following a link via a click is mapped to a xhr operation of FF^+. The rationale is that we do not trust clicks as, in fact, they might have been performed by malicious Javascript code. Moreover, it is unrealistic to assume that the user carefully checks every single link before clicking on it. Even though a user click could correspond to a load event, we decided to treat it as an xhr_req event and apply a more conservative security policy, whereby SessInt strips all authentication cookies before sending cross-origin requests, so as to prevent cross-origin request forgeries from malicious scripts.

3) Implicit loads: Implicit loads from a page or script correspond to xhr operations in FF^+, since we cannot trust these events. Hence, SessInt strips all authentication cookies before sending cross-origin requests.

4) Passwords: In order to prevent passwords from being leaked by malicious Javascript code, we sandbox login forms into an isolated popup. SessInt implements a password manager, which checks that the input password is correct before sending it. If the password is not yet in the password manager, the user is asked for confirmation and, in case of a positive answer, the password is stored and associated to the page and action URLs, to enforce the runtime discipline adopted by FF^+. Notice that the Chrome API does not allow to inspect the page content before inline scripts are executed. However, if these scripts modify the action URL before the extension creates the sandboxed form, the password manager will detect it by a comparison with the stored action URL and will warn the user before proceeding.

5) Cookies: SessInt performs a taint tracking over the open network connections, exactly as FF^+: cookies are only updated when they are received over an untainted connection. SessInt marks any authentication cookie received by the browser on HTTP connections as `HttpOnly`; authentication cookies received over HTTPS, instead, are marked as both `HttpOnly` and `Secure`. This prevents leakage from malicious Javascript programs and protects cookies in case HTTP links are injected into HTTPS websites. To preserve functionality, SessInt forces a redirection on HTTPS for the entire website when a login form is submitted over HTTPS: indeed, if the website contains some hard-coded HTTP links, marking some authentication cookies as `Secure` would break the session when navigating these links. As done by other extensions [15], [20], [33], authentication cookies are detected based on standard naming conventions (e.g., `PHPSESSID`) and a heuristics that measures the degree of entropy of the cookie value. In a recent paper we show how the authentication cookie detection process

can be improved significantly using machine learning techniques [16].

B. Protection vs usability

There are a few situations where the security policy of FF^+ would break too many websites, hence we have to slightly relax it in SESSINT. We discuss these situations and the implications on browser security and usability.

1) HTTP sessions: Some websites only support HTTPS for a subset of their pages. If some portions of the website do not provide support for HTTPS, SESSINT selectively allows a fallback to HTTP, with the proviso that cookies which have been previously promoted to `Secure` by the extension must be included to preserve the session [15]. If a HTTPS connection times out, we do not force the fallback, to prevent a network attacker intercepting HTTPS traffic from forcing SESSINT into leaking over HTTP authentication cookies of websites normally providing HTTPS support. Of course, we cannot provide session security against network attacks for websites which only partially support HTTPS, and the user is warned when this is the case.

2) Redirection to HTTPS: Many websites redirect the browser to HTTPS when an HTTP access is requested. However, SESSINT would not include authentication cookies upon these HTTPS redirections, since these redirects could as well be exploited by a network attacker to point the browser to a sensitive HTTPS URL and carry out a forgery: this cookie stripping breaks many websites, e.g., Facebook. To regain functionality, in the specific case of a protocol upgrading *with unmodified URL*, the user is asked (once for each site) to confirm that the redirection is expected, so that authentication cookies can be sent to the website. If the redirection looks suspicious, the user can block it.

3) HTTPS login forms into HTTP pages: It is common to find websites where HTTPS login forms are embedded (e.g., as iframes) into HTTP pages. This is insecure, as an attacker can change the HTTP page so as to redirect forms to a server that he controls, but it is a very common practice and we need to let it work. Our choice is to warn the user when this happens, then the password manager will give an extra warning in case the password is going to be sent to a URL that is not yet known. The combination of the two warnings should make the user well aware of a possible attack.

4) Subdomains and external sites: It is common that secure sessions link to subdomains or to external sites as, e.g., in e-payments. Navigating to a different domain would normally strip authentication cookies. To avoid breaking websites, SESSINT by default sends authentication cookies when moving into a subdomain, even though this could sometimes be exploited by a web attacker with scripting capabilities in the subdomain [14], [20]. We are investigating how to extend this behaviour to external (trusted) websites. A simple idea might be to include a white-list of trusted sites, e.g., for e-payments, that are needed to be reached by other websites, so that when the navigation comes back to the original site, authentication cookies are correctly sent and the session is preserved. We also plan to study to which extent we can engineer in SESSINT previous proposals aimed at supporting useful collaborative web scenarios [19].

C. Experiments

We tested SESSINT on existing vulnerable web applications, such as OWASP Mutillidae and Damn Vulnerable Web Application. We believe this is important to confirm that the more relaxed security policy adopted in SESSINT does not sacrifice too much of the bullet-proof security of FF^+. Here, we report some simple, but significant examples of attacks prevented by SESSINT:

1) CSRF: A link to domain A from a different domain B, that performs an action inside an active session with A (cf. Figure 1 (a)). With SESSINT authentication cookies are stripped and the action has no effect.

2) Cookie stealing via XSS: An XSS attack can access the JavaScript object `document.cookie` and leak an authentication cookie (cf. Figure 1 (b)). With SESSINT all authentication cookies are set `HttpOnly` and the attack is prevented.

3) Local CSRF: Domain A is vulnerable to XSS. The attacker injects a payload into A that redirects to a location, still in A, that performs an action inside the session (cf. Figure 1 (c)). The attack is prevented by SESSINT, since authentication cookies are stripped when a redirection happens over a tainted connection.

V. CONCLUSION

We introduced a novel notion of web session integrity and we showed that our definition is both general and amenable for client-side enforcement. We then proposed FF^+, a security-enhanced model of a web browser that provides a full-fledged and provably sound enforcement of web session integrity. Based on that, we developed SESSINT, a proof-of-concept browser extension which implements the security checks formalized in FF^+. We discussed the effectiveness of our solution and we presented some design choices we made to foster usability.

As a future work, we would like to further engineer SESSINT, trying to support more complicated collaborative web scenarios. We also plan to investigate how to enforce web session integrity in a browser supporting information flow control policies like FlowFox [24].

Acknowledgements: We would like to thank the anonymous referees for their valuable comments and our shepherd Tamara Rezk for her guidance in improving the original submission of this paper. This work was partially supported by the following MIUR Projects: PON ADAPT, PRIN CINA and PRIN Security Horizons.

REFERENCES

[1] HTTP state management mechanism. http://tools.ietf.org/html/rfc6265.

[2] Provably Sound Browser-Based Enforcement of Web Session Integrity (full version). http://www.dais.unive.it/~calzavara/csf14-full.pdf.

[3] The Transport Layer Security (TLS) protocol. http://tools.ietf.org/html/rfc4346.

[4] B. Adida. Sessionlock: securing web sessions against eavesdropping. In *International Conference on the World Wide Web (WWW)*, pages 517–524, 2008.

[5] D. Akhawe, A. Barth, P. E. Lam, J. C. Mitchell, and D. Song. Towards a formal foundation of web security. In *IEEE Computer Security Foundations Symposium (CSF)*, pages 290–304, 2010.

[6] A. Armando, R. Carbone, L. Compagna, J. Cuéllar, G. Pellegrino, and A. Sorniotti. An authentication flaw in browser-based single sign-on protocols: Impact and remediations. *Computers & Security*, 33:41–58, 2013.

[7] 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 *Formal Methods in Security Engineering (FMSE)*, pages 1–10, 2008.

[8] C. Bansal, K. Bhargavan, A. Delignat-Lavaud, and S. Maffeis. Keys to the cloud: Formal analysis and concrete attacks on encrypted web storage. In *Principles of Security and Trust (POST)*, pages 126–146, 2013.

[9] C. Bansal, K. Bhargavan, and S. Maffeis. Discovering concrete attacks on website authorization by formal analysis. In *IEEE Computer Security Foundations Symposium (CSF)*, pages 247–262, 2012.

[10] A. Barth, C. Jackson, and J. C. Mitchell. Robust defenses for cross-site request forgery. In *ACM Conference on Computer and Communications Security (CCS)*, pages 75–88, 2008.

[11] A. Bohannon. *Foundations of webscript security*. PhD thesis, University of Pennsylvania, 2012.

[12] A. Bohannon and B. C. Pierce. Featherweight Firefox: formalizing the core of a web browser. In *USENIX Conference on Web Application Development (WebApps)*, pages 1–12, 2010.

[13] A. Bohannon, B. C. Pierce, V. Sjöberg, S. Weirich, and S. Zdancewic. Reactive noninterference. In *ACM Conference on Computer and Communications Security (CCS)*, pages 79–90, 2009.

[14] A. Bortz, A. Barth, and A. Czeskis. Origin cookies: Session integrity for web applications. In *Web 2.0 Security & Privacy (W2SP)*, 2011.

[15] M. Bugliesi, S. Calzavara, R. Focardi, and W. Khan. Automatic and robust client-side protection for cookie-based sessions. In *Engineering Secure Software and Systems (ESSoS)*, pages 161–178, 2014.

[16] S. Calzavara, G. Tolomei, M. Bugliesi, and S. Orlando. Quite a mess in my cookie jar! Leveraging machine learning to protect web authentication. In *International Conference on World Wide Web (WWW)*, pages 189–200, 2014.

[17] I. Dacosta, S. Chakradeo, M. Ahamad, and P. Traynor. One-time cookies: Preventing session hijacking attacks with stateless authentication tokens. *ACM Transactions on Internet Technology*, 12(1):1, 2012.

[18] P. De Ryck, L. Desmet, T. Heyman, F. Piessens, and W. Joosen. Csfire: Transparent client-side mitigation of malicious cross-domain requests. In *Engineering Secure Software and Systems (ESSoS)*, pages 18–34, 2010.

[19] P. De Ryck, L. Desmet, W. Joosen, and F. Piessens. Automatic and precise client-side protection against CSRF attacks. In *European Symposium on Research in Computer Security (ESORICS)*, pages 100–116, 2011.

[20] P. De Ryck, N. Nikiforakis, L. Desmet, F. Piessens, and W. Joosen. Serene: Self-reliant client-side protection against session fixation. In *Distributed Applications and Interoperable Systems (DAIS)*, pages 59–72, 2012.

[21] S. Fogie, J. Grossman, R. Hansen, A. Rager, and P. D. Petkov. *XSS Attacks: Cross Site Scripting Exploits and Defense*. Syngress Publishing, 2007.

[22] C. Fournet and T. Rezk. Cryptographically sound implementations for typed information-flow security. In *Principles of Programming Languages (POPL)*, pages 323–335, 2008.

[23] B. S. Y. Fung and P. P. C. Lee. A privacy-preserving defense mechanism against request forgery attacks. In *International Conference on Trust, Security and Privacy in Computing and Communications (TRUSTCOM)*, pages 45–52, 2011.

[24] W. D. Groef, D. Devriese, N. Nikiforakis, and F. Piessens. FlowFox: a web browser with flexible and precise information flow control. In *ACM Conference on Computer and Communications Security (CCS)*, pages 748–759, 2012.

[25] P. A. Hallgren, D. T. Mauritzson, and A. Sabelfeld. Glasstube: A lightweight approach to web application integrity. In *Workshop on Programming Languages and Analysis for Security (PLAS)*, pages 71–82, 2013.

[26] C. Jackson and A. Barth. ForceHTTPS: protecting high-security web sites from network attacks. In *International Conference on World Wide Web (WWW)*, pages 525–534, 2008.

[27] M. Johns, B. Braun, M. Schrank, and J. Posegga. Reliable protection against session fixation attacks. In *ACM Symposium on Applied Computing (SAC)*, pages 1531–1537, 2011.

[28] M. Johns, S. Lekies, B. Braun, and B. Flesch. BetterAuth: web authentication revisited. In *Annual Computer Security Applications Conference (ACSAC)*, pages 169–178, 2012.

[29] M. Johns and J. Winter. RequestRodeo: client side protection against session riding. *Proceedings of the OWASP Europe Conference*, pages 5–17, 2006.

[30] E. Kirda, C. Krügel, G. Vigna, and N. Jovanovic. Noxes: a client-side solution for mitigating cross-site scripting attacks. In *ACM Symposium on Applied Computing (SAC)*, pages 330–337, 2006.

[31] Z. Mao, N. Li, and I. Molloy. Defeating cross-site request forgery attacks with browser-enforced authenticity protection. In *Financial Cryptography (FC)*, pages 238–255, 2009.

[32] A. C. Myers, A. Sabelfeld, and S. Zdancewic. Enforcing robust declassification and qualified robustness. *Journal of Computer Security*, 14(2):157–196, 2006.

[33] N. Nikiforakis, W. Meert, Y. Younan, M. Johns, and W. Joosen. SessionShield: Lightweight protection against session hijacking. In *Engineering Secure Software and Systems (ESSoS)*, pages 87–100, 2011.

[34] N. Nikiforakis, Y. Younan, and W. Joosen. Hproxy: Client-side detection of ssl stripping attacks. In *International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment (DIMVA)*, pages 200–218, 2010.

[35] S. Tang, N. Dautenhahn, and S. T. King. Fortifying web-based applications automatically. In *ACM Conference on Computer and Communications Security (CCS)*, pages 615–626, 2011.

[36] Y. Zhou and D. Evans. Why aren't HTTP-Only cookies more widely deployed. In *Web 2.0 Security and Privacy Workshop (W2SP'10)*, 2010.

APPENDIX A
ENCODING ATTACKS

Here, we include three additional attacks which are captured by our definition of session integrity.

Password theft (Figure 2 (a)): In this scenario, the browser requests a login page over an HTTP connection to S. In the unattacked trace, S would respond with the page and the trace would be concluded with the authentication step, where the password is sent to S over an HTTPS connection. In the attacked trace, instead, the attacker at E intercepts the login page on HTTP and responds to the browser with a fake page of its own, masquerading as S. As a result, the attacker may steal the S-level password and start its own authenticated session with S, thus violating the integrity condition for the trace. Notice that the attack is not reported as a confidentiality leak (when the password is inadvertently passed to E), but rather as an integrity violation that arises from E using *pwd* to start a new session on behalf of the user.

Login CSRF (Figure 2 (b)): Again, the trace starts with the browser authenticating with S, and continues with a request for a page at E in a new browser tab. Later on, the user enters a secondary password *xyz* for future accesses to S, that is stored in the clear. In the unattacked trace, this last step would include the cookie c set by S and thus store the credentials on the user's account at S. In the attacked trace, instead, the attacker at E forces the browser to silently authenticate at S with the *attacker's* password, starting his own session associated with the new cookie $\hat{c} : \mathsf{E}$. At this stage, the subsequent request by the browser includes this new cookie registered in the browser, thus continuing the attacker session at S (rather than resuming the intended user session). As a result, the user's password *xyz* is stored at the attacker's account, who may later use it to start a new session at S on behalf of the user. Again, it is this last step (rather than the leakage of the user's credentials) that breaks the integrity condition on the attacked trace.

Session fixation (Figure 2 (c)): The attacker at E injects a malicious script on S through an XSS vulnerability, which registers in the browser a cookie $\hat{c}$ chosen by the attacker. This cookie is not refreshed by S when the user authenticates, rather it is endorsed by the user password and grants access to the session associated to the user's credentials. The attacker can then arbitrarily use $\hat{c}$ to hijack the user's session, violating the integrity condition. Without the attacker intervention, no redirection would have occurred in the trace, and the login step would not have included any cookie. Clearly, the problem would be easily rectified had S refreshed the cookie upon receiving the credential *pwd*.

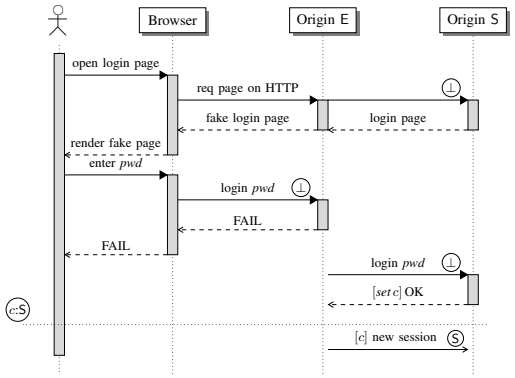

(a) Password theft (*pwd* : S)

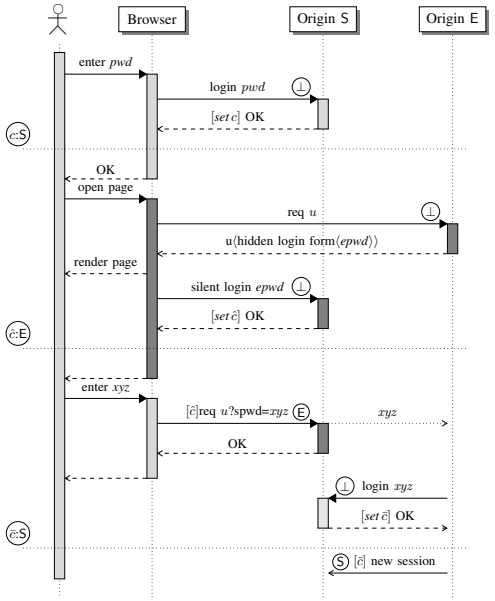

(b) Login CSRF (*epwd* : E, *pwd* : S, *xyz* : S)

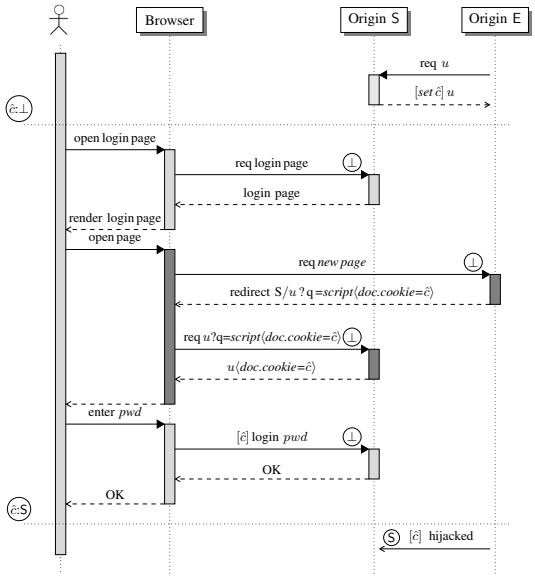

(c) Session fixation (*pwd* : S)

Figure 2: Other violations of session integrity

We now provide a formal encoding of the attacks to web session integrity presented in the paper. To keep the notation lighter, we omit from the syntax of events the network connection identifiers (in doc_resp and xhr_resp events) and the page identifiers (in text events). We use the underscore '_' to stand for syntactic elements which are not strictly needed to understand the examples. When denoting attacked traces, we make explicit the input stream which is actually consumed by the browser: this is a simple representation of what the attacker is doing along the derivation corresponding to the attacked trace semantics. This information is precisely tracked in the small-step semantics we define to carry out our proofs (see Appendix C).

Cross-Site Request Forgery (CSRF): Let $\mathsf{S} = \mathsf{https}(d_1)$ be the honest origin. The unattacked trace (I, O) up to dummy outputs can be encoded as follows:

$$\begin{array}{rcl} I & = & [\mathsf{load}(u_1), \mathsf{doc\_resp}(\{\}, u_1, \mathsf{blank}, \{k \mapsto \lambda x.\mathsf{auth}(u_1', x)\}, \_), \mathsf{text}(k, pwd), \mathsf{doc\_resp}(ck_1, \_), \mathsf{load}(u_2)] \\ O & = & [\mathsf{doc\_req}(\{\}, u_1), \mathsf{login}(\{\}, u_1', pwd), \mathsf{doc\_req}(\{\}, u_2)] \end{array}$$

Assume that ck_1 is the authentication cookie chosen by S and let $\mathsf{E} = \mathsf{http}(d_2)$ be a web attacker, where d_2 is the domain of u_2. The attacked trace (E, I', O') looks as follows:

$$\begin{array}{rcl} I' & = & [I, \mathsf{doc\_resp}(\{\}, u_2, \mathsf{blank}, \_, \mathsf{xhr}(u_1'', \_))] \\ O' & = & [O, \mathsf{xhr\_req}(ck_1, u_1'')] \end{array}$$

Let τ' be defined as $\tau_\perp$ updated after the login event in O, which assigns to the output event $\mathsf{xhr\_req}(ck_1, u_1'')$ a trust level of S. Since the opponent has a lower level E, this output event leads to a security violation.

Reflected XSS: Let $\mathsf{S} = \mathsf{https}(d_1)$ be the honest origin and let $ck_1 = \{k_1 \mapsto (n, \mathsf{S})\}$ be the authentication cookie chosen by S. The unattacked trace (I, O) up to dummy outputs can be encoded as follows:

$$\begin{array}{rcl} I & = & [\mathsf{load}(u_1), \mathsf{doc\_resp}(\{\}, u_1, \mathsf{blank}, \{k \mapsto \mathsf{auth}(u_1', x)\}, \_), \mathsf{text}(k, pwd), \mathsf{doc\_resp}(ck_1, u_1', \mathsf{blank}, \_, \_), \\ & & \mathsf{load}(u_2), \mathsf{doc\_resp}(\{\}, u_1'', \mathsf{blank}, \_, \mathsf{let}\ x = k_1?\ \mathsf{in}\ \mathsf{xhr}((\mathsf{http}, d_2, x), \_))] \\ O & = & [\mathsf{doc\_req}(\{\}, u_1), \mathsf{login}(\{\}, u_1', pwd), \mathsf{doc\_req}(\{\}, u_2)] \end{array}$$

Notice that I contains a high input from u_1'', which delivers a script stealing the authentication cookie from the browser: this input produces a dummy output in the unattacked trace, since no request has been made to u_1'', but the attacker can leverage it in the attacked run. This is precisely the way we model XSS vulnerabilities.

Let $\mathsf{E} = \mathsf{http}(d_2)$ be a web attacker, where d_2 is the domain of u_2. The attacked trace (E, I', O') looks as follows:

$$\begin{array}{rcl} I' & = & [\mathsf{load}(u_1), \mathsf{doc\_resp}(\{\}, u_1, \mathsf{blank}, \{k \mapsto \mathsf{auth}(u_1', x)\}), \mathsf{text}(k, pwd), \mathsf{doc\_resp}(ck_1, u_1', \mathsf{blank}, \_, \_), \\ & & \mathsf{load}(u_2), \mathsf{doc\_resp}(\{\}, u_2, u_1'', \_, \_), \mathsf{doc\_resp}(\{\}, u_1'', \mathsf{blank}, \_, \mathsf{let}\ x = k_1?\ \mathsf{in}\ \mathsf{xhr}((\mathsf{http}, d_2, x), \_))] \\ O' & = & [O, \mathsf{doc\_req}(\{\}, u_1''), \mathsf{xhr\_req}(\{\}, (\mathsf{http}, d_2, n)), \mathsf{doc\_req}(ck_1, u_1)] \end{array}$$

The previous output stream leads to a violation of our security notion, since the event $\mathsf{doc\_req}(ck_1, u_1)$ has a trust level of S after the login event in O, but the attacker has a lower level E.

Password theft: Let $\mathsf{S} = \mathsf{https}(d_1)$ be the honest origin. The unattacked trace (I, O) up to dummy outputs can be encoded as follows:

$$\begin{array}{rcl} I & = & [\mathsf{load}(u_1), \mathsf{doc\_resp}(\{\}, u_1, \mathsf{blank}, \{k \mapsto \mathsf{auth}(u_1', x)\}, \_), \mathsf{text}(k, pwd), \mathsf{doc\_resp}(ck_1, u_1', \mathsf{blank}, \_, \_)] \\ O & = & [\mathsf{doc\_req}(\{\}, u_1), \mathsf{login}(\{\}, u_1', pwd)] \end{array}$$

Assume the login page is sent in clear, i.e., let $u_1 = (\mathsf{http}, d_1, v_1)$, while the login form is sent encrypted, i.e., let $u_1' = (\mathsf{https}, d_1, v_1')$. This is a standard setup for many existing websites.

Let $\mathsf{E} = \mathsf{net}$ be a network attacker. The attacked trace (E, I', O') looks as follows:

$$\begin{array}{rcl} I' & = & [\mathsf{load}(u_1), \mathsf{doc\_resp}(\{\}, u_1, \mathsf{blank}, \{k \mapsto \mathsf{auth}((\mathsf{http}, d_2, \_), x)\}, \_), \mathsf{text}(k, pwd)] \\ O' & = & [\mathsf{doc\_req}(\{\}, u_1), \mathsf{login}(\{\}, (\mathsf{http}, d_2, \_), pwd), \mathsf{login}(ck_2, u_1'', pwd), \mathsf{doc\_req}(ck_2, u_1')] \end{array}$$

In the attacked trace we leverage rule (A-FIX) to let the attacker choose a known cookie ck_2 and endorse it with the user's password pwd, which was previously communicated over HTTP. (This is convenient for modelling the case, even though in practice this cookie is not necessarily fixated by the attacker.) The last output event in O' breaks session integrity.

Login CSRF: Let $\mathsf{S} = \mathsf{https}(d_1)$ be the honest origin and let $u_1'' = (\mathsf{https}, d_1, xyz)$ be the URL encoding the request to store the secondary password xyz on the honest server. The unattacked trace (I, O) up to dummy outputs can be encoded as follows:

$$\begin{array}{rcl} I & = & [\mathsf{load}(u_1), \mathsf{doc\_resp}(\{\}, u_1, \mathsf{blank}, \{k \mapsto \mathsf{auth}(u_1', x)\}, \_), \mathsf{text}(k, pwd), \mathsf{doc\_resp}(ck_1, u_1', \mathsf{blank}, \_, \_), \\ & & \mathsf{load}(u_2), \mathsf{doc\_resp}(ck_2, u_1', \mathsf{blank}, \_, \_), \mathsf{load}(u_1'')] \\ O & = & [\mathsf{doc\_req}(\{\}, u_1), \mathsf{login}(\{\}, u_1', pwd), \mathsf{doc\_req}(\{\}, u_2), \mathsf{doc\_req}(ck_1, u_1'')] \end{array}$$

Notice that the last doc_resp event in I has no import in the unattacked trace, since it produces a dummy output, but it will be leveraged by the attacker in the attacked trace. In particular, the attacker will force the browser into overwriting the authentication cookie $ck_1 = \{k' \mapsto (n_1, f_1)\}$ associated to the user's credentials with the new cookie $ck_2 = \{k' \mapsto (n_2, f_2)\}$ associated to the attacker's credentials.

Let $evil$ be the attacker's password and let $\mathsf{E} = \mathsf{http}(d_2)$ be a web attacker, where d_2 is the domain of u_2. The attacked trace (E, I', O') looks as follows:

$$\begin{array}{rcl} I' & = & [\mathsf{load}(u_1), \mathsf{doc\_resp}(\{\}, u_1, \mathsf{blank}, \{k \mapsto \mathsf{auth}(u_1', x)\}, \_), \mathsf{text}(k, pwd), \mathsf{doc\_resp}(ck_1, u_1', \mathsf{blank}, \_, \_), \\ & & \mathsf{load}(u_2), \mathsf{doc\_resp}(\{\}, u_2, \mathsf{blank}, \_, \mathsf{auth}(u_1', evil)), \mathsf{doc\_resp}(ck_2, u_1', \mathsf{blank}, \_, \_), \mathsf{load}(u_1'')] \\ O' & = & [\mathsf{doc\_req}(\{\}, u_1), \mathsf{login}(\{\}, u_1', pwd), \mathsf{doc\_req}(\{\}, u_2), \mathsf{login}(\{\}, u_1', evil), \mathsf{doc\_req}(ck_2, u_1''), \\ & & \mathsf{login}(ck', \hat{u}_1, xyz), \mathsf{doc\_req}(ck', u_1)] \end{array}$$

The second login event in O' grants trust evil to network requests including the cookie ck_2, which is set by the input event $\mathsf{doc\_resp}(ck_2, u_1', \mathsf{blank}, \_, \_)$. Hence, the output event $\mathsf{doc\_req}(ck_2, u_1'')$ leaks the secondary password xyz and the intruder can perform $\mathsf{login}(ck', \hat{u}_1, xyz)$ using (A-FIX). This login event grants trust S to $\mathsf{doc\_req}(ck', u_1)$, which eventually leads to a security violation.

Session fixation: Let $\mathsf{S} = \mathsf{https}(d_1)$ be the honest origin. The unattacked trace (I, O) up to dummy outputs can be encoded as follows:

$$\begin{array}{rcl} I & = & [\mathsf{load}(u_1), \mathsf{doc\_resp}(ck_1, u_1, \mathsf{blank}, \{k \mapsto \mathsf{auth}(u_1', x)\}, \_), \mathsf{load}(u_2), \mathsf{doc\_resp}(\{\}, u_1', \mathsf{blank}, \_, k'!\langle m, \bot\rangle), \\ & & \mathsf{text}(k, pwd), \mathsf{doc\_resp}(\{\}, u_1', \mathsf{blank}, \_, \_), \mathsf{load}(u_1'')] \\ O & = & [\mathsf{doc\_req}(\{\}, u_1), \mathsf{doc\_req}(\{\}, u_2), \mathsf{login}(ck_1, u_1', pwd), \mathsf{doc\_req}(ck_1, u_1'')] \end{array}$$

The input event $\mathsf{doc\_resp}(\{\}, u_1', \mathsf{blank}, \_, k'!\langle m, \bot\rangle)$ has no import in the unattacked trace, but it will be leveraged by the attacker to overwrite the original cookie $ck_1 = \{k' \mapsto (n, \bot)\}$ with the fixated cookie $ck_2 = \{k' \mapsto (m, \bot\}$.

Let $\mathsf{E} = \mathsf{http}(d_2)$ be a web attacker, where d_2 is the domain of u_2. The attacked trace (E, I', O') looks as follows:

$$\begin{array}{rcl} I' & = & [\mathsf{load}(u_1), \mathsf{doc\_resp}(ck_1, u_1, \mathsf{blank}, \{k \mapsto \mathsf{auth}(u_1', x)\}, \_), \mathsf{load}(u_2), \mathsf{doc\_resp}(\{\}, u_2, u_1', \_, \_), \\ & & \mathsf{doc\_resp}(\{\}, u_1', \mathsf{blank}, \_, k'!\langle m, \bot\rangle, \mathsf{text}(k, pwd), \mathsf{doc\_resp}(\{\}, u_1', \mathsf{blank}, \_, \_), \mathsf{load}(u_1'')] \\ O' & = & [\mathsf{doc\_req}(\{\}, u_1), \mathsf{doc\_req}(\{\}, u_2), \mathsf{doc\_req}(ck_1, u_1'), \mathsf{login}(ck_2, u_1', pwd), \mathsf{doc\_req}(ck_2, u_1'')] \end{array}$$

The login event in O' grants trust S_1 to ck_2 and the last output event breaks session integrity.

Local CSRF: Let $\mathsf{S} = \mathsf{https}(d_1)$ be the honest origin. The unattacked trace (I, O) up to dummy outputs can be encoded as follows:

$$\begin{array}{rcl} I & = & [\mathsf{load}(u_1), \mathsf{doc\_resp}(\{\}, u_1, \mathsf{blank}, \{k \mapsto \mathsf{auth}(u_1', x)\}, \_), \mathsf{text}(k, pwd), \mathsf{doc\_resp}(ck_1, u_1', \mathsf{blank}, \_, \_), \\ & & \mathsf{load}(u_2), \mathsf{doc\_resp}(\{\}, u_1'', \mathsf{blank}, \_, \mathsf{xhr}(u_1, \_))] \\ O & = & [\mathsf{doc\_req}(\{\}, u_1), \mathsf{login}(\{\}, u_1', pwd), \mathsf{doc\_req}(\{\}, u_2)] \end{array}$$

The attack starts like the reflected XSS: in particular, notice that the last doc_resp event in I has no import in the unattacked trace, but it will be fed to the browser in the attacked trace.

Let $\mathsf{E} = \mathsf{http}(d_2)$ be a web attacker, where d_2 is the domain of u_2. The attacked trace (E, I', O') looks as follows:

$$\begin{array}{rcl} I' & = & [\mathsf{load}(u_1), \mathsf{doc\_resp}(\{\}, u_1, \mathsf{blank}, \{k \mapsto \mathsf{auth}(u_1', x)\}), \mathsf{text}(k, pwd), \mathsf{doc\_resp}(ck_1, u_1', \mathsf{blank}, \_, \_), \\ & & \mathsf{load}(u_2), \mathsf{doc\_resp}(\{\}, u_2, u_1'', \_, \_), \mathsf{doc\_resp}(\{\}, u_1'', \mathsf{blank}, \_, \mathsf{xhr}(u_1, \_))] \\ O' & = & [O, \mathsf{doc\_req}(\{\}, u_1''), \mathsf{xhr\_req}(ck_1, u_1)] \end{array}$$

The last output event breaks session integrity.

APPENDIX B
FLYWEIGHT FIREFOX: FORMAL SEMANTICS

In this section we describe the full formal semantics of FF. We first introduce the standard cookie operations available in web browsers, then we give the reactive semantics for input and output events. The semantics rules in this section will only be used to show how standard web browsers fail at enforcing web session integrity: they are never referenced in the formal proofs, which instead are based on the semantics rules of FF^+.

Cookie operations: Given a cookie store K and a URL u, we define the result of the partial function $\mathit{get\_ck}(K, u)$ as the least map M such that:

$$M(k) = \begin{cases} (n, f) & \text{if } u = (\mathsf{https}, d, v) \text{ and } \exists ck : K(d) = ck \wedge ck(k) = (n, f) \\ (n, f) & \text{if } u = (\mathsf{http}, d, v) \text{ and } \exists ck : K(d) = ck \wedge ck(k) = (n, f) \wedge f \in \{\bot, \mathtt{H}\} \end{cases}$$

When $u = \mathsf{blank}$, we stipulate that $\mathit{get\_ck}(C, u)$ is not defined.

We also define a function to update cookies in a cookie store as follows:

$$\mathit{upd\_ck}(K, d, ck) = \begin{cases} K \uplus \{d \mapsto ck\} & \text{if } d \notin \mathit{dom}(K) \\ K' \uplus \{d \mapsto (ck' \lhd ck)\} & \text{if } K = K' \uplus \{d \mapsto ck'\} \end{cases}$$

Semantics: inputs: The reactive semantics for input events $C \xrightarrow{i} P$ is given in Table 5.

TABLE 5 Reactive semantics of FF: inputs

(I-LOAD)
$$\frac{ck = \mathit{get\_ck}(K, u)}{\langle W, K, N, \{\}, [\,]\rangle \xrightarrow{\mathsf{load}(u)} \langle W, K, N \uplus \{n \mapsto (u, ())\}, \{\}, \mathsf{doc\_req}(ck, u)\rangle}$$

(I-TEXT)
$$\frac{W(p) = (u, h, h') \qquad h(k) = \lambda x.e}{\langle W, K, N, \{\}, [\,]\rangle \xrightarrow{\mathsf{text}(p,k,n)} \langle W, K, N, \{p \mapsto e\{n/x\}\}, [\,]\rangle}$$

(I-DOCRESP)
$$\frac{d = \mathit{domain}(u) \qquad K' = \mathit{upd\_ck}(K, d, ck)}{\langle W, K, N \uplus \{n \mapsto (u, ())\}, \{\}, [\,]\rangle \xrightarrow{\mathsf{doc\_resp}(n,ck,u,\mathsf{blank},h,e)} \langle W \uplus \{p \mapsto (u, h, \{\})\}, K', N, \{p \mapsto e\}, [\,]\rangle}$$

(I-DOCREDIR)
$$\frac{d = \mathit{domain}(u) \qquad K' = \mathit{upd\_ck}(K, d, ck) \qquad ck' = \mathit{get\_ck}(K', u')}{\langle W, K, N \uplus \{n \mapsto (u, ())\}, \{\}, [\,]\rangle \xrightarrow{\mathsf{doc\_resp}(n,ck,u,u',h,e)} \langle W, K', N \uplus \{n \mapsto (u', ())\}, \{\}, \mathsf{doc\_req}(ck', u')\rangle}$$

(I-XHRRESP)
$$\frac{d = \mathit{domain}(u) \qquad K' = \mathit{upd\_ck}(K, d, ck) \qquad h' = h'' \uplus \{n \mapsto \lambda x.e\} \qquad W' = W \uplus \{p \mapsto (u', h, h'')\}}{\langle W \uplus \{p \mapsto (u', h, h')\}, K, N \uplus \{n \mapsto (u', p)\}, \{\}, [\,]\rangle \xrightarrow{\mathsf{xhr\_resp}(n,ck,u,\mathsf{blank},v)} \langle W', K', N', \{p \mapsto e\{v/x\}\}, [\,]\rangle}$$

(I-XHRREDIR)
$$\frac{d = \mathit{domain}(u) \qquad K' = \mathit{upd\_ck}(K, d, ck) \qquad ck' = \mathit{get\_ck}(K', u')}{\langle W, K, N \uplus \{n \mapsto (u, p)\}, \{\}, [\,]\rangle \xrightarrow{\mathsf{xhr\_resp}(n,ck,u,u',v)} \langle W, K', N \uplus \{n \mapsto (u', p)\}, \{\}, \mathsf{xhr\_req}(ck', u')\rangle}$$

(I-MIRROR)
$$\frac{C \overset{i}{\mapsto} P}{C \xrightarrow{i} P}$$

(I-COMPLETE)
$$\frac{\langle W, K, N, \{\}, [\,]\rangle \overset{i}{\not\mapsto}}{\langle W, K, N, \{\}, [\,]\rangle \xrightarrow{i} \langle W, K, N, \{\}, \bullet\rangle}$$

Notation: we write $C \overset{i}{\not\mapsto}$ whenever there does not exist P such that $C \overset{i}{\mapsto} P$.

Semantics: outputs: The reactive semantics for output events $P \xrightarrow{o} Q$ is given in Table 6.

TABLE 6 Reactive semantics of FF: outputs

(O-APP)
$$\frac{}{\langle W, K, N, \{p \mapsto (\lambda x.e)\ v\}, [\,]\rangle \overset{\bullet}{\mapsto} \langle W, K, N, \{p \mapsto e\{v/x\}\}, [\,]\rangle}$$

(O-LETCTX)
$$\frac{\langle W, K, N, \{p \mapsto e'\}, [\,]\rangle \overset{o}{\mapsto} \langle W', K', N', \{p \mapsto e''\}, [\,]\rangle}{\langle W, K, N, \{p \mapsto \mathsf{let}\ x = e'\ \mathsf{in}\ e\}, [\,]\rangle \overset{o}{\mapsto} \langle W', K', N', \{p \mapsto \mathsf{let}\ x = e''\ \mathsf{in}\ e\}, [\,]\rangle}$$

(O-LET)
$$\frac{}{\langle W, K, N, \{p \mapsto \mathsf{let}\ x = v\ \mathsf{in}\ e\}, [\,]\rangle \overset{\bullet}{\mapsto} \langle W, K, N, \{p \mapsto e\{v/x\}\}, [\,]\rangle}$$

(O-GET)
$$\frac{W(p) = (u, h, h') \qquad d = \mathit{domain}(u) \qquad \exists ck : K(d) = ck \wedge ck(k) = (n, f) \wedge f \in \{\bot, \mathtt{S}\}}{\langle W, K, N, \{p \mapsto k?\}, [\,]\rangle \overset{\bullet}{\mapsto} \langle W, K, N, \{p \mapsto n\}, [\,]\rangle}$$

(O-GETFAIL)
$$\frac{W(p) = (u, h, h') \qquad d = \mathit{domain}(u) \qquad \neg\exists ck : K(d) = ck \wedge ck(k) = (n, f) \wedge f \in \{\bot, \mathtt{S}\}}{\langle W, K, N, \{p \mapsto k?\}, [\,]\rangle \overset{\bullet}{\mapsto} \langle W, K, N, \{p \mapsto ()\}, [\,]\rangle}$$

(O-SET)
$$\frac{\begin{array}{c} W(p) = (u, h, h') \qquad d = \mathit{domain}(u) \\ \neg\exists ck : K(d) = ck \wedge ck(k) = (\_, f') \wedge f' \in \{\mathtt{H}, \top\} \qquad K' = \mathit{upd\_ck}(K, d, \{k \mapsto (n, f)\}) \end{array}}{\langle W, K, N, \{p \mapsto k!\langle n, f\rangle\}, [\,]\rangle \overset{\bullet}{\mapsto} \langle W, K', N, \{p \mapsto ()\}, [\,]\rangle}$$

(O-SETFAIL)
$$\frac{W(p) = (u, h, h') \qquad d = \mathit{domain}(u) \qquad \exists ck : K(d) = ck \wedge ck(k) = (\_, f') \wedge f' \in \{\mathtt{H}, \top\}}{\langle W, K, N, \{p \mapsto k!\langle n, f\rangle\}, [\,]\rangle \overset{\bullet}{\mapsto} \langle W, K, N, \{p \mapsto ()\}, [\,]\rangle}$$

(O-XHR)
$$\frac{h'' = h' \uplus \{n \mapsto \lambda x.e\} \qquad W' = W \uplus \{p \mapsto (u', h, h'')\} \qquad ck = \mathit{get\_ck}(K, u)}{\langle W \uplus \{p \mapsto (u', h, h')\}, K, N, \{p \mapsto \mathsf{xhr}(u, \lambda x.e)\}, [\,]\rangle \xrightarrow{\mathsf{xhr\_req}(ck,u)} \langle W', K, N \uplus \{n \mapsto (u, p)\}, \{p \mapsto ()\}, [\,]\rangle}$$

(O-LOGIN)
$$\frac{ck = \mathit{get\_ck}(K, u)}{\langle W, K, N, \{p \mapsto \mathsf{auth}(u, c)\}, [\,]\rangle \xrightarrow{\mathsf{login}(ck,u,c)} \langle W, K, N \uplus \{n \mapsto (u, p)\}, \{p \mapsto ()\}, [\,]\rangle}$$

(O-FLUSH)
$$\frac{}{\langle W, K, N, T, o\rangle \overset{o}{\mapsto} \langle W, K, N, T, [\,]\rangle}$$

(O-MIRROR)
$$\frac{P \overset{o}{\mapsto} Q}{P \xrightarrow{o} Q}$$

(O-COMPLETE)
$$\frac{\langle W, K, N, \{p \mapsto e\}, [\,]\rangle \not\mapsto}{\langle W, K, N, \{p \mapsto e\}, [\,]\rangle \xrightarrow{\bullet} \langle W, K, N, \{\}, [\,]\rangle}$$

Notation: we write $P \not\mapsto$ whenever there do not exist o and Q such that $P \overset{o}{\mapsto} Q$.

APPENDIX C
PROOFS

A. Preliminaries

An *extended state* is a triple $\xi = \langle Q, I, \tau \rangle$ which includes a state Q of the reactive system, an input stream I and a trust function τ. We define a small-step semantics for extended states, which corresponds to our previous notion of trace.

$$\text{(S-IN)}\ \frac{C \xrightarrow{i} P}{\langle C, i :: I, \tau \rangle \xrightarrow{i} \langle P, I, \tau \rangle} \qquad \text{(S-OUT)}\ \frac{P \xrightarrow{o} Q \qquad \tau \xrightarrow{o} \tau'}{\langle P, I, \tau \rangle \xrightarrow{o} \langle Q, I, \tau' \rangle}$$

An *attacked extended state* is a quadruple $\sigma = \langle Q, I, \tau, M \rangle$ which includes a state Q of the reactive system, an input stream I, a trust function τ and a set of events M intercepted/overheard by the attacker. We define a small-step semantics for attacked extended states, which corresponds to our previous notion of attacked trace.

$$\text{(AS-IN)}\ \frac{C \xrightarrow{i} P}{l \vdash \langle C, i :: I, \tau, M \rangle \xrightarrow{i} \langle P, I, \tau, M \rangle} \qquad \text{(AS-OUT)}\ \frac{P \xrightarrow{o} Q \qquad \tau \xrightarrow{o} \tau'}{l \vdash \langle P, I, \tau, M \rangle \xrightarrow{o} \langle Q, I, \tau', M \rangle}$$

$$\text{(AS-GETIN)}\ \frac{\tau, l \dagger i}{l \vdash \langle Q, i :: I, \tau, M \rangle \xrightarrow{\bullet} \langle Q, I, \tau, M \cup \{i\} \rangle}$$

$$\text{(AS-GETOUT)}\ \frac{P \xrightarrow{o} Q \qquad \tau, l \dagger o}{l \vdash \langle P, I, \tau, M \rangle \xrightarrow{\bullet} \langle Q, I, \tau, M \cup \{o\} \rangle} \qquad \text{(AS-HEARIN)}\ \frac{\tau, l \,?\, i}{l \vdash \langle Q, i :: I, \tau, M \rangle \xrightarrow{\bullet} \langle Q, i :: I, \tau, M \cup \{i\} \rangle}$$

$$\text{(AS-HEAROUT)}\ \frac{P \xrightarrow{o} Q \qquad \tau \xrightarrow{o} \tau' \qquad \tau, l \,?\, o}{l \vdash \langle P, I, \tau, M \rangle \xrightarrow{o} \langle Q, I, \tau', M \cup \{o\} \rangle} \qquad \text{(AS-SYNIN)}\ \frac{C \xrightarrow{i} P \qquad \tau, l, M \Vdash i}{l \vdash \langle C, I, \tau, M \rangle \xrightarrow{i} \langle P, I, \tau, M \rangle}$$

$$\text{(AS-SYNOUT)}\ \frac{\tau, l, M \Vdash o \qquad \tau \xrightarrow{o} \tau'}{l \vdash \langle Q, I, \tau, M \rangle \xrightarrow{o} \langle Q, I, \tau', M \rangle}$$

The formal correspondence between the big-step semantics and the small-step semantics is given by the two lemmas below.

Lemma 1 (Small-step Trace)**.** *If $\tau \vdash Q(I) \rightsquigarrow O$, then we have:*

$$\langle Q_0, I_0, \tau_0 \rangle \xrightarrow{\alpha_1} \langle Q_1, I_1, \tau_1 \rangle \xrightarrow{\alpha_2} \ldots \xrightarrow{\alpha_n} \langle Q_n, I_n, \tau_n \rangle,$$

where $Q_0 = Q$, $I_0 = I$, $\tau_0 = \tau$ and $O = [(o_i, \tau_{i-1}(o_i))_{1 \le i \le n} \mid o_i = \alpha_i]$.

Proof. By induction on the derivation of $\tau \vdash Q(I) \rightsquigarrow O$. □

Lemma 2 (Small-step Attacked Trace)**.** *If $\tau, l, M \vdash Q(I) \rightsquigarrow O$, then we have:*

$$l \vdash \langle Q_0, I_0, \tau_0, M_0 \rangle \xrightarrow{\alpha_1} \langle Q_1, I_1, \tau_1, M_1 \rangle \xrightarrow{\alpha_2} \ldots \xrightarrow{\alpha_n} \langle Q_n, I_n, \tau_n, M_n \rangle,$$

where $Q_0 = Q$, $I_0 = I$, $\tau_0 = \tau$, $M_0 = M$ and $O = [(o_i, \tau_{i-1}(o_i))_{1 \le i \le n} \mid o_i = \alpha_i]$.

Proof. By induction on the derivation of $\tau, l, M \vdash Q(I) \rightsquigarrow O$. □

B. Typing: safety results

We define a dynamic typing regime on the browser, which ensures a number of invariants which are needed to prove session integrity. These invariants must be preserved also when interacting with the opponent.

Notation 1 (Cookie Values)**.** *We let $ck\_vals(\alpha)$ stand for $ck\_vals(ck)$ whenever α is any input/output network event including the cookies ck.*

Notation 2 (Cookie Label). *Given a domain name d and a cookie flag f, we write $cookie\_label(d, f)$ for the security label defined as follows:*

$$cookie\_label(d, f) = \begin{cases} \mathsf{https}(d) & \text{if } f = \top \\ \mathsf{http}(d) & \text{if } f = \mathtt{H} \\ \bot & \text{otherwise.} \end{cases}$$

Definition 7 (Browser Typing). *Let $Q = \langle W, K, N, T, O\rangle$. We write $l \models Q$ if and only if:*

1) $\forall p \in dom(W) : W(p) = (u, h, h', q) \wedge url\_label(u) \sqsubseteq l \Rightarrow \forall n \in fn(h) \cup fn(h') : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$;
2) $\forall p \in dom(W) : W(p) = (u, h, h', q) \wedge url\_label(u) \not\sqsubseteq l \Rightarrow \forall n \in fn(h) : n \in \mathcal{N}_{\bot}$;
3) $\forall p \in dom(W) : W(p) = (u, h, h', q) \wedge url\_label(u) = \hat{l} \not\sqsubseteq l \Rightarrow \forall n \in fn(h') : \exists l' : (l' \sqsubseteq \hat{l} \vee l' \sqsubseteq l) \wedge n \in \mathcal{N}_{l'}$;
4) $\forall d \in dom(K) : K(d) = ck \wedge ck(k) = (m, f) \wedge cookie\_label(d, f) \sqsubseteq l \Rightarrow \forall n \in \{k, m\} : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$;
5) $\forall d \in dom(K) : K(d) = ck \wedge ck(k) = (n, f) \wedge cookie\_label(d, f) = l' \not\sqsubseteq l \Rightarrow n \in \mathcal{N}_{l'}$;
6) $\forall n \in dom(N) : N(n) = (u, p, q) \wedge W(p) = (u', h, h', q') \wedge q = \checkmark \Rightarrow url\_label(u) = url\_label(u')$;
7) $\forall p \in dom(T) : T(p) = (e, \hat{l}) \wedge \hat{l} \sqsubseteq l \Rightarrow \forall n \in fn(e) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$;
8) $\forall p \in dom(T) : T(p) = (e, \hat{l}) \wedge \hat{l} \not\sqsubseteq l \Rightarrow \forall n \in fn(e) : \exists l' : (l' \sqsubseteq \hat{l} \vee l' \sqsubseteq l) \wedge n \in \mathcal{N}_{l'}$;
9) $\forall o \in O : ev\_label(o) \sqsubseteq l \Rightarrow \forall n \in fn(o) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$;
10) $\forall o \in O : ev\_label(o) = l' \not\sqsubseteq l \Rightarrow ck\_vals(o) \subseteq \mathcal{N}_{l'}$;
11) $\forall o \in O : o \neq \mathsf{login}(ck, u, c)$.

Lemma 3 (Low HTTP Cookies). *Let K be a cookie jar such that:*

$$\forall d \in dom(K) : K(d) = ck \wedge ck(k) = (m, f) \wedge cookie\_label(d, f) \sqsubseteq l \Rightarrow \forall n \in \{k, m\} : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}.$$

If $url\_label(u) \sqsubseteq l$, then for any $n \in fn(get\_http\_ck(K, u))$ there exists $l' \sqsubseteq l$ such that $n \in \mathcal{N}_{l'}$.

Proof. Let $u = (\pi, d, v)$. We observe that, if $\{k \mapsto (n, f)\} \in get\_ck(K, u)$, then there exists ck such that $K(d) = ck$ and $ck(k) = (n, f)$, and in all cases $cookie\_label(d, f) = url\_label(u)$. Since we know that $url\_label(u) \sqsubseteq l$, we have $cookie\_label(d, f) \sqsubseteq l$ and we can conclude by our hypothesis on K. □

Lemma 4 (Output Secrecy). *Let $l \models P$ and $P \xrightarrow{o} Q$. If $ev\_label(o) \sqsubseteq l$ and $n \in fn(o)$, then $n \in \mathcal{N}_{l'}$ for some $l' \sqsubseteq l$.*

Proof. We observe that $P \xrightarrow{o} Q$ can be proved only by rule (O-MIRROR) or rule (O-COMPLETE). The only non-trivial case corresponds to (O-MIRROR), so we need to prove a similar result where $P \xrightarrow{o} Q$ has been replaced by $P \overset{o}{\mapsto} Q$. The proof is by induction on the derivation of $P \overset{o}{\mapsto} Q$:

Case (O-LETCTX): in this case we know that:

1) $P = \langle W, K, N, \{p \mapsto (\mathsf{let}\ x = e'\ \mathsf{in}\ e, \hat{l})\}, [\,]\rangle$;
2) $\langle W, K, N, \{p \mapsto (e', \hat{l})\}, [\,]\rangle \overset{o}{\mapsto} \langle W', K', N', \{p \mapsto (e'', \hat{l})\}, [\,]\rangle$;
3) $Q = \langle W', K', N', \{p \mapsto (\mathsf{let}\ x = e''\ \mathsf{in}\ e, \hat{l})\}, [\,]\rangle$.

Since $fn(e') \subseteq fn(\mathsf{let}\ x = e'\ \mathsf{in}\ e)$, it is easy to show that $l \models P$ implies $l \models \langle W, K, N, \{p \mapsto (e', \hat{l})\}, [\,]\rangle$. The conclusion thus follows by induction hypothesis;

Case (O-XHR): in this case we know that:

1) $P = \langle W \uplus \{p \mapsto (u', h, h', q)\}, K, N, \{p \mapsto (\mathsf{xhr}(u, \lambda x.e), \hat{l})\}, [\,]\rangle$;
2) $o = \mathsf{xhr\_req}(ck, u)$;
3) $q = \checkmark \wedge url\_label(u') = url\_label(u) \Rightarrow ck = get\_http\_ck(K, u)$;
4) $q = \times \vee url\_label(u') \neq url\_label(u) \Rightarrow ck = \{\}$;
5) $\hat{l} \neq \bot \Rightarrow \hat{l} = url\_label(u) = url\_label(u')$.

Let $ev\_label(o) = url\_label(u) \sqsubseteq l$, we want to show that for any name $n \in fn(o) = fn(ck) \cup fn(u)$ there exists a label $l' \sqsubseteq l$ such that $n \in \mathcal{N}_{l'}$. We first show the property for $fn(u)$, by distinguishing two cases:

- if $\hat{l} \sqsubseteq l$, we know that $\forall n \in fn(\mathsf{xhr}(u, \lambda x.e)) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$ by condition 7 of Definition 7. In particular, this implies that $\forall n \in fn(u) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$;
- if $\hat{l} \not\sqsubseteq l$, we know that $\forall n \in fn(\mathsf{xhr}(u, \lambda x.e)) : \exists l' : (l' \sqsubseteq \hat{l} \vee l' \sqsubseteq l) \wedge n \in \mathcal{N}_{l'}$ by condition 8 of Definition 7. In particular, this implies that $\forall n \in fn(u) : \exists l' : (l' \sqsubseteq \hat{l} \vee l' \sqsubseteq l) : n \in \mathcal{N}_{l'}$. Notice that it must be the case that

$\hat{l} \neq \bot$, hence we know that $\hat{l} = \mathit{url\_label}(u) = \mathit{ev\_label}(o)$ by assumption 5 above. Since $\mathit{ev\_label}(o) \sqsubseteq l$, we conclude that $\forall n \in \mathit{fn}(u) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$.

Now we focus on ck and we distinguish two cases. If $q = \times$ or $\mathit{url\_label}(u') \neq \mathit{url\_label}(u)$, then $ck = \{\}$ by assumption 4 above and we are done. Otherwise, let $q = \checkmark$ and $\mathit{url\_label}(u') = \mathit{url\_label}(u)$, then $ck = \mathit{get\_http\_ck}(K, u)$ by assumption 3 above. Since $l \models P$ holds true, we know that:

$$\forall d \in \mathit{dom}(K) : K(d) = ck \wedge ck(k) = (m, f) \wedge \mathit{cookie\_label}(d, f) \sqsubseteq l \Rightarrow \forall n \in \{k, m\} : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'},$$

by condition 4 of Definition 7. Since $\mathit{ev\_label}(o) = \mathit{url\_label}(u) \sqsubseteq l$ by hypothesis and $ck = \mathit{get\_http\_ck}(K, u)$, the desired conclusion follows by Lemma 3;

Case (O-LOGIN): in this case we know that:

1) $P = \langle W, K, N, \{p \mapsto (\mathsf{auth}(u, c), \hat{l})\}, [\,]\rangle$;
2) $W(p) = (u', h, h', \checkmark)$;
3) $o = \mathsf{login}(ck, u, c)$ with $ck = \mathit{get\_http\_ck}(K, u)$;
4) $\rho(c) = \mathit{url\_label}(u)$;
5) $\hat{l} = \mathit{url\_label}(u) = \mathit{url\_label}(u')$.

Let $\mathit{ev\_label}(o) = \mathit{url\_label}(u) \sqsubseteq l$, we want to show that for any name $n \in \mathit{fn}(o) = \mathit{fn}(ck) \cup \mathit{fn}(u) \cup \{c\}$ there exists a label $l' \sqsubseteq l$ such that $n \in \mathcal{N}_{l'}$. We first show the property for $\mathit{fn}(u) \cup \{c\}$, by distinguishing two cases:

- if $\hat{l} \sqsubseteq l$, we know that $\forall n \in \mathit{fn}(\mathsf{auth}(u, c)) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$ by condition 7 of Definition 7. In particular, this implies that $\forall n \in \mathit{fn}(u) \cup \{c\} : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$;
- if $\hat{l} \not\sqsubseteq l$, we know that $\forall n \in \mathit{fn}(\mathsf{auth}(u, c)) : \exists l' : (l' \sqsubseteq \hat{l} \vee l' \sqsubseteq l) : n \in \mathcal{N}_{l'}$ by condition 8 of Definition 7. In particular, this implies that $\forall n \in \mathit{fn}(u) \cup \{c\} : \exists l' : (l' \sqsubseteq \hat{l} \vee l' \sqsubseteq l) \wedge n \in \mathcal{N}_{l'}$. Now we know that $\hat{l} = \mathit{url\_label}(u) = \mathit{ev\_label}(o)$ by assumption 5 above. Since $\mathit{ev\_label}(o) \sqsubseteq l$, we conclude that $\forall n \in \mathit{fn}(u) \cup \{c\} : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$.

Now we focus on ck: since $l \models P$ holds true, we know that:

$$\forall d \in \mathit{dom}(K) : K(d) = ck \wedge ck(k) = (m, f) \wedge \mathit{cookie\_label}(d, f) \sqsubseteq l \Rightarrow \forall n \in \{k, m\} : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'},$$

by condition 4 of Definition 7. Since $\mathit{ev\_label}(o) = \mathit{url\_label}(u) \sqsubseteq l$ by hypothesis and $ck = \mathit{get\_http\_ck}(K, u)$, the desired conclusion follows by Lemma 3;

Case (O-FLUSH): in this case we know that:

1) $P = \langle W, K, N, T, o\rangle$;
2) $Q = \langle W, K, N, T, [\,]\rangle$.

Since $l \models P$ holds true, we know that $\mathit{ev\_label}(o) \sqsubseteq l \Rightarrow \forall n \in \mathit{fn}(o) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$ by condition 9 of Definition 7. Hence, the conclusion is immediate.

□

Lemma 5 (High HTTP Cookies)**.** *Let K be a cookie jar such that:*

$$\forall d \in \mathit{dom}(K) : K(d) = ck \wedge ck(k) = (n, f) \wedge \mathit{cookie\_label}(d, f) = l' \not\sqsubseteq l \Rightarrow n \in \mathcal{N}_{l'}.$$

If $\mathit{url\_label}(u) = l' \not\sqsubseteq l$, then for any $n \in \mathit{ck\_vals}(\mathit{get\_http\_ck}(K, u))$ we have $n \in \mathcal{N}_{l'}$.

Proof. Let $u = (\pi, d, v)$. We observe that, if $\{k \mapsto (n, f)\} \in \mathit{get\_http\_ck}(K, u)$, then there exists ck such that $K(d) = ck$ and $ck(k) = (n, f)$, and in all cases $\mathit{cookie\_label}(d, f) = \mathit{url\_label}(u)$. Since we know that $\mathit{url\_label}(u) = l' \not\sqsubseteq l$, we have $\mathit{cookie\_label}(d, f) = l' \not\sqsubseteq l$ and we can conclude by our hypothesis on K. □

Lemma 6 (Preventing Fixation)**.** *If $l \models P$ and $P \xrightarrow{o} Q$ with $\mathit{ev\_label}(o) = l' \not\sqsubseteq l$, then $\mathit{ck\_vals}(o) \subseteq \mathcal{N}_{l'}$.*

Proof. We observe that $P \xrightarrow{o} Q$ can be proved only by rule (O-MIRROR) or rule (O-COMPLETE). The only non-trivial case corresponds to (O-MIRROR), so we need to prove a similar result where $P \xrightarrow{o} Q$ has been replaced by $P \overset{o}{\mapsto} Q$. The proof is by induction on the derivation of $P \overset{o}{\mapsto} Q$:

Case (O-LETCTX): in this case we know that:

1) $P = \langle W, K, N, \{p \mapsto (\mathsf{let}\ x = e'\ \mathsf{in}\ e, \hat{l})\}, [\,]\rangle$;
2) $\langle W, K, N, \{p \mapsto (e', \hat{l})\}, [\,]\rangle \overset{o}{\mapsto} \langle W', K', N', \{p \mapsto (e'', \hat{l})\}, [\,]\rangle$;

3) $Q = \langle W', K', N', \{p \mapsto (\mathsf{let}\ x = e''\ \mathsf{in}\ e, \hat{l})\}, [\,]\rangle$.

Since $fn(e') \subseteq fn(\mathsf{let}\ x = e'\ \mathsf{in}\ e)$, it is easy to show that $l \models P$ implies $l \models \langle W, K, N, \{p \mapsto (e', \hat{l})\}, [\,]\rangle$. The conclusion thus follows by induction hypothesis;

Case (O-XHR): in this case we know that:

1) $P = \langle W \uplus \{p \mapsto (u', h, h', q)\}, K, N, \{p \mapsto (\mathsf{xhr}(u, \lambda x.e), \hat{l})\}, [\,]\rangle$;
2) $o = \mathsf{xhr\_req}(ck, u)$;
3) $q = \checkmark \wedge \mathit{url\_label}(u') = \mathit{url\_label}(u) \Rightarrow ck = \mathit{get\_http\_ck}(K, u)$;
4) $q = \times \vee \mathit{url\_label}(u') \neq \mathit{url\_label}(u) \Rightarrow ck = \{\}$;
5) $\hat{l} \neq \bot \Rightarrow \hat{l} = \mathit{url\_label}(u) = \mathit{url\_label}(u')$.

Let $\mathit{ev\_label}(o) = \mathit{url\_label}(u) = l' \not\sqsubseteq l$, we want to show that $\mathit{ck\_vals}(ck) \subseteq \mathcal{N}_{l'}$. We distinguish two cases. If $\mathit{url\_label}(u') \neq \mathit{url\_label}(u)$, then $ck = \{\}$ by assumption 4 above and we are done. Otherwise, let $\mathit{url\_label}(u') = \mathit{url\_label}(u)$, then $ck = \mathit{get\_http\_ck}(K, u)$ by assumption 3 above. Since $l \models P$ holds true and $\mathit{url\_label}(u) = l' \not\sqsubseteq l$, the conclusion follows by Lemma 5;

Case (O-LOGIN): in this case we know that:

1) $P = \langle W, K, N, \{p \mapsto (\mathsf{auth}(u, c), \hat{l})\}, [\,]\rangle$;
2) $W(p) = (u', h, h', \checkmark)$;
3) $o = \mathsf{login}(ck, u, c)$ with $ck = \mathit{get\_http\_ck}(K, u)$;
4) $\rho(c) = \mathit{url\_label}(u)$;
5) $\hat{l} = \mathit{url\_label}(u) = \mathit{url\_label}(u')$.

Let $\mathit{ev\_label}(o) = \mathit{url\_label}(u) = l' \not\sqsubseteq l$, we want to show that $\mathit{ck\_vals}(ck) \subseteq \mathcal{N}_{l'}$. Since $l \models P$ holds true and $\mathit{url\_label}(u) = l' \not\sqsubseteq l$, the conclusion follows by Lemma 5;

Case (O-FLUSH): in this case we know that:

1) $P = \langle W, K, N, T, o\rangle$;
2) $Q = \langle W, K, N, T, [\,]\rangle$.

Since $l \models P$ holds true, we know that $\mathit{ev\_label}(o) = l' \not\sqsubseteq l \Rightarrow \mathit{ck\_vals}(o) \subseteq \mathcal{N}_{l'}$ by condition 9 of Definition 7. Hence, the conclusion is immediate.

□

Lemma 7 (Evil Output)**.** *Let $l \models P$ and let τ be a trust function such that:*

$$\forall o \in \mathcal{O} : \tau(o) = \mathsf{evil} \Rightarrow \exists n \in \mathit{ck\_vals}(o) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}.$$

If $P \xrightarrow{o} Q$ and $\tau(o) = \mathsf{evil}$, then $\mathit{ev\_label}(o) \sqsubseteq l$.

Proof. Since $\tau(o) = \mathsf{evil}$, by our hypothesis on τ we know that there exists $n \in \mathit{ck\_vals}(o)$ such that $n \in \mathcal{N}_{l'}$ for some $l' \sqsubseteq l$. Let us assume by contradiction that $\mathit{ev\_label}(o) = \hat{l} \not\sqsubseteq l$, then by Lemma 6 we know that $\mathit{ck\_vals}(o) \subseteq \mathcal{N}_{\hat{l}}$. This implies that $n \notin \mathit{ck\_vals}(o)$, hence we get a contradiction and we conclude. □

Lemma 8 (Preventing Login CSRF)**.** *If $l \models P$ and $P \xrightarrow{\mathsf{login}(ck,u,c)} Q$, then $\rho(c) \neq \mathsf{evil}$.*

Proof. We observe that $P \xrightarrow{o} Q$ can be proved only by rule (O-MIRROR) or rule (O-COMPLETE). The only non-trivial case corresponds to (O-MIRROR), so we need to prove a similar result where $P \xrightarrow{o} Q$ has been replaced by $P \overset{o}{\mapsto} Q$. The proof is by induction on the derivation of $P \overset{o}{\mapsto} Q$:

Case (O-LETCTX): in this case we know that:

1) $P = \langle W, K, N, \{p \mapsto (\mathsf{let}\ x = e'\ \mathsf{in}\ e, \hat{l})\}, [\,]\rangle$;
2) $\langle W, K, N, \{p \mapsto (e', \hat{l})\}, [\,]\rangle \xmapsto{\mathsf{login}(ck,u,c)} \langle W', K', N', \{p \mapsto (e'', \hat{l})\}, [\,]\rangle$;
3) $Q = \langle W', K', N', \{p \mapsto (\mathsf{let}\ x = e''\ \mathsf{in}\ e, \hat{l})\}, [\,]\rangle$.

Since $fn(e') \subseteq fn(\mathsf{let}\ x = e'\ \mathsf{in}\ e)$, it is easy to show that $l \models P$ implies $l \models \langle W, K, N, \{p \mapsto (e', \hat{l})\}, [\,]\rangle$. The conclusion thus follows by induction hypothesis;

Case (O-LOGIN): in this case we know that:

1) $P = \langle W, K, N, \{p \mapsto (\mathsf{auth}(u, c), \hat{l})\}, [\,]\rangle$;
2) $W(p) = (u', h, h', \checkmark)$;

3) $o = \mathsf{login}(ck, u, c)$ with $ck = \mathit{get\_http\_ck}(K, u)$;
4) $\rho(c) = \mathit{url\_label}(u)$;
5) $l = \mathit{url\_label}(u) = \mathit{url\_label}(u')$.

The conclusion is immediate, since $\rho(c) = \mathit{url\_label}(u) \neq \mathsf{evil}$;

Case (O-FLUSH): in this case we know that:

1) $P = \langle W, K, N, T, \mathsf{login}(ck, u, c)\rangle$;
2) $Q = \langle W, K, N, T, [\,]\rangle$.

But this scenario is excluded by the assumption $l \models P$ (see condition 11 of Definition 7). Hence, the case is vacuous and we conclude.

□

C. Typing: subject reduction

We now show that typing is preserved at runtime. We first prove a subject reduction lemma for the browser, then we use it to prove a more general subject reduction theorem over extended attacked states.

Lemma 9 (Cookie Flags)**.** *The following statements hold true:*

1) if $\{k \mapsto (n, f)\} \in (ck \nearrow \mathsf{http})$*, then* $f = \mathtt{H}$*;*
2) if $\{k \mapsto (n, f)\} \in (ck \nearrow \mathsf{https})$*, then* $f = \top$.

Proof. By induction on the structure of ck. □

Lemma 10 (Secure Low/Low Update)**.** *Let K be a cookie jar such that:*

$$\forall d \in \mathit{dom}(K) : K(d) = ck \wedge ck(k) = (m, f) \wedge \mathit{cookie\_label}(d, f) \sqsubseteq l \Rightarrow \forall n \in \{k, m\} : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}.$$

If $\mathit{url\_label}(u) \sqsubseteq l$ *and* $K' = \mathit{sec\_upd\_ck}(K, u, ck')$ *for some* ck' *such that* $\forall n \in \mathit{fn}(ck') : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$*, then:*

$$\forall d \in \mathit{dom}(K') : K'(d) = ck \wedge ck(k) = (m, f) \wedge \mathit{cookie\_label}(d, f) \sqsubseteq l \Rightarrow \forall n \in \{k, m\} : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}.$$

Proof. Let $u = (\pi, d, v)$. By using Lemma 9 we observe that, whenever $\{k \mapsto (n, f)\} \in (ck' \nearrow \pi)$, we have $\mathit{cookie\_label}(d, f) = \mathit{url\_label}(u)$, hence the conclusion follows by our assumptions on K and ck. □

Lemma 11 (Secure Low/High Update)**.** *Let K be a cookie jar such that:*

$$\forall d \in \mathit{dom}(K) : K(d) = ck \wedge ck(k) = (n, f) \wedge \mathit{cookie\_label}(d, f) = l' \not\sqsubseteq l \Rightarrow n \in \mathcal{N}_{l'}.$$

If $\mathit{url\_label}(u) \sqsubseteq l$ *and* $K' = \mathit{sec\_upd\_ck}(K, u, ck')$ *for some* ck'*, then:*

$$\forall d \in \mathit{dom}(K') : K'(d) = ck \wedge ck(k) = (n, f) \wedge \mathit{cookie\_label}(d, f) = l' \not\sqsubseteq l \Rightarrow n \in \mathcal{N}_{l'}.$$

Proof. Let $u = (\pi, d, v)$. By using Lemma 9 we observe that, whenever $\{k \mapsto (n, f)\} \in (ck' \nearrow \pi)$, we have $\mathit{cookie\_label}(d, f) = \mathit{url\_label}(u)$, hence the conclusion follows by our assumption on K. □

Lemma 12 (Secure High/Low Update)**.** *Let K be a cookie jar such that:*

$$\forall d \in \mathit{dom}(K) : K(d) = ck \wedge ck(k) = (m, f) \wedge \mathit{cookie\_label}(d, f) \sqsubseteq l \Rightarrow \forall n \in \{k, m\} : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}.$$

If $\mathit{url\_label}(u) = l' \not\sqsubseteq l$ *and* $K' = \mathit{sec\_upd\_ck}(K, u, ck')$ *for some* ck'*, then:*

$$\forall d \in \mathit{dom}(K') : K'(d) = ck \wedge ck(k) = (m, f) \wedge \mathit{cookie\_label}(d, f) \sqsubseteq l \Rightarrow \forall n \in \{k, m\} : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}.$$

Proof. Let $u = (\pi, d, v)$. By using Lemma 9 we observe that, whenever $\{k \mapsto (n, f)\} \in (ck' \nearrow \pi)$, we have $\mathit{cookie\_label}(d, f) = \mathit{url\_label}(u)$, hence the conclusion follows by our assumption on K. □

Lemma 13 (Secure High/High Update)**.** *Let K be a cookie jar such that:*

$$\forall d \in \mathit{dom}(K) : K(d) = ck \wedge ck(k) = (n, f) \wedge \mathit{cookie\_label}(d, f) = l' \not\sqsubseteq l \Rightarrow n \in \mathcal{N}_{l'}.$$

If $\mathit{url\_label}(u) = l' \not\sqsubseteq l$ *and* $K' = \mathit{sec\_upd\_ck}(K, u, ck')$ *for some* ck' *such that* $\mathit{ck\_vals}(ck') \subseteq \mathcal{N}_{l'}$*, then:*

$$\forall d \in \mathit{dom}(K') : K'(d) = ck \wedge ck(k) = (n, f) \wedge \mathit{cookie\_label}(d, f) = l' \not\sqsubseteq l \Rightarrow n \in \mathcal{N}_{l'}.$$

Proof. Let $u = (\pi, d, v)$. By using Lemma 9 we observe that, whenever $\{k \mapsto (n, f)\} \in (ck' \nearrow \pi)$, we have $cookie\_label(d, f) = url\_label(u)$, hence the conclusion follows by our assumptions on K and ck. □

The next definition is a generalization of the concept of well-formed input event. Intuitively, since in the attacked run the browser has to deal with tainted inputs created by the intruder, we want to constrain the structure of these inputs.

Definition 8 (Consistent Input)**.** *An input i is* consistent *with respect to a security label l, written $l \vdash_\diamond i$, if either $\vdash_\diamond i$, or $\forall n \in fn(i) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$ and $ev\_label(i) \sqsubseteq l$.*

Lemma 14 (Well-formed Input Secrecy)**.** *If $\vdash_\diamond i$ and $n \in fn(i)$, then $n \in \mathcal{N}_l$ for some $l \sqsubseteq ev\_label(i)$.*

Proof. By a case analysis on the rule applied to prove $\vdash_\diamond i$. □

The next subject reduction lemma guarantees that the browser preserves typing upon reduction. Notice that the lemma does not hold for arbitrary input events, but all the events which are present in a well-formed trace and/or can be generated by the opponent do allow to preserve typing.

Lemma 15 (Browser Subject Reduction)**.** *The following statements hold true:*

1) if $l \models C$ and $C \xrightarrow{i} P$ and $l \vdash_\diamond i$, then $l \models P$;
2) if $l \models P$ and $P \xrightarrow{o} Q$, then $l \models Q$.

Proof. We start by proving the first point. If $C \xrightarrow{i} P$ was derived by (I-COMPLETE), the conclusion is immediate, since the internal state of the browser does not change, but for the output buffer, which will only contain the dummy output ("•"). Let then $C \xrightarrow{i} P$ be proved by (I-MIRROR), we need to prove a variant of the statement where $C \xrightarrow{i} P$ has been replaced by $C \overset{i}{\mapsto} P$. The proof proceeds by a case analysis on the rule applied to prove $C \overset{i}{\mapsto} P$:

Case (I-LOAD): in this case we know that:
1) $C = \langle W, K, N, \{\}, [\,]\rangle$;
2) $i = \mathsf{load}(u)$ with $\vdash_\diamond \mathsf{load}(u)$;
3) $P = \langle W, K, N \uplus \{n \mapsto (u, (), \checkmark)\}, \{\}, \mathsf{doc\_req}(ck, u)\rangle$ with $ck = get\_http\_ck(K, u)$.

We only have to preserve the invariants on the output buffer. Specifically, we have to prove that:

$$url\_label(u) \sqsubseteq l \Rightarrow \forall n \in fn(u) \cup fn(ck) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'} \tag{1}$$
$$url\_label(u) = l' \not\sqsubseteq l \Rightarrow ck\_vals(ck) \subseteq \mathcal{N}_{l'} \tag{2}$$

We first show (1). Let $u = (\pi, d, v)$, since $\vdash_\diamond \mathsf{load}(u)$ holds true we know that $d \in \mathcal{N}_\perp$ and there exists $\hat{l} \sqsubseteq url\_label(u) : v \in \mathcal{N}_{\hat{l}}$. Thus, assuming $url\_label(u) \sqsubseteq l$, we have that $\hat{l} \sqsubseteq l$ and we proved that $\forall n \in fn(u) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$. To conclude, we notice that $\forall n \in fn(ck) = get\_http\_ck(K, u) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$ by Lemma 3. We now focus on (2). Since we assume that $url\_label(u) = l' \not\sqsubseteq l$, we know that $ck\_vals(ck) = ck\_vals(get\_http\_ck(K, u)) \subseteq \mathcal{N}_{l'}$ by Lemma 5 and we are done;

Case (I-TEXT): in this case we know that:
1) $C = \langle W, K, N, \{\}, [\,]\rangle$;
2) $i = \mathsf{text}(p, k, n)$ with $\vdash_\diamond \mathsf{text}(p, k, n)$;
3) $P = \langle W, K, N, \{p \mapsto (e\{n/x\}, \rho(n))\}, [\,]\rangle$;
4) $W(p) = (u, h, h', q)$;
5) $h(k) = \lambda x.e$.

We only have to preserve the invariants on the task list. Specifically, we have to prove that:

$$\rho(n) \sqsubseteq l \Rightarrow \forall m \in fn(e\{n/x\}) : \exists l' \sqsubseteq l : m \in \mathcal{N}_{l'} \tag{3}$$
$$\rho(n) \not\sqsubseteq l \Rightarrow \forall m \in fn(e\{n/x\}) : \exists l' : (l' \sqsubseteq \rho(n) \vee l' \sqsubseteq l) \wedge m \in \mathcal{N}_{l'} \tag{4}$$

We first observe that $\vdash_\diamond \mathsf{text}(p, k, n)$ implies $n \in \mathcal{N}_{\rho(n)}$, then we show the previous implications separately. Let $\rho(n) \sqsubseteq l$, we distinguish two cases to prove (3):

- if $url\_label(u) \sqsubseteq l$, by condition 1 of Definition 7 we know that $\forall n \in fn(h) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$. In particular, this implies that $\forall n \in fn(e) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$. Since we know that $n \in \mathcal{N}_{\rho(n)}$ and $\rho(n) \sqsubseteq l$, we conclude that $\forall m \in fn(e\{n/x\}) : \exists l' \sqsubseteq l : m \in \mathcal{N}_{l'}$;
- if $url\_label(u) \not\sqsubseteq l$, by condition 2 of Definition 7 we know that $\forall n \in fn(h) : n \in \mathcal{N}_{\bot}$. In particular, this implies that $\forall n \in fn(e) : n \in \mathcal{N}_{\bot}$. Since we know that $n \in \mathcal{N}_{\rho(n)}$ and $\rho(n) \sqsubseteq l$, we conclude that $\forall m \in fn(e\{n/x\}) : \exists l' \sqsubseteq l : m \in \mathcal{N}_{l'}$.

Let now $\rho(n) \not\sqsubseteq l$, we distinguish again two cases to prove (4):

- if $url\_label(u) \sqsubseteq l$, by condition 1 of Definition 7 we know that $\forall n \in fn(h) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$. In particular, this implies that $\forall n \in fn(e) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$. Since we know that $n \in \mathcal{N}_{\rho(n)}$, we conclude that $\forall m \in fn(e\{n/x\}) : \exists l' : (l' \sqsubseteq \rho(n) \vee l' \sqsubseteq l) \wedge m \in \mathcal{N}_{l'}$;
- if $url\_label(u) \not\sqsubseteq l$, by condition 2 of Definition 7 we know that $\forall n \in fn(h) : n \in \mathcal{N}_{\bot}$. In particular, this implies that $\forall n \in fn(e) : n \in \mathcal{N}_{\bot}$. Since we know that $n \in \mathcal{N}_{\rho(n)}$, we conclude that $\forall m \in fn(e\{n/x\}) : \exists l' \sqsubseteq \rho(n) : m \in \mathcal{N}_{l'}$;

Case (I-DocResp): in this case we know that:

1) $C = \langle W, K, N \uplus \{n \mapsto (u, (), q)\}, \{\}, [\,]\rangle$;
2) $i = \mathsf{doc\_resp}(n, ck, u, \mathsf{blank}, h, e)$ with $l \vdash_{\diamond} i$;
3) $P = \langle W \uplus \{p \mapsto (u, h, \{\}, q)\}, K', N, \{p \mapsto (e, \bot)\}, [\,]\rangle$;
4) $q = \checkmark \Rightarrow K' = \mathit{sec\_upd\_ck}(K, u, ck)$;
5) $q = \times \Rightarrow K' = K$.

We have to prove the invariant for the new window, the updated cookie jar and the new task. We distinguish two cases, based on why $l \vdash_{\diamond} i$ holds true:

- in the first case we know that $\vdash_{\diamond} i$ holds true. This implies that $fn(h) \subseteq \mathcal{N}_{\bot}$ and thus both conditions 1 and 2 hold true. Condition 3 is trivial, thus we proved the invariant for the new window.
  We now move to the cookie jar: if $K' = K$, we do not need to prove anything, so we assume $K' = \mathit{sec\_upd\_ck}(K, u, ck)$. If $url\_label(u) \sqsubseteq l$, we have that $\forall n \in fn(ck) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$ by Lemma 14, hence we can preserve conditions 4 and 5 by Lemmas 10 and 11 respectively. Otherwise, let $url\_label(u) = l' \not\sqsubseteq l$: since $\vdash_{\diamond} i$ holds true, we have that $ck\_vals(ck) \subseteq \mathcal{N}_{l'}$, hence we can preserve conditions 4 and 5 by Lemmas 12 and 13 respectively.
  Finally, we consider the task list. Since $\vdash_{\diamond} i$ holds true, we know that $\forall n \in fn(e) : n \in \mathcal{N}_{\bot}$, hence both conditions 7 and 8 are satisfied and we conclude;
- in the second case we know that $\forall n \in fn(i) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$ and $ev\_label(i) = url\_label(u) \sqsubseteq l$. Condition 1 immediately follows, while conditions 2 and 3 are trivially satisfied.
  We now move to the cookie jar: if $K' = K$, we do not need to prove anything, so we assume $K' = \mathit{sec\_upd\_ck}(K, u, ck)$. Since we know that $\forall n \in fn(ck) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$ and $url\_label(u) \sqsubseteq l$, we can preserve conditions 4 and 5 by Lemmas 10 and 11 respectively.
  Finally, we consider the task list. Since we know that $\forall n \in fn(e) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$, condition 7 is satisfied. Condition 8 is trivial, since $\bot$ is the lowest security label;

Case (I-DocRedir): in this case we know that:

1) $C = \langle W, K, N \uplus \{n \mapsto (u, (), q)\}, \{\}, [\,]\rangle$;
2) $i = \mathsf{doc\_resp}(n, ck, u, u', h, e)$ with $l \vdash_{\diamond} i$;
3) $P = \langle W, K', N \uplus \{n \mapsto (u', (), q')\}, \{\}, \mathsf{doc\_req}(ck'', u')\rangle$;
4) $q = \checkmark \Rightarrow K' = \mathit{sec\_upd\_ck}(K, u, ck)$;
5) $q = \times \Rightarrow K' = K$;
6) $q = \checkmark \wedge url\_label(u) = url\_label(u') \Rightarrow ck'' = \mathit{get\_http\_ck}(K', u') \wedge q' = \checkmark$;
7) $q = \times \vee url\_label(u) \neq url\_label(u') \Rightarrow ck'' = \{\} \wedge q' = \times$.

We have to preserve the invariant for the the updated cookie jar and the new output buffer. We distinguish two cases, based on why $l \vdash_{\diamond} i$ holds true:

- in the first case we know that $\vdash_{\diamond} i$ holds true. If $K' = K$, we do not need to prove anything, so we assume $K' = \mathit{sec\_upd\_ck}(K, u, ck)$. If $url\_label(u) \sqsubseteq l$, we have that $\forall n \in fn(ck) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$ by Lemma 14, hence we can preserve conditions 4 and 5 by Lemmas 10 and 11 respectively. Otherwise,

let $url\_label(u) = l' \not\sqsubseteq l$: since $\vdash_\diamond i$ holds true, we have that $ck\_vals(ck) \subseteq \mathcal{N}_{l'}$, hence we can preserve conditions 4 and 5 by Lemmas 12 and 13 respectively.
We then move to the output buffer. Let $ev\_label(\mathsf{doc\_req}(ck'', u')) = url\_label(u') \sqsubseteq l$, we want to show that $\forall n \in fn(\mathsf{doc\_req}(ck'', u')) = fn(ck'') \cup fn(u') : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$ to prove condition 9. As to $fn(u')$, we know that $\forall n \in fn(u') : \exists l' \sqsubseteq url\_label(u') : n \in \mathcal{N}_{l'}$ by the assumption $\vdash_\diamond u'$ (in the premises of $\vdash_\diamond i$). We then focus on $fn(ck'')$: if $q = \times$ or $url\_label(u) \neq url\_label(u')$, we have $ck'' = \{\}$, hence $fn(ck'') = \emptyset$ and we are done; if $q = \checkmark$ and $url\_label(u) = url\_label(u')$, we have $ck'' = \textit{get\_http\_ck}(K', u')$ and we have that $\forall n \in fn(ck') : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$ by Lemma 3. Thus, we proved condition 9.
Let now $ev\_label(\mathsf{doc\_req}(ck'', u')) = url\_label(u') = l' \not\sqsubseteq l$, we want to show that $ck\_vals(ck'') \subseteq \mathcal{N}_{l'}$ to prove condition 10. If $q = \times$ or $url\_label(u) \neq url\_label(u')$, we have $ck'' = \{\}$, hence $ck\_vals(ck'') = \emptyset$ and we are done. If $q = \checkmark$ and $url\_label(u) = url\_label(u')$, we have $ck'' = \textit{get\_http\_ck}(K', u')$ and we have that $ck\_vals(ck') \subseteq \mathcal{N}_{l'}$ by Lemma 5. Thus, we proved condition 10. Condition 11 is immediate.
- in the second case we know that $\forall n \in fn(i) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$ and $ev\_label(i) = url\_label(u) \sqsubseteq l$. We start with the cookie jar: if $K' = K$, we do not need to prove anything, so we assume $K' = \textit{sec\_upd\_ck}(K, u, ck)$. Since we know that $\forall n \in fn(ck) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$ and $url\_label(u) \sqsubseteq l$, we can preserve conditions 4 and 5 by Lemmas 10 and 11 respectively.
We then move to the output buffer. Let $ev\_label(\mathsf{doc\_req}(ck'', u')) = url\_label(u') \sqsubseteq l$, we want to show that $\forall n \in fn(\mathsf{doc\_req}(ck'', u')) = fn(ck'') \cup fn(u') : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$ to prove condition 9. As to $fn(u') \subseteq fn(i)$, we already know that $\forall n \in fn(u') : \exists l' \sqsubseteq l : n \subseteq \mathcal{N}_{l'}$. We then focus on $fn(ck'')$: if $q = \times$ or $url\_label(u) \neq url\_label(u')$, we have $ck'' = \{\}$, hence $fn(ck'') = \emptyset$ and we are done; if $q = \checkmark$ and $url\_label(u) = url\_label(u')$, we have $ck'' = \textit{get\_http\_ck}(K', u')$ and we have that $\forall n \in fn(ck') : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$ by Lemma 5. Thus, we proved condition 9.
Let now $ev\_label(\mathsf{doc\_req}(ck'', u')) = url\_label(u') = l' \not\sqsubseteq l$, we want to show that $ck\_vals(ck'') \subseteq \mathcal{N}_{l'}$ to prove condition 10. If $q = \times$ or $url\_label(u) \neq url\_label(u')$, we have $ck'' = \{\}$, hence $ck\_vals(ck'') = \emptyset$ and we are done. If $q = \checkmark$ and $url\_label(u) = url\_label(u')$, we get a contradiction, since we know that $url\_label(u) \sqsubseteq l$ and we trivially conclude. Thus, we proved condition 10. Condition 11 is immediate;

Case (I-XHRRESP): in this case we know that:

1) $C = \langle W \uplus \{p \mapsto (u', h, h' \uplus \{n \mapsto \lambda x.e\}, q)\}, K, N \uplus \{n \mapsto (u, p, q)\}, \{\}, [\,]\rangle$;
2) $i = \mathsf{xhr\_resp}(n, ck, u, \mathsf{blank}, v)$ with $l \vdash_\diamond i$;
3) $P = \langle W \uplus \{p \mapsto (u', h, h', q)\}, K', N, \{p \mapsto (e\{v/x\}, \hat{l})\}, [\,]\rangle$;
4) $q = \checkmark \Rightarrow K' = \textit{sec\_upd\_ck}(K, u, ck)$;
5) $q = \times \Rightarrow K' = K$;
6) $\hat{l} = url\_label(u')$.

The invariants on windows and network connections are preserved, since we are just removing an handler/a network connection, while the properties of the cookie jar are shown exactly as in case (I-DOCRESP). The only interesting properties to show are related to the task list:

$$\hat{l} \sqsubseteq l \Rightarrow \forall n \in fn(e\{v/x\}) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'} \tag{5}$$

$$\hat{l} \not\sqsubseteq l \Rightarrow \forall n \in fn(e\{v/x\}) : \exists l' : (l' \sqsubseteq \hat{l} \vee l' \sqsubseteq l) \wedge n \in \mathcal{N}_{l'} \tag{6}$$

To carry out our reasoning, we distinguish two cases, based on why $l \vdash_\diamond i$ holds true:

- in the first case we know that $\vdash_\diamond i$ holds true, hence we have $fn(v) \subseteq \mathcal{N}_\perp$. We first prove (5): let $\hat{l} = url\_label(u') \sqsubseteq l$, then by condition 1 of Definition 7 we know that $\forall n \in fn(h') : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$. In particular, this implies that $\forall n \in fn(e) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$. Thus, we conclude that $\forall n \in fn(e\{v/x\}) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$. We now prove (6): let $\hat{l} = url\_label(u') \not\sqsubseteq l$, then by condition 3 of Definition 7 we know that $\forall n \in fn(h') : \exists l' : (l' \sqsubseteq \hat{l} \vee l' \sqsubseteq l) \wedge n \in \mathcal{N}_{l'}$. Hence, this condition must be true in particular for e. Since we know that $fn(v) \subseteq \mathcal{N}_\perp$, we conclude that $\forall n \in fn(e\{v/x\}) : \exists l' : (l' \sqsubseteq \hat{l} \vee l' \sqsubseteq l) \wedge n \in \mathcal{N}_{l'}$;
- in the second case we know that $\forall n \in fn(i) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$ and $ev\_label(i) = url\_label(u) \sqsubseteq l$. We first prove (5): let $\hat{l} = url\_label(u') \sqsubseteq l$, then by condition 1 of Definition 7 we know that $\forall n \in fn(h') : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$. In particular, this implies that $\forall n \in fn(e) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$. Since $fn(v) \subseteq fn(i)$, we can conclude that $\forall n \in fn(e\{v/x\}) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$. We now prove (6): let $\hat{l} = url\_label(u') \not\sqsubseteq l$,

then by condition 3 of Definition 7 we know that $\forall n \in \mathit{fn}(h') : \exists l' : (l' \sqsubseteq \hat{l} \vee l' \sqsubseteq l) \wedge n \in \mathcal{N}_{l'}$. Hence, this condition must be true in particular for e. Since we know that $\mathit{fn}(v) \subseteq \mathit{fn}(i)$, we conclude that $\forall n \in \mathit{fn}(e\{v/x\}) : \exists l' : (l' \sqsubseteq \hat{l} \vee l' \sqsubseteq l) \wedge n \in \mathcal{N}_{l'}$.

Case (I-XHRREDIR): in this case we know that:

1) $C = \langle W, K, N \uplus \{n \mapsto (u, p, q)\}, \{\}, [\,]\rangle$;
2) $i = \mathsf{xhr\_resp}(n, ck, u, u', v)$ with $l \vdash_{\diamond} i$;
3) $P = \langle W, K', N \uplus \{n \mapsto (u', p, q')\}, \{\}, \mathsf{doc\_req}(ck'', u')\rangle$;
4) $q = \checkmark \Rightarrow K' = \mathit{sec\_upd\_ck}(K, u, ck)$;
5) $q = \times \Rightarrow K' = K$;
6) $q = \checkmark \wedge \mathit{url\_label}(u) = \mathit{url\_label}(u') \Rightarrow ck'' = \mathit{get\_http\_ck}(K', u') \wedge q' = \checkmark$;
7) $q = \times \vee \mathit{url\_label}(u) \neq \mathit{url\_label}(u') \Rightarrow ck'' = \{\} \wedge q' = \times$.

The case is analogous to (I-DOCREDIR), but we additionally have to prove that condition 6 of Definition 7 is preserved, which is ensured by assumption 6 above. Specifically, we observe that for the new network connection $\{n \mapsto (u', p, q')\}$ we have $q' = \checkmark$ only if $\mathit{url\_label}(u) = \mathit{url\_label}(u')$, hence the invariant follows.

We now prove the second point. If $P \xrightarrow{o} Q$ was derived by (O-COMPLETE), the conclusion is easy, since the internal state of the browser does not change, but for the task list, which will be empty. Let then $P \xrightarrow{o} Q$ be proved by (O-MIRROR), we need to prove a variant of the statement where $P \xrightarrow{o} Q$ has been replaced by $P \overset{o}{\mapsto} Q$. The proof is by induction on the derivation of $P \overset{o}{\mapsto} Q$:

Case (O-APP): we have:

1) $P = \langle W, K, N, \{p \mapsto ((\lambda x.e)\ v, \hat{l})\}, [\,]\rangle$;
2) $Q = \langle W, K, N, \{p \mapsto (e\{v/x\}, \hat{l})\}, [\,]\rangle$.

The only invariants which are not trivial to preserve are on the task list. Since $\mathit{fn}(e\{v/x\}) = \mathit{fn}(e) \cup \mathit{fn}(v)$, both conditions 7 and 8 follow by the very same conditions of the typing assumption;

Case (O-LETCTX): we have:

1) $P = \langle W, K, N, \{p \mapsto (\mathsf{let}\ x = e'\ \mathsf{in}\ e, \hat{l})\}, [\,]\rangle$;
2) $Q = \langle W', K', N', \{p \mapsto (\mathsf{let}\ x = e''\ \mathsf{in}\ e, \hat{l})\}, [\,]\rangle$;
3) $\langle W, K, N, \{p \mapsto (e', \hat{l})\}, [\,]\rangle \overset{o}{\mapsto} \langle W', K', N', \{p \mapsto (e'', \hat{l})\}, [\,]\rangle$.

By conditions 7 and 8 of the typing assumption, we know that:

$$\hat{l} \sqsubseteq l \Rightarrow \forall n \in \mathit{fn}(\mathsf{let}\ x = e'\ \mathsf{in}\ e) = \mathit{fn}(e') \cup \mathit{fn}(e) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$$
$$\hat{l} \not\sqsubseteq l \Rightarrow \forall n \in \mathit{fn}(\mathsf{let}\ x = e'\ \mathsf{in}\ e) = \mathit{fn}(e') \cup \mathit{fn}(e) : \exists l' : (l' \sqsubseteq \hat{l} \vee l' \sqsubseteq l) \wedge n \in \mathcal{N}_{l'}$$

Since $\mathit{fn}(e') \subseteq \mathit{fn}(\mathsf{let}\ x = e'\ \mathsf{in}\ e)$, it is easy to show that $l \models \langle W, K, N, \{p \mapsto (e', \hat{l})\}, [\,]\rangle$ holds true. We apply the induction hypothesis and we get $l \models \langle W', K, N', \{p \mapsto (e'', \hat{l})\}, [\,]\rangle$. This implies that:

$$\hat{l} \sqsubseteq l \Rightarrow \forall n \in \mathit{fn}(e'') : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$$
$$\hat{l} \not\sqsubseteq l \Rightarrow \forall n \in \mathit{fn}(e'') : \exists l' : (l' \sqsubseteq \hat{l} \vee l' \sqsubseteq l) : n \in \mathcal{N}_{l'}$$

We can show that $l \models Q$ holds true by using the implications above. Indeed, we have:

$$\hat{l} \sqsubseteq l \Rightarrow \forall n \in \mathit{fn}(e'') \cup \mathit{fn}(e) = \mathit{fn}(\mathsf{let}\ x = e''\ \mathsf{in}\ e) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$$
$$\hat{l} \not\sqsubseteq l \Rightarrow \forall n \in \mathit{fn}(e'') \cup \mathit{fn}(e) = \mathit{fn}(\mathsf{let}\ x = e''\ \mathsf{in}\ e) : \exists l' : (l' \sqsubseteq \hat{l} \vee l' \sqsubseteq l) \wedge n \in \mathcal{N}_{l'}$$

Case (O-LET): analogous to case (O-APP);

Case (O-GET): we have:

1) $P = \langle W, K, N, \{p \mapsto (k?, \hat{l})\}, [\,]\rangle$;
2) $Q = \langle W, K, N, \{p \mapsto (n, \hat{l})\}, [\,]\rangle$;
3) $W(p) = (u, h, h', ck')$ with $u = (\pi, d, v)$ and $\exists ck : K(d) = ck \wedge ck(k) = (n, f) \wedge f \in \{\bot, \mathtt{S}\}$.

Since $\mathit{cookie\_label}(d, f) = \bot$ for $f \in \{\bot, \mathtt{S}\}$, by condition 4 of the typing assumption we know that $n \in \mathcal{N}_{l'}$ for some $l' \sqsubseteq l$. Hence, both conditions 7 and 8 of Definition 7 are satisfied and we are done;

Case (O-GETFAIL): we have:

1) $P = \langle W, K, N, \{p \mapsto (k?, \hat{l})\}, [\,]\rangle$;
2) $Q = \langle W, K, N, \{p \mapsto ((), \hat{l})\}, [\,]\rangle$.

The conclusion immediately follows by the assumption $l \models P$, since the expression $()$ has no free name;

Case (O-SET): we have:

1) $P = \langle W, K, N, \{p \mapsto (k!\langle n, f\rangle, \bot)\}, [\,]\rangle$ with $f \in \{\bot, \mathtt{S}\}$;
2) $Q = \langle W, K', N, \{p \mapsto ((), \bot)\}, [\,]\rangle$ with $K' = \mathit{upd\_ck}(K, d, \{k \mapsto (n, f)\})$;
3) there does not exist ck such that $K(d) = ck$ and $ck(k) = (m, f')$ with $f' \in \{\mathtt{H}, \top\}$.

The only interesting property to show is on the cookie jar. Since $\bot \sqsubseteq l$, we know that $\forall m \in \mathit{fn}(k!\langle n, f\rangle) : \exists l' \sqsubseteq l : m \in \mathcal{N}_{l'}$ by condition 7 of Definition 7. Since $\mathit{cookie\_label}(d, f) = \bot$ for $f \in \{\bot, \mathtt{S}\}$, the (possible) inclusion of the new binding $\{k \mapsto (n, f)\}$ in K' cannot disrupt condition 4 of Definition 7 and we conclude;

Case (O-SETFAIL): analogous to case (O-GETFAIL);

Case (O-XHR): we have:

1) $P = \langle W \uplus \{p \mapsto (u', h, h', q)\}, K, N, \{p \mapsto (\mathsf{xhr}(u, \lambda x.e), \hat{l})\}, [\,]\rangle$;
2) $Q = \langle W \uplus \{p \mapsto (u', h, h' \uplus \{n \mapsto \lambda x.e\}, q)\}, K, N \uplus \{n \mapsto (u, p, q')\}, \{p \mapsto ((), \hat{l})\}, [\,]\rangle$;
3) $\hat{l} \neq \bot \Rightarrow \hat{l} = \mathit{url\_label}(u) = \mathit{url\_label}(u') \wedge q = \checkmark$;
4) $q = \checkmark \wedge \mathit{url\_label}(u') = \mathit{url\_label}(u) \Rightarrow ck = \mathit{get\_http\_ck}(K, u) \wedge q' = \checkmark$;
5) $q = \times \vee \mathit{url\_label}(u') \neq \mathit{url\_label}(u) \Rightarrow ck = \{\} \wedge q' = \times$.

Most of the properties to show are immediate, the only interesting points are proving the invariant for the new XHR handler $\{n \mapsto \lambda x.e\}$ and the new network connection $\{n \mapsto (u, p, q')\}$. We start with the XHR handler, we have to show:

$$\mathit{url\_label}(u') \sqsubseteq l \Rightarrow \forall m \in \mathit{fn}(\{n \mapsto \lambda x.e\}) : \exists l' \sqsubseteq l : m \in \mathcal{N}_{l'} \tag{7}$$

$$\mathit{url\_label}(u') = l'' \not\sqsubseteq l \Rightarrow \forall m \in \mathit{fn}(\{n \mapsto \lambda x.e\}) : \exists l' : (l' \sqsubseteq l'' \vee l' \sqsubseteq l) \wedge m \in \mathcal{N}_{l'} \tag{8}$$

Without loss of generality, we assume $n \in \mathcal{N}_\bot$. Let $\mathit{url\_label}(u') \sqsubseteq l$, we distinguish two cases. If $\hat{l} = \bot$, then by condition 7 of Definition 7 we know that $\forall n \in \mathit{fn}(\mathsf{xhr}(u, \lambda x.e)) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$, hence we have that $\forall m \in \mathit{fn}(\{n \mapsto \lambda x.e\}) : \exists l' \sqsubseteq l : m \in \mathcal{N}_{l'}$. Otherwise, assume that $\hat{l} \neq \bot$, then we must have $\mathit{url\_label}(u) = \mathit{url\_label}(u') = \hat{l}$ by assumption 3 above. Hence, we know that $\hat{l} \sqsubseteq l$ and by condition 7 of Definition 7 we know that $\forall n \in \mathit{fn}(\mathsf{xhr}(u, \lambda x.e)) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$. Thus, we have again that $\forall m \in \mathit{fn}(\{n \mapsto \lambda x.e\}) : \exists l' \sqsubseteq l : m \in \mathcal{N}_{l'}$ and we conclude the proof of (7).

Let now $\mathit{url\_label}(u') = l'' \not\sqsubseteq l$. If $\hat{l} = \bot$, then by condition 7 of Definition 7 we know that $\forall n \in \mathit{fn}(\mathsf{xhr}(u, \lambda x.e)) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$, hence we have that $\forall m \in \mathit{fn}(\{n \mapsto \lambda x.e\}) : \exists l' \sqsubseteq l : m \in \mathcal{N}_{l'}$. Otherwise, assume that $\hat{l} \neq \bot$, then we must have $\mathit{url\_label}(u) = \mathit{url\_label}(u') = \hat{l}$ by assumption 3 above. Hence, we know that $\hat{l} = l'' \not\sqsubseteq l$ and by condition 8 of Definition 7 we know that $\forall n \in \mathit{fn}(\mathsf{xhr}(u, \lambda x.e)) : \exists l' : (l' \sqsubseteq l'' \vee l' \sqsubseteq l) \wedge n \in \mathcal{N}_{l'}$. Thus, we have that $\forall m \in \mathit{fn}(\{n \mapsto \lambda x.e\}) : \exists l' : (l' \sqsubseteq l'' \vee l' \sqsubseteq l) \wedge m \in \mathcal{N}_{l'}$ and we conclude the proof of (8).

Now we focus on the new network connection $\{n \mapsto (u, p, q')\}$. We have to show that $q' = \checkmark$ implies $\mathit{url\_label}(u) = \mathit{url\_label}(u')$, which is ensured by assumption 4 above. Specifically, we observe that for the new network connection $\{n \mapsto (u, p, q')\}$ we have $q' = \checkmark$ only if $\mathit{url\_label}(u) = \mathit{url\_label}(u')$, hence the invariant follows;

Case (O-LOGIN): we have:

1) $P = \langle W, K, N, \{p \mapsto (\mathsf{auth}(u, c), \hat{l})\}, [\,]\rangle$;
2) $W(p) = (u', h, h', \checkmark)$;
3) $\hat{l} = \mathit{url\_label}(u) = \mathit{url\_label}(u')$;
4) $Q = \langle W, K, N \uplus \{n \mapsto (u, (), \checkmark)\}, \{p \mapsto ((), \hat{l})\}, [\,]\rangle$.

The conclusion immediately follows by the assumption $l \models P$, since the new network connection cannot violate the invariant and the new running expression $()$ has no free name;

Case (O-FLUSH): we have:

1) $P = \langle W, K, N, T, o\rangle$;
2) $Q = \langle W, K, N, T, [\,]\rangle$.

The conclusion immediately follows by the assumption $l \models P$, since only the output buffer changes and becomes empty.

□

We now generalize typing from the browser to attacked extended states, to include a number of invariants on the input stream, the trust function, and the attacker power.

Definition 9 (Typing)**.** *Let $\sigma = \langle Q, I, \tau, M\rangle$. We write $l \models \sigma$ if and only if:*

1) $l \models Q$;
2) $\forall i \in I : \vdash_\diamond i$;
3) $\forall o \in \mathcal{O} : \tau(o) = l' \not\sqsubseteq l \Rightarrow o \in \{*\_\mathsf{req}(ck, u) \mid ck\_vals(ck) \cap \mathcal{N}_{l'} \neq \emptyset\}$;
4) $\forall o \in \mathcal{O} : \tau(o) = \mathsf{evil} \Rightarrow \exists n \in ck\_vals(o) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$;
5) $\forall o \in M : ev\_label(o) = l' \not\sqsubseteq l \Rightarrow ck\_vals(o) \subseteq \mathcal{N}_{l'}$;
6) $\forall \alpha \in M : ev\_label(\alpha) \sqsubseteq l \Rightarrow \forall n \in fn(\alpha) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$;
7) $\forall o \in M : \tau(o) = \mathsf{evil} \Rightarrow \forall n \in fn(o) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$.

We introduce a new definition for the set of events intercepted/overheard by the attacker, which ensures that the secrecy of "high" names is preserved. This is convenient to state a number of the next results.

Definition 10 (Consistent Set of Events)**.** *We write $\tau, l \vdash_\diamond M$ if and only if:*

1) $\forall \alpha \in M : ev\_label(\alpha) \sqsubseteq l \Rightarrow \forall n \in fn(\alpha) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$;
2) $\forall o \in M : \tau(o) = \mathsf{evil} \Rightarrow \forall n \in fn(o) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$.

Lemma 16 (Tainted Names)**.** *If $\tau, l \vdash_\diamond M$ and $\tau, l, M \Vdash n$, then $n \in \mathcal{N}_{l'}$ for some $l' \sqsubseteq l$.*

Proof. By a case analysis on the rule applied to prove $\tau, l, M \Vdash n$. □

Lemma 17 (Tainted Input)**.** *If $\tau, l \vdash_\diamond M$ and $\tau, l, M \Vdash i$, then $ev\_label(i) \sqsubseteq l \wedge \forall n \in fn(i) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$.*

Proof. By a case analysis on the rule applied to prove $\tau, l, M \Vdash i$, using Lemma 16. □

Lemma 18 (Tainted Output)**.** *If $\tau, l \vdash_\diamond M$ and $\tau, l, M \Vdash o$, then $\forall n \in fn(o) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$.*

Proof. By a case analysis on the rule applied to prove $\tau, l, M \Vdash o$, using Lemma 16. □

Lemma 19 (Trusted Login)**.** *Let $l \models P$ and $P \xrightarrow{o} Q$. If $\tau \xrightarrow{o} \tau'$ and:*

$$\forall o' \in \mathcal{O} : \tau(o') = l' \not\sqsubseteq l \Rightarrow o' \in \{*\_\mathsf{req}(ck, u) \mid ck\_vals(ck) \cap \mathcal{N}_{l'} \neq \emptyset\},$$

then:

$$\forall o' \in \mathcal{O} : \tau'(o') = l' \not\sqsubseteq l \Rightarrow o' \in \{*\_\mathsf{req}(ck, u) \mid ck\_vals(ck) \cap \mathcal{N}_{l'} \neq \emptyset\}.$$

Proof. By a case analysis on the rule applied to prove $\tau \xrightarrow{o} \tau'$:

Case (A-FIX): we have $\tau \xrightarrow{\mathsf{login}(ck,u,c)} \tau'$, where:

1) there exists $k = \kappa(u)$ such that $ck(k) = (n, \_)$ for some name n;
2) $\rho(c) = url\_label(u)$ or $\rho(c) = \mathsf{evil}$;
3) $\tau' = \tau \sqcup \tau_{u,n,c}$;
4) $\tau_{u,n,c}(o') = \rho(c)$ iff $o' \in \{*\_\mathsf{req}(ck', u') \mid domain(u) = domain(u') \wedge ck'(\kappa(u)) = n \wedge \tau(o) \sqsubseteq \rho(c)\}$.

If $\rho(c) \sqsubseteq l$, the conclusion is immediate, since any change of trust is bounded above by l. Let then $\rho(c) = l' \not\sqsubseteq l$. Since we assume that $l \models P$, we know that $\rho(c) \neq \mathsf{evil}$ by Lemma 8, thus we must have $\rho(c) = l' = url\_label(u)$ by condition 2 above. By Lemma 6 we then know that $n \in \mathcal{N}_{l'}$ and we can conclude;

Case (A-SRV): we have $\tau \xrightarrow{\mathsf{login}(ck,u,c)} \tau'$, where:

1) the server picks a name $n \in \mathcal{N}_{\rho(c)}$;
2) $\rho(c) = url\_label(u)$ or $\rho(c) = \mathsf{evil}$;
3) $\tau' = \tau \sqcup \tau_{u,n,c}$;
4) $\tau_{u,n,c}(o') = \rho(c)$ iff $o \in \{*\_\mathsf{req}(ck', u') \mid domain(u) = domain(u') \wedge ck'(\kappa(u)) = n \wedge \tau(o) \sqsubseteq \rho(c)\}$.

If $\rho(c) \sqsubseteq l$, the conclusion is immediate, since any change of trust is bounded above by l. Let then $\rho(c) = l' \not\sqsubseteq l$. Since we assume that $l \models P$, we know that $\rho(c) \neq \mathsf{evil}$ by Lemma 8, thus we must have $\rho(c) = l' = url\_label(u)$ by condition 2 above. Hence, we know that $n \in \mathcal{N}_{l'}$ by condition 1 above and we can conclude;

Case (A-NIL): we have $\tau' = \tau$ and the result is trivial.

□

Lemma 20 (Login CSRF Prevention and Trust)**.** *Let $l \models P$ and $P \xrightarrow{o} Q$. If $\tau \xrightarrow{o} \tau'$ and there exists o' such that $\tau'(o') = \mathsf{evil}$, then $\tau(o') = \mathsf{evil}$.*

Proof. By a case analysis on the rule applied to prove $\tau \xrightarrow{o} \tau'$, using Lemma 8. Indeed, the lemma states that $\rho(c) \neq \mathsf{evil}$ for any event $\mathsf{login}(ck, u, c)$ output by the browser. Since any login operation may change the trust of a given output event only by raising it to $\rho(c)$, we can conclude. □

Lemma 21 (Tainted Login)**.** *Let $\tau, l \vdash_\diamond M$ and assume that $\tau, l, M \Vdash o$. If $\tau \xrightarrow{o} \tau'$ and there exists o' such that $\tau'(o') = l' \not\sqsubseteq l$, then $\tau(o') = l'$.*

Proof. By a case analysis on the rule applied to prove $\tau \xrightarrow{o} \tau'$. If the reduction rule is (A-NIL), we have $\tau' = \tau$ and the conclusion is trivial. Otherwise, let $o = \mathsf{login}(ck, u, c)$ and assume that either (A-FIX) or (A-SRV) was the applied reduction rule. We first observe that the assumptions $\tau, l \vdash_\diamond M$ and $\tau, l, M \Vdash o$ imply that $\forall n \in \mathit{fn}(o) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$ by Lemma 18. Hence, we know in particular that $c \in \mathcal{N}_{l'}$ for some $l' \sqsubseteq l$, which implies that $\rho(c) \sqsubseteq l$ by the assumption that the function ρ respects the partitioning of names. Since the increase of trust is bounded above by $\rho(c) \sqsubseteq l$, the desired conclusion follows. □

Lemma 22 (Evil Login)**.** *Let $\tau, l \vdash_\diamond M$ and assume that $\tau, l, M \Vdash o$. Let τ be a trust function such that:*

$$\forall o \in \mathcal{O} : \tau(o) = \mathsf{evil} \Rightarrow \exists n \in \mathit{ck\_vals}(o) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}.$$

If $\tau \xrightarrow{o} \tau'$, then we have:

$$\forall o \in \mathcal{O} : \tau'(o) = \mathsf{evil} \Rightarrow \exists n \in \mathit{ck\_vals}(o) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}.$$

Proof. By a case analysis on the rule applied to prove $\tau \xrightarrow{o} \tau'$:

Case (A-FIX): we have $\tau \xrightarrow{\mathsf{login}(ck,u,c)} \tau'$, where:

1) there exists $k = \kappa(u)$ such that $ck(k) = (n, \_)$ for some name n;
2) $\rho(c) = \mathit{url\_label}(u)$ or $\rho(c) = \mathsf{evil}$;
3) $\tau' = \tau \sqcup \tau_{u,n,c}$;
4) $\tau_{u,n,c}(o') = \rho(c)$ iff $o \in \{*\_\mathsf{req}(ck', u') \mid \mathit{domain}(u) = \mathit{domain}(u') \wedge ck'(\kappa(u)) = n \wedge \tau(o) \sqsubseteq \rho(c)\}$.

If $\rho(c) \neq \mathsf{evil}$, the conclusion immediately follows by our assumption on τ, since any login operation may change the trust of a given output event only by raising it to $\rho(c)$. Let then $\rho(c) = \mathsf{evil}$. Since $\tau, l \vdash_\diamond M$ holds true, by Lemma 18 we know that $\forall m \in \mathit{ck\_vals}(ck) : \exists l' \sqsubseteq l : m \in \mathcal{N}_{l'}$. In particular, this implies that there exists $l' \sqsubseteq l$ such that $n \in \mathcal{N}_{l'}$. Given that the reduction step may set to evil the trust of an output event o' only if $n \in \mathit{ck\_vals}(o')$, the conclusion follows;

Case (A-SRV): we have $\tau \xrightarrow{\mathsf{login}(ck,u,c)} \tau'$, where:

1) the server picks a name $n \in \mathcal{N}_{\rho(c)}$;
2) $\rho(c) = \mathit{url\_label}(u)$ or $\rho(c) = \mathsf{evil}$;
3) $\tau' = \tau \sqcup \tau_{u,n,c}$;
4) $\tau_{u,n,c}(o') = \rho(c)$ iff $o \in \{*\_\mathsf{req}(ck', u') \mid \mathit{domain}(u) = \mathit{domain}(u') \wedge ck'(\kappa(u)) = n \wedge \tau(o) \sqsubseteq \rho(c)\}$.

If $\rho(c) \neq \mathsf{evil}$, the conclusion immediately follows by our assumption on τ, since any login operation may change the trust of a given output event only by raising it to $\rho(c)$. Let then $\rho(c) = \mathsf{evil}$, by assumption 1 above we must have $n \in \mathcal{N}_{\mathsf{evil}}$. Since $\tau, l \vdash_\diamond M$ holds true, by Lemma 18 we know that there exists $l' \sqsubseteq l$ such that $c \in \mathcal{N}_{l'}$. Since we assume that the function ρ respects the partitioning of names, we have $\rho(c) = \mathsf{evil} \sqsubseteq l' \sqsubseteq l$. Given that the reduction step may set to evil the trust of an output event o' only if $n \in \mathit{ck\_vals}(o')$, the conclusion follows;

Case (A-NIL): we have $\tau' = \tau$ and the result is trivial.

□

Lemma 23 (Preventing Compromise)**.** *Let $\tau, l \vdash_\diamond M$ and assume that:*

$$\forall o \in M : \mathit{ev\_label}(o) = l' \not\sqsubseteq l \Rightarrow \mathit{ck\_vals}(o) \subseteq \mathcal{N}_{l'}.$$

If $\tau \xrightarrow{o} \tau'$ for some o such that $\tau, l, M \Vdash o$, then we have:

$$\forall o \in M : \tau'(o) = \mathsf{evil} \Rightarrow \forall n \in \mathit{fn}(o) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}.$$

Proof. First, we unfold the notation $\tau, l \vdash_\diamond M$ to the corresponding two conditions:

$$\forall \alpha \in M : \mathit{ev\_label}(\alpha) \sqsubseteq l \Rightarrow \forall n \in \mathit{fn}(\alpha) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}, \tag{9}$$

$$\forall o \in M : \tau(o) = \mathsf{evil} \Rightarrow \forall n \in \mathit{fn}(o) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}. \tag{10}$$

We then proceed by a case analysis on the rule applied to prove $\tau \xrightarrow{o} \tau'$:

Case (A-FIX): we have $\tau \xrightarrow{\mathsf{login}(ck,u,c)} \tau'$, where:

1) there exists $k = \kappa(u)$ such that $ck(k) = (n, \_)$ for some name n;
2) $\rho(c) = \mathit{url\_label}(u)$ or $\rho(c) = \mathsf{evil}$;
3) $\tau' = \tau \sqcup \tau_{u,n,c}$;
4) $\tau_{u,n,c}(o') = \rho(c)$ iff $o \in \{*\_\mathsf{req}(ck', u') \mid \mathit{domain}(u) = \mathit{domain}(u') \wedge ck'(\kappa(u)) = n \wedge \tau(o) \sqsubseteq \rho(c)\}$.

Let $o' \in M$, we distinguish two cases. If $\mathit{ev\_label}(o') \sqsubseteq l$, the conclusion is immediate by (9). Let then $\mathit{ev\_label}(o') = l' \not\sqsubseteq l$, by our assumption on the output events in M we know that $\mathit{ck\_vals}(o') \subseteq \mathcal{N}_{l'}$. Since we know that $\tau, l, M \Vdash \mathsf{login}(ck, u, c)$ holds true, by Lemma 18 we know that for all $m \in \mathit{ck\_vals}(ck)$ there exists $\hat{l} \sqsubseteq l$ such that $m \in \mathcal{N}_{\hat{l}}$. Hence, we know that $\mathit{ck\_vals}(ck) \cap \mathit{ck\_vals}(o') = \emptyset$ and the trust of o' cannot change after this login operation, so the conclusion follows by (10);

Case (A-SRV): we have $\tau \xrightarrow{\mathsf{login}(ck,u,c)} \tau'$, where:

1) the server picks a name $n \in \mathcal{N}_{\rho(c)}$;
2) $\rho(c) = \mathit{url\_label}(u)$ or $\rho(c) = \mathsf{evil}$;
3) $\tau' = \tau \sqcup \tau_{u,n,c}$;
4) $\tau_{u,n,c}(o') = \rho(c)$ iff $o \in \{*\_\mathsf{req}(ck', u') \mid \mathit{domain}(u) = \mathit{domain}(u') \wedge ck'(\kappa(u)) = n \wedge \tau(o) \sqsubseteq \rho(c)\}$.

Let $o' \in M$, we distinguish two cases. If $\mathit{ev\_label}(o') \sqsubseteq l$, the conclusion is immediate by (9). Let then $\mathit{ev\_label}(o') = l' \not\sqsubseteq l$, by our assumption on the output events in M we know that $\mathit{ck\_vals}(o') \subseteq \mathcal{N}_{l'}$. Since we know that $\tau, l, M \Vdash \mathsf{login}(ck, u, c)$ holds true, by Lemma 18 we know that there exists $\hat{l} \sqsubseteq l$ such that $c \in \mathcal{N}_{\hat{l}}$. Given that the function ρ respects the partitioning of names, we know that $\rho(c) \sqsubseteq \hat{l} \sqsubseteq l$, hence we must have $\rho(c) \neq l'$. Since we know that $n \in \mathcal{N}_{\rho(c)}$ by assumption 1 above, we observe that $n \notin \mathit{ck\_vals}(o')$ and the trust of o' cannot change after this login operation, so the conclusion follows by (10);

Case (A-NIL): we have $\tau' = \tau$ and the result is trivial by (10).

□

Theorem 2 (Subject Reduction)**.** *If $l \models \sigma$ and $l \vdash \sigma \xrightarrow{\alpha} \sigma'$, then $l \models \sigma'$.*

Proof. By a case analysis on the rule applied to prove $l \vdash \sigma \xrightarrow{\alpha} \sigma'$:

Case (AS-IN): we have $\sigma = \langle C, i :: I, \tau, M\rangle$ and $\sigma' = \langle P, I, \tau, M\rangle$ with $\alpha = i$ and $C \xrightarrow{i} P$. By the assumption $l \models \sigma$, we have:

1) $l \models C$;
2) $\forall i' \in i :: I : \vdash_\diamond i'$;
3) $\forall o \in \mathcal{O} : \tau(o) = l' \not\sqsubseteq l \Rightarrow o \in \{*\_\mathsf{req}(ck, u) \mid \mathit{ck\_vals}(ck) \cap \mathcal{N}_{l'} \neq \emptyset\}$;
4) $\forall o \in \mathcal{O} : \tau(o) = \mathsf{evil} \Rightarrow \exists n \in \mathit{ck\_vals}(o) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$;
5) $\forall o \in M : \mathit{ev\_label}(o) = l' \not\sqsubseteq l \Rightarrow \mathit{ck\_vals}(o) \subseteq \mathcal{N}_{l'}$;
6) $\forall \alpha \in M : \mathit{ev\_label}(\alpha) \sqsubseteq l \Rightarrow \forall n \in \mathit{fn}(\alpha) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$;
7) $\forall o \in M : \tau(o) = \mathsf{evil} \Rightarrow \forall n \in \mathit{fn}(o) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$.

We want to show that $l \models \sigma'$. The only non-trivial condition to preserve is 1, i.e., we have to prove that $l \models P$. Since we know that $l \models C$ by condition 1 and $l \vdash_\diamond i$ by condition 2, the desired conclusion follows by Lemma 15;

Case (AS-OUT): we have $\sigma = \langle P, I, \tau, M\rangle$ and $\sigma' = \langle Q, I, \tau', M\rangle$ with $\alpha = o$ and $P \xrightarrow{o} Q$ and $\tau \xrightarrow{o} \tau'$. By the assumption $l \models \sigma$, we have:

1) $l \models P$;
2) $\forall i \in I : \vdash_\diamond i$;

3) $\forall o \in \mathcal{O} : \tau(o) = l' \not\sqsubseteq l \Rightarrow o \in \{*\_\mathsf{req}(ck, u) \mid ck\_vals(ck) \cap \mathcal{N}_{l'} \neq \emptyset\}$;
4) $\forall o \in \mathcal{O} : \tau(o) = \mathsf{evil} \Rightarrow \exists n \in ck\_vals(o) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$;
5) $\forall o \in M : ev\_label(o) = l' \not\sqsubseteq l \Rightarrow ck\_vals(o) \subseteq \mathcal{N}_{l'}$;
6) $\forall \alpha \in M : ev\_label(\alpha) \sqsubseteq l \Rightarrow \forall n \in fn(\alpha) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$;
7) $\forall o \in M : \tau(o) = \mathsf{evil} \Rightarrow \forall n \in fn(o) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$.

We want to show that $l \models \sigma'$. We have four non-trivial conditions to preserve, specifically 1, 3, 4 and 7: we show them separately. We start by condition 1, i.e., we have to prove that $l \models Q$: this follows by the assumption $l \models P$, using Lemma 15. We now move to proving condition 3, i.e., we have to show that:

$$\forall o \in \mathcal{O} : \tau'(o) = l' \not\sqsubseteq l \Rightarrow o \in \{*\_\mathsf{req}(ck, u) \mid ck\_vals(ck) \cap \mathcal{N}_{l'} \neq \emptyset\},$$

which follows immediately by Lemma 19. Finally, we focus on conditions 4 and 7, i.e., we have to prove:

$$\forall o \in \mathcal{O} : \tau'(o) = \mathsf{evil} \Rightarrow \exists n \in ck\_vals(o) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$$
$$\forall o \in M : \tau'(o) = \mathsf{evil} \Rightarrow \forall n \in fn(o) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'},$$

which follow immediately by Lemma 20, given conditions 4 and 7 above;

Case (AS-GETIN): we have $\sigma = \langle Q, i :: I, \tau, M \rangle$ and $\sigma' = \langle Q, I, \tau, M \cup \{i\} \rangle$ with $\alpha = \bullet$ and $\tau, l \dagger i$ being true by the premise of the rule. By the assumption $l \models \sigma$, we have:

1) $l \models Q$;
2) $\forall i' \in i :: I : \vdash_\diamond i'$;
3) $\forall o \in \mathcal{O} : \tau(o) = l' \not\sqsubseteq l \Rightarrow o \in \{*\_\mathsf{req}(ck, u) \mid ck\_vals(ck) \cap \mathcal{N}_{l'} \neq \emptyset\}$;
4) $\forall o \in \mathcal{O} : \tau(o) = \mathsf{evil} \Rightarrow \exists n \in ck\_vals(o) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$;
5) $\forall o \in M : ev\_label(o) = l' \not\sqsubseteq l \Rightarrow ck\_vals(o) \subseteq \mathcal{N}_{l'}$;
6) $\forall \alpha \in M : ev\_label(\alpha) \sqsubseteq l \Rightarrow \forall n \in fn(\alpha) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$;
7) $\forall o \in M : \tau(o) = \mathsf{evil} \Rightarrow \forall n \in fn(o) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$.

We want to show that $l \models \sigma'$. The only non-trivial condition to preserve is 6. Specifically, we have to prove:

$$\forall \alpha \in M \cup \{i\} : ev\_label(\alpha) \sqsubseteq l \Rightarrow \forall n \in fn(\alpha) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}.$$

Due to condition 6 above, the only event in $M \cup \{i\}$ which can break the desired property is i itself. Since we have $\vdash_\diamond i$ by condition 2 above, we can conclude by Lemma 14;

Case (AS-GETOUT): we have $\sigma = \langle P, I, \tau, M \rangle$ and $\sigma' = \langle Q, I, \tau, M \cup \{o\} \rangle$ with $\alpha = \bullet$ and $P \xrightarrow{o} Q$. By the premises of the reduction rule, we also know that $\tau, l \dagger o$. By the assumption $l \models \sigma$, we have:

1) $l \models P$;
2) $\forall i \in I : \vdash_\diamond i$;
3) $\forall o \in \mathcal{O} : \tau(o) = l' \not\sqsubseteq l \Rightarrow o \in \{*\_\mathsf{req}(ck, u) \mid ck\_vals(ck) \cap \mathcal{N}_{l'} \neq \emptyset\}$;
4) $\forall o \in \mathcal{O} : \tau(o) = \mathsf{evil} \Rightarrow \exists n \in ck\_vals(o) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$;
5) $\forall o \in M : ev\_label(o) = l' \not\sqsubseteq l \Rightarrow ck\_vals(o) \subseteq \mathcal{N}_{l'}$;
6) $\forall \alpha \in M : ev\_label(\alpha) \sqsubseteq l \Rightarrow \forall n \in fn(\alpha) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$;
7) $\forall o \in M : \tau(o) = \mathsf{evil} \Rightarrow \forall n \in fn(o) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$.

We want to show that $l \models \sigma'$. Conditions 2, 3 and 4 are trivially preserved, we show the remaining conditions separately. We start by condition 1, i.e., we have to prove that $l \models Q$: this follows by the assumption $l \models P$, using Lemma 15. As to condition 5, we have to show that:

$$\forall o' \in M \cup \{o\} : ev\_label(o') = l' \not\sqsubseteq l \Rightarrow ck\_vals(o') \subseteq \mathcal{N}_{l'}.$$

Due to condition 5 above, the only output event in $M \cup \{o\}$ which can break the desired property is o itself, hence the conclusion follows by Lemma 6. We then move to condition 6, i.e., we have to prove that:

$$\forall \alpha \in M \cup \{o\} : ev\_label(\alpha) \sqsubseteq l \Rightarrow \forall n \in fn(\alpha) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}.$$

Due to condition 6 above, the only output event in $M \cup \{o\}$ which can break the desired property is o itself, hence the conclusion follows by Lemma 4. Finally, we need to show condition 7, i.e., we have to prove that:

$$\forall o' \in M \cup \{o\} : \tau(o') = \mathsf{evil} \Rightarrow \forall n \in fn(o') : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}.$$

Due to condition 7 above, the only output event in $M \cup \{o\}$ which can break the desired property is o itself. Assume then that $\tau(o) = \mathsf{evil}$, then by Lemma 7 we know that $ev\_label(o) \sqsubseteq l$. Hence, the desired conclusion follows by Lemma 4;

Case (AS-HEARIN): analogous to case (AS-GETIN);

Case (AS-HEAROUT): the proof combines the reasoning performed for cases (AS-OUT) and (AS-GETOUT);

Case (AS-SYNIN): we have $\sigma = \langle C, I, \tau, M\rangle$ and $\sigma' = \langle P, I, \tau, M\rangle$ with $\alpha = i$ and $C \xrightarrow{i} P$. By the premises of the reduction rule, we also know that $\tau, l, M \Vdash i$. By the assumption $l \models \sigma$, we have:

1) $l \models C$;
2) $\forall i \in I : \vdash_\diamond i$;
3) $\forall o \in \mathcal{O} : \tau(o) = l' \not\sqsubseteq l \Rightarrow o \in \{*\_\mathsf{req}(ck, u) \mid ck\_vals(ck) \cap \mathcal{N}_{l'} \neq \emptyset\}$;
4) $\forall o \in \mathcal{O} : \tau(o) = \mathsf{evil} \Rightarrow \exists n \in ck\_vals(o) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$;
5) $\forall o \in M : ev\_label(o) = l' \not\sqsubseteq l \Rightarrow ck\_vals(o) \subseteq \mathcal{N}_{l'}$;
6) $\forall \alpha \in M : ev\_label(\alpha) \sqsubseteq l \Rightarrow \forall n \in fn(\alpha) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$;
7) $\forall o \in M : \tau(o) = \mathsf{evil} \Rightarrow \forall n \in fn(o) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$.

We want to show that $l \models \sigma'$. The only non-trivial condition to preserve is 1, i.e., we have to prove that $l \models P$. Since we know that $l \vdash_\diamond i$ by Lemma 17 and $l \models C$ by condition 1, the desired conclusion follows by Lemma 15;

Case (AS-SYNOUT): we have $\sigma = \langle Q, I, \tau, M\rangle$ and $\sigma' = \langle Q, I, \tau', M\rangle$ with $\alpha = o$ and $\tau \xrightarrow{o} \tau'$. By the premises of the reduction rule, we also know that $\tau, l, M \Vdash o$. By the assumption $l \models \sigma$, we have:

1) $l \models Q$;
2) $\forall i \in I : \vdash_\diamond i$;
3) $\forall o \in \mathcal{O} : \tau(o) = l' \not\sqsubseteq l \Rightarrow o \in \{*\_\mathsf{req}(ck, u) \mid ck\_vals(ck) \cap \mathcal{N}_{l'} \neq \emptyset\}$;
4) $\forall o \in \mathcal{O} : \tau(o) = \mathsf{evil} \Rightarrow \exists n \in ck\_vals(o) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$;
5) $\forall o \in M : ev\_label(o) = l' \not\sqsubseteq l \Rightarrow ck\_vals(o) \subseteq \mathcal{N}_{l'}$;
6) $\forall \alpha \in M : ev\_label(\alpha) \sqsubseteq l \Rightarrow \forall n \in fn(\alpha) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$;
7) $\forall o \in M : \tau(o) = \mathsf{evil} \Rightarrow \forall n \in fn(o) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$.

We want to show that $l \models \sigma'$. We have three non-trivial conditions to preserve, specifically 3, 4 and 7. We first focus on condition 3, so we have to prove that:

$$\forall o \in \mathcal{O} : \tau'(o) = l' \not\sqsubseteq l \Rightarrow o \in \{*\_\mathsf{req}(ck, u) \mid ck\_vals(ck) \cap \mathcal{N}_{l'} \neq \emptyset\},$$

which follows immediately by Lemma 21, given condition 3 above. Now we move to condition 4, so we have to show that:

$$\forall o \in \mathcal{O} : \tau'(o) = \mathsf{evil} \Rightarrow \exists n \in ck\_vals(o) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'},$$

which follows immediately by Lemma 22, given condition 4 above. Finally, we move to condition 7, i.e., we have to show that:

$$\forall o \in M : \tau'(o) = \mathsf{evil} \Rightarrow \forall n \in fn(o) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'},$$

which follows immediately by Lemma 23.

□

D. Simulation

We now introduce a "simulation" between the attacked run and the original run, which is enough to prove session integrity for any well-formed trace.

Definition 11 (Tainted and Untainted Events)**.** *We define the following predicates over events:*

1) $l \vdash tainted(i)$ *if and only if* $ev\_label(i) \sqsubseteq l \wedge \forall n \in fn(i) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$*;*
2) $l \vdash untainted(i)$ *if and only if* $ev\_label(i) \not\sqsubseteq l \wedge \vdash_\diamond i$*;*
3) $l \vdash tainted(o)$ *if and only if* $\forall n \in fn(o) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$*;*
4) $l \vdash untainted(o)$ *if and only if* $\exists n \in fn(o) : \exists l' \not\sqsubseteq l : n \in \mathcal{N}_{l'}$.

Definition 12 (Corresponding Browsers)**.** *We say that* $Q = \langle W, K, N, T, O\rangle$ *and* $Q' = \langle W', K', N', T', O'\rangle$ *are* corresponding *for a security label* l*, written* $Q \sim_l Q'$*, if and only if:*

1) $l \models Q$ *and* $l \models Q'$*;*
2) $\mathit{url\_label}(u) \not\sqsubseteq l \Rightarrow (W(p) = (u, h, h', \checkmark) \Leftrightarrow W'(p) = (u, h, h', \checkmark))$*;*
3) $\mathit{cookie\_label}(d, f) \not\sqsubseteq l \Rightarrow (K(d) = ck \wedge ck(k) = (n, f) \Leftrightarrow K'(d) = ck' \wedge ck'(k) = (n, f))$*;*
4) $\mathit{url\_label}(u) \not\sqsubseteq l \Rightarrow (N(n) = (u, v, \checkmark) \Leftrightarrow N'(n) = (u, v, \checkmark))$*;*
5) $W(p) = (u, h, h', \checkmark) \wedge \mathit{url\_label}(u) \not\sqsubseteq l \Rightarrow (p \in \mathit{dom}(T) \cup \mathit{dom}(T') \Rightarrow (T(p) = T'(p) \wedge O = O' = [\,]))$*;*
6) $l \vdash \mathit{untainted}(o) \Rightarrow (O = o \Leftrightarrow O' = o)$*.*

Lemma 24 (Integrity for High HTTP Cookies)**.** *Let* K *and* K' *be two cookie jars such that:*

$$\mathit{cookie\_label}(d, f) \not\sqsubseteq l \Rightarrow (K(d) = ck \wedge ck(k) = (n, f) \Leftrightarrow K'(d) = ck' \wedge ck'(k) = (n, f)).$$

If $\mathit{url\_label}(u) = l' \not\sqsubseteq l$*, then* $\mathit{get\_http\_ck}(K, u) = \mathit{get\_http\_ck}(K', u)$*.*

Proof. Let $u = (\pi, d, v)$. We observe that $\{k \mapsto (n, f)\} \in \mathit{get\_http\_ck}(K, u)$ iff there exists ck such that $K(d) = ck$ and $ck(k) = (n, f)$ with $\mathit{cookie\_label}(d, f) = \mathit{url\_label}(u)$. Since we know that $\mathit{url\_label}(u) = l' \not\sqsubseteq l$, we have $\mathit{cookie\_label}(d, f) = l' \not\sqsubseteq l$, hence we conclude by our hypothesis on K and K'. □

Lemma 25 (Integrity for Secure High Updates)**.** *Let* K_1 *and* K_2 *be two cookie jars such that:*

$$\mathit{cookie\_label}(d, f) \not\sqsubseteq l \Rightarrow (K_1(d) = ck \wedge ck(k) = (n, f) \Leftrightarrow K_2(d) = ck' \wedge ck'(k) = (n, f)).$$

Let $\mathit{url\_label}(u) \not\sqsubseteq l$*. If* $K_1' = \mathit{sec\_upd\_ck}(K_1, u, \hat{ck})$ *and* $K_2' = \mathit{sec\_upd\_ck}(K_2, u, \hat{ck})$*, then:*

$$\mathit{cookie\_label}(d, f) \not\sqsubseteq l \Rightarrow (K_1'(d) = ck \wedge ck(k) = (n, f) \Leftrightarrow K_2'(d) = ck' \wedge ck'(k) = (n, f)).$$

Proof. Let $u = (\pi, d, v)$. By using Lemma 9 we observe that, whenever $\{k \mapsto (n, f)\} \in (\hat{ck} \nearrow \pi)$, we have $\mathit{cookie\_label}(d, f) = \mathit{url\_label}(u)$. Since the secure update operation fails in storing a new cookie only if the cookie jar already contains a cookie with a higher label, the storage of (some of) the cookies in $\hat{ck}$ fails in K_1 if and only if it fails in K_2, hence the conclusion follows by our assumptions on K_1 and K_2. □

Lemma 26 (Integrity for Secure Low Updates)**.** *Let* K_1 *and* K_2 *be two cookie jars such that:*

$$\mathit{cookie\_label}(d, f) \not\sqsubseteq l \Rightarrow (K_1(d) = ck \wedge ck(k) = (n, f) \Leftrightarrow K_2(d) = ck' \wedge ck'(k) = (n, f)).$$

Let $\mathit{url\_label}(u) \sqsubseteq l$*. If* $K_1' = \mathit{sec\_upd\_ck}(K_1, u, \hat{ck})$*, then:*

$$\mathit{cookie\_label}(d, f) \not\sqsubseteq l \Rightarrow (K_1'(d) = ck \wedge ck(k) = (n, f) \Leftrightarrow K_2(d) = ck' \wedge ck'(k) = (n, f)).$$

Proof. Let $u = (\pi, d, v)$. By using Lemma 9 we observe that, whenever $\{k \mapsto (n, f)\} \in (\hat{ck} \nearrow \pi)$, we have $\mathit{cookie\_label}(d, f) = \mathit{url\_label}(u)$. Since the secure update operation does not allow to overwrite cookies with a higher label, the conclusion follows by our assumptions on K_1 and K_2. □

Lemma 27 (Untainted Output)**.** *Let* $P = \langle W, K, N, \{p \mapsto (e, \hat{l})\}, [\,]\rangle$*. If* $l \models P$ *and* $P \xrightarrow{o} Q$ *with* $l \vdash \mathit{untainted}(o)$*, then* $W(p) = (u, h, h', \checkmark)$ *and* $\mathit{url\_label}(u) \not\sqsubseteq l$*.*

Proof. Since $l \vdash \mathit{untainted}(o)$ holds true, we observe that $P \xrightarrow{o} Q$ must be derived by (O-MIRROR), i.e., by the assumption $P \overset{o}{\mapsto} Q$. We then prove a similar statement where $P \xrightarrow{o} Q$ has been replaced by $P \overset{o}{\mapsto} Q$. The proof is by induction on the derivation of $P \overset{o}{\mapsto} Q$:

Case (O-LETCTX): in this case we know that:
1) $P = \langle W, K, N, \{p \mapsto (\mathsf{let}\ x = e'\ \mathsf{in}\ e, \hat{l})\}, [\,]\rangle$;
2) $Q = \langle W', K', N', \{p \mapsto (\mathsf{let}\ x = e''\ \mathsf{in}\ e, \hat{l})\}, [\,]\rangle$;
3) $\langle W, K, N, \{p \mapsto (e', \hat{l})\}, [\,]\rangle \overset{o}{\mapsto} \langle W', K', N', \{p \mapsto (e'', \hat{l})\}, [\,]\rangle$.

Since $\mathit{fn}(e') \subseteq \mathit{fn}(\mathsf{let}\ x = e'\ \mathsf{in}\ e)$, it is easy to show that $l \models P$ implies $l \models \langle W, K, N, \{p \mapsto (e', \hat{l})\}, [\,]\rangle$. The conclusion thus follows by induction hypothesis;

Case (O-XHR): in this case we know that:
1) $P = \langle W \uplus \{p \mapsto (u', h, h', q)\}, K, N, \{p \mapsto (\mathsf{xhr}(u, \lambda x.e), \hat{l})\}, [\,]\rangle$;
2) $o = \mathsf{xhr\_req}(ck, u)$;
3) $Q = \langle W \uplus \{p \mapsto (u', h, h' \uplus \{n \mapsto \lambda x.e\}, q)\}, K, N \uplus \{n \mapsto (u, p, q')\}, \{p \mapsto ((), l)\}, [\,]\rangle$;

4) $\hat{l} \neq \bot \Rightarrow \hat{l} = \mathit{url\_label}(u) = \mathit{url\_label}(u') \wedge q = \checkmark$;
5) $q = \checkmark \wedge \mathit{url\_label}(u) = \mathit{url\_label}(u') \Rightarrow ck = \mathit{get\_http\_ck}(K, u) \wedge q' = \checkmark$;
6) $q = \times \vee \mathit{url\_label}(u) \neq \mathit{url\_label}(u') \Rightarrow ck = \{\} \wedge q' = \times$.

Since $l \vdash \mathit{untainted}(o)$ holds true, we know that there exists $n \in \mathit{fn}(o) = \mathit{fn}(ck) \cup \mathit{fn}(u)$ such that $n \in \mathcal{N}_{l'}$ for some $l' \not\sqsubseteq l$. We distinguish two cases:

- if $n \in \mathit{fn}(ck)$, then we know that $q = \checkmark$ and $\mathit{url\_label}(u) = \mathit{url\_label}(u')$. Since $n \in \mathcal{N}_{l'}$ with $l' \not\sqsubseteq l$ and $n \in \mathit{fn}(\mathit{get\_http\_ck}(K, u))$, we know that $\mathit{url\_label}(u) \not\sqsubseteq l$ by Lemma 3. Hence, we have $\mathit{url\_label}(u') \not\sqsubseteq l$;
- if $n \in \mathit{fn}(u)$, we know that $\hat{l} \not\sqsubseteq l$ by the typing assumption (see condition 7 of Definition 7). Hence, by assumption 4 above, we must have $\hat{l} = \mathit{url\_label}(u') \not\sqsubseteq l$ and $q = \checkmark$;

Case (O-LOGIN): in this case we know that:

1) $P = \langle W, K, N, \{p \mapsto (\mathsf{auth}(u, c), \hat{l})\}, [\,]\rangle$;
2) $o = \mathsf{login}(ck, u, c)$ with $ck = \mathit{get\_http\_ck}(K, u)$;
3) $Q = \langle W, K, N \uplus \{n \mapsto (u, (), \checkmark)\}, \{p \mapsto ((), \hat{l})\}, [\,]\rangle$;
4) $W(p) = (u', h, h', \checkmark)$;
5) $\rho(c) = \mathit{url\_label}(u)$;
6) $\hat{l} = \mathit{url\_label}(u) = \mathit{url\_label}(u')$;

Since $l \vdash \mathit{untainted}(o)$ holds true, we know that there exists $n \in \mathit{fn}(o) = \mathit{fn}(ck) \cup \mathit{fn}(u) \cup \{c\}$ such that $n \in \mathcal{N}_{l'}$ for some $l' \not\sqsubseteq l$. We distinguish two cases:

- let $n \in \mathit{fn}(ck)$. Since $n \in \mathcal{N}_{l'}$ with $l' \not\sqsubseteq l$ and $n \in \mathit{fn}(\mathit{get\_http\_ck}(K, u))$, we know that $\mathit{url\_label}(u) \not\sqsubseteq l$ by Lemma 3. Hence, we have $\mathit{url\_label}(u') \not\sqsubseteq l$ by assumption 6 above;
- let $n \in \mathit{fn}(u) \cup \{c\}$, we know that $\hat{l} \not\sqsubseteq l$ by the typing assumption (see condition 7 of Definition 7). Hence, we conclude $\hat{l} = \mathit{url\_label}(u') \not\sqsubseteq l$ by assumption 6 above.

□

Lemma 28 (Window Update)**.** *Let $P = \langle W, K, N, T, O\rangle$ with $W(p) = (u, h, h', q)$. If $P \xrightarrow{o} Q$ and $Q = \langle W', K', N', T', O\rangle$, then $W'(p) = (u, h, h'', q)$ for some h''.*

Proof. If $P \xrightarrow{o} Q$ was proved by (O-COMPLETE), the conclusion is immediate. Otherwise, let $P \xrightarrow{o} Q$ be derived by (O-MIRROR), we prove a similar statement where $P \xrightarrow{o} Q$ has been replaced by $P \overset{o}{\mapsto} Q$. The proof is by induction on the derivation of $P \xrightarrow{o} Q$. Most of the cases are immediate, since the windows in W are not updated in W'. Case (O-LETCTX) follows by induction hypothesis, while case (O-XHR) follows by a simple inspection of the premises of the reduction rule. □

Lemma 29 (Browser Simulation)**.** *The following statements hold true:*

1) if $Q \sim_l Q'$, then $Q' \sim_l Q$;
2) if $C \sim_l C'$ and $C \xrightarrow{i} P$ with $l \vdash \mathit{untainted}(i)$, then $C' \xrightarrow{i} P'$ and $P \sim_l P'$;
3) if $C \sim_l C'$ and $C \xrightarrow{i} P$ with $l \vdash \mathit{tainted}(i)$, then $P \sim_l C'$;
4) if $P \sim_l C$ and $P \xrightarrow{o} Q$, then $l \vdash \mathit{tainted}(o)$ and $Q \sim_l C$;
5) if $P \sim_l P'$ and $P \xrightarrow{o} Q$ with $l \vdash \mathit{untainted}(o)$, then $P' \xrightarrow{o} Q'$ and $Q \sim_l Q'$.

Proof. We show the five points separately:

1) this follows directly by the definition of $Q \sim_l Q'$;
2) by a case analysis on the rule applied to prove $C \xrightarrow{i} P$. If the applied reduction rule is (I-COMPLETE), we perform a further case analysis on i:

   Case $i = \mathsf{load}(u)$: for any consumer state C there exists P such that $C \overset{i}{\mapsto} P$, hence rule (I-COMPLETE) can never be applied and the case is trivial;

   Case $i = \mathsf{text}(p, k, n)$: let $C = \langle W, K, N, \{\}, [\,]\rangle \xrightarrow{i} \langle W, K, N, \{\}, \bullet\rangle = P$ and let $C' = \langle W', K', N', \{\}, [\,]\rangle$. Let $C \overset{i}{\not\mapsto}$, we distinguish two cases:

   - if $C' \overset{i}{\not\mapsto}$, then we have:

$$\frac{\text{(I-COMPLETE)}}{C' = \langle W', K', N', \{\}, [\,]\rangle \xrightarrow{i} \langle W', K', N', \{\}, \bullet\rangle = P'}$$

with $P \sim_l P'$;

- otherwise, assume that there exists p such that $W'(p) = (u', h'_1, h'_2, q')$ and $h'_1(k) = \lambda x.e$ for some expression e. We must have either $url\_label(u') \sqsubseteq l$ or $q' = \mathsf{X}$ by the hypothesis $C \sim_l C'$. Thus, we have:

$$\frac{}{C' = \langle W', K', N', \{\}, [\,]\rangle \xrightarrow{i} \langle W', K', N', \{p \mapsto (e\{n/x\}, \rho(n)\}, [\,]\rangle = P''} \text{(I-TEXT)}$$

with $P \sim_l P''$, since we know that either $url\_label(u') \sqsubseteq l$ or $q' = \mathsf{X}$ and thus the new running expression $e\{n/x\}$ cannot break the invariant;

Case $i = \mathsf{doc\_resp}(n, ck, u, \mathsf{blank}, h, e)$: let $C = \langle W, K, N, \{\}, [\,]\rangle \xrightarrow{i} \langle W, K, N, \{\}, \bullet\rangle = P$ and assume that $C' = \langle W', K', N', \{\}, [\,]\rangle$. Since we know that $C \not\xrightarrow{i}$, we know that there does not exist $n \in \mathit{dom}(N)$ such that $N(n) = (u, (), q)$. We then distinguish two cases:

- assume that $N' \neq N'' \uplus \{n \mapsto (u, (), q')\}$, then we have:

$$\frac{}{C' = \langle W', K', N', \{\}, [\,]\rangle \xrightarrow{i} \langle W', K', N', \{\}, \bullet\rangle = P'} \text{(I-COMPLETE)}$$

with $P \sim_l P'$;

- otherwise, assume that $N' = N'' \uplus \{n \mapsto (u, (), q')\}$. Since we are assuming that $l \vdash untainted(i)$ holds true, we know that $url\_label(u) \not\sqsubseteq l$, hence we have $q' = \mathsf{X}$ by the hypothesis $C \sim_l C'$. Thus, we have:

$$\frac{}{C' = \langle W', K', N', \{\}, [\,]\rangle \xrightarrow{i} \langle W' \uplus \{p \mapsto (u, h, \{\}, \mathsf{X})\}, K', N'', \{p \mapsto (e, \bot)\}, [\,]\rangle = P''} \text{(I-DOCRESP)}$$

with $P \sim_l P''$, since the new page p is tainted, the removed network connection n is tainted and the expression e runs in a tainted page;

Case $i = \mathsf{doc\_resp}(n, ck, u, u', h, e)$ with $u' \neq \mathsf{blank}$: let $C = \langle W, K, N, \{\}, [\,]\rangle \xrightarrow{i} \langle W, K, N, \{\}, \bullet\rangle = P$ and assume that $C' = \langle W', K', N', \{\}, [\,]\rangle$. Since we know that $C \not\xrightarrow{i}$, we know that there does not exist $n \in \mathit{dom}(N)$ such that $N(n) = (u, (), q)$. We then distinguish two cases:

- assume that $N' \neq N'' \uplus \{n \mapsto (u, (), q')\}$, then we have:

$$\frac{}{C' = \langle W', K', N', \{\}, [\,]\rangle \xrightarrow{i} \langle W', K', N', \{\}, \bullet\rangle = P'} \text{(I-COMPLETE)}$$

with $P \sim_l P'$;

- otherwise, assume that $N' = N'' \uplus \{n \mapsto (u, (), q')\}$. Since we are assuming that $l \vdash untainted(i)$ holds true, we know that $url\_label(u) \not\sqsubseteq l$, hence we have $q' = \mathsf{X}$ by the hypothesis $C \sim_l C'$. Thus, we have:

$$\frac{}{C' = \langle W', K', N', \{\}, [\,]\rangle \xrightarrow{i} \langle W', K', N' \uplus \{m \mapsto (u, (), \mathsf{X})\}, \{\}, [\,]\rangle = P''} \text{(I-DOCREDIR)}$$

with $P \sim_l P''$, since we just updated a tainted network connection;

Case $i = \mathsf{xhr\_resp}(n, ck, u, \mathsf{blank}, v)$: let $C = \langle W, K, N, \{\}, [\,]\rangle \xrightarrow{i} \langle W, K, N, \{\}, \bullet\rangle = P$ and assume that $C' = \langle W', K', N', \{\}, [\,]\rangle$. Let $C \not\xrightarrow{i}$, we distinguish two cases:

- if $C' \not\xrightarrow{i}$, then we have:

$$\frac{}{C' = \langle W', K', N', \{\}, [\,]\rangle \xrightarrow{i} \langle W', K', N', \{\}, \bullet\rangle = P'} \text{(I-COMPLETE)}$$

with $P \sim_l P'$;

- otherwise, assume that there exists P' such that $C' \overset{i}{\mapsto} P'$:

(I-XHRRESP)
$$\frac{W' = \hat{W} \uplus \{p \mapsto (u', h, h' \uplus \{n \mapsto \lambda x.e\}, q)\} \qquad N' = N'' \uplus \{n \mapsto (u, p, q)\} \\ W'' = \hat{W} \uplus \{p \mapsto (u', h, h', q)\} \qquad q = \checkmark \Rightarrow K'' = \mathit{sec\_upd\_ck}(K', u, ck) \\ q = \times \Rightarrow K'' = K' \qquad l' = \mathit{url\_label}(u')}{C' = \langle W', K', N', \{\}, [\,]\rangle \xrightarrow{i} \langle W'', K'', N'', \{p \mapsto (e\{v/x\}, l'\}, [\,]\rangle = P''}$$

Since we are assuming that $l \vdash \mathit{untainted}(i)$ holds true, we know that $\mathit{url\_label}(u) \not\sqsubseteq l$, hence we must have $q = \times$ by the hypothesis $C \sim_l C'$. Thus, $K'' = K'$ and we trivially preserve the invariant on the cookie jar. Since $q = \times$, we are just removing a tainted network connection n and we preserve the invariant on the connection store; similarly, the change in the page p and the introduction of the new expression $e\{v/x\}$ cannot break the invariant, since p is tainted, so we conclude $P \sim_l P''$;

Case $i = \mathsf{xhr\_resp}(n, ck, u, u', v)$ with $u' \neq \mathsf{blank}$: similar to the case $i = \mathsf{doc\_resp}(n, ck, u, u', h, e)$.

If the applied reduction rule is (I-MIRROR), we prove a similar statement where $\xrightarrow{i}$ has been replaced by $\overset{i}{\mapsto}$.

The proof is by a case analysis on the rule applied to prove $C \overset{i}{\mapsto} P$:

Case (I-LOAD): we have:

1. $C = \langle W, K, N, \{\}, [\,]\rangle$;
2. $i = \mathsf{load}(u)$ with $\vdash_\diamond i$;
3. $P = \langle W, K, N \uplus \{n \mapsto (u, (), \checkmark)\}, \{\}, \mathsf{doc\_req}(ck, u)\rangle$ with $ck = \mathit{get\_http\_ck}(K, u)$.

Let $C' = \langle W', K', N', \{\}, [\,]\rangle$, we have:

(I-LOAD)
$$\frac{ck' = \mathit{get\_http\_ck}(K', u)}{C' = \langle W', K', N', \{\}, [\,]\rangle \xrightarrow{\mathsf{load}(u)} \langle W', K', N' \uplus \{n \mapsto (u, (), \checkmark)\}, \{\}, \mathsf{doc\_req}(ck', u)\rangle = P'}$$

To show $P \sim_l P'$ we only need to focus on the output buffers. Specifically, we have to show that the two new document requests are both tainted or both untainted and, whenever they are untainted, they must coincide. We perform a case distinction:

- let $\mathit{url\_label}(u) \sqsubseteq l$. Since we know that $\vdash_\diamond \mathsf{load}(u)$ holds true, we must have that $\forall n \in \mathit{fn}(u) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$ by Lemma 14. By Lemma 3 we know that $\forall n \in \mathit{fn}(ck) \cup \mathit{fn}(ck') : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$. Hence, we have $l \vdash \mathit{tainted}(\mathsf{doc\_req}(ck, u))$ and $l \vdash \mathit{tainted}(\mathsf{doc\_req}(ck', u))$ and we are done;
- let $\mathit{url\_label}(u) \not\sqsubseteq l$. By Lemma 24 we have $ck = ck'$, which is enough to conclude;

Case (I-TEXT): we have:

1. $C = \langle W, K, N, \{\}, [\,]\rangle$;
2. $i = \mathsf{text}(p, k, n)$ with $\vdash_\diamond i$;
3. $P = \langle W, K, N, \{p \mapsto (e\{n/x\}, \rho(n))\}, [\,]\rangle$;
4. $W(p) = (u, h_1, h_2, q)$ and $h_1(k) = \lambda x.e$.

Let $C' = \langle W', K', N', \{\}, [\,]\rangle$. If $\mathit{url\_label}(u) \sqsubseteq l$ or $q = \times$, then the text input event may succeed or not in C' by our hypothesis $C \sim_l C'$. If it succeeds, then $W'(p) = (u', h'_1, h'_2, q')$ with $h'_1(k) = \lambda x.e'$ for some expression e'. Moreover, we know that either $\mathit{url\_label}(u') \sqsubseteq l$ or $q' = \times$ again by the hypothesis $C \sim_l C'$. Thus, we have:

(I-TEXT)
$$\frac{W'(p) = (u', h'_1, h'_2, q') \qquad h'_1(k) = \lambda x.e'}{C' = \langle W', K', N', \{\}, [\,]\rangle \xrightarrow{i} \langle W', K', N', \{p \mapsto (e'\{n/x\}, \rho(n))\}, [\,]\rangle = P'}$$

with $P \sim_l P'$, since we know that either $\mathit{url\_label}(u) \sqsubseteq l$ or $q = \times$, and either $\mathit{url\_label}(u') \sqsubseteq l$ or $q' = \times$, thus the new running expressions $e\{n/x\}$ and $e'\{n/x\}$ cannot break the invariant. If instead the text input fails, we have:

(I-COMPLETE)
$$\frac{C' \overset{i}{\not\mapsto}}{C' = \langle W', K', N', \{\}, [\,]\rangle \xrightarrow{i} \langle W', K', N', \{\}, \bullet\rangle = P''}$$

with $P \sim_l P''$, since again we know that $url\_label(u) \sqsubseteq l$ or $q = \times$. Finally, let $url\_label(u) \not\sqsubseteq l$ and $q = \checkmark$, then the text input event must succeed in C' by our hypothesis $C \sim_l C'$. Specifically, we have:

(I-TEXT)
$$\frac{W'(p) = (u, h_1, h_2, \checkmark) \qquad h_1(k) = \lambda x.e}{C' = \langle W', K', N', \{\}, [\,]\rangle \xrightarrow{i} \langle W', K', N', \{p \mapsto (e\{n/x\}, \rho(n))\}, [\,]\rangle = \hat{P}}$$

with $P \sim_l \hat{P}$, since the new running expressions are the same;

Case (I-DOCRESP): we have:

1. $C = \langle W, K, N \uplus \{n \mapsto (u, (), q)\}, \{\}, [\,]\rangle$;
2. $i = \mathsf{doc\_resp}(n, ck, u, \mathsf{blank}, h, e)$ with $l \vdash untainted(i)$;
3. $P = \langle W \uplus \{p \mapsto (u, h, \{\}, q)\}, \hat{K}, N, \{p \mapsto (e, \bot)\}, [\,]\rangle$;
4. $q = \checkmark \Rightarrow \hat{K} = \mathit{sec\_upd\_ck}(K, u, ck)$;
5. $q = \times \Rightarrow \hat{K} = K$.

Let $C' = \langle W', K', N', \{\}, [\,]\rangle$. Since $l \vdash untainted(i)$ holds true, we know that $url\_label(u) \not\sqsubseteq l$. If $q = \times$, then $P = \langle W \uplus \{p \mapsto (u, h, \{\}, q)\}, K, N, \{p \mapsto (e, \bot)\}, [\,]\rangle$ and the document response event may succeed or not in C' by our hypothesis $C \sim_l C'$. If it succeeds, then $N' = N'' \uplus \{n \mapsto (u, (), \times)\}$ and we have:

(I-DOCRESP)
$$C' = \langle W', K', N'' \uplus \{n \mapsto (u, (), \times)\}, \{\}, [\,]\rangle \overset{i}{\mapsto} \langle W' \uplus \{p \mapsto (u, h, \{\}, \times)\}, K', N'', \{p \mapsto (e, \bot)\}, [\,]\rangle = P'$$

with $P \sim_l P'$. If instead the document response fails, we have:

(I-COMPLETE)
$$\frac{C' \overset{i}{\not\mapsto}}{C' = \langle W', K', N', \{\}, [\,]\rangle \xrightarrow{i} \langle W', K', N', \{\}, \bullet\rangle = P''}$$

with $P \sim_l P''$, since we know that $q = \times$. Finally, let $q = \checkmark$, then we have:

$$P = \langle W \uplus \{p \mapsto (u, h, \{\}, q)\}, \hat{K}, N, \{p \mapsto (e, \bot)\}, [\,]\rangle,$$

with $\hat{K} = \mathit{sec\_upd\_ck}(K, u, ck)$ and the document response event must succeed in C' by our hypothesis $C \sim_l C'$. In particular, we have $N' = N'' \uplus \{n \mapsto (u, (), \checkmark)\}$ and we can prove:

(I-DOCRESP)
$$\frac{K'' = \mathit{sec\_upd\_ck}(K', u, ck)}{C' = \langle W', K', N'' \uplus \{n \mapsto (u, (), \checkmark)\}, \{\}, [\,]\rangle \overset{i}{\mapsto} \langle W' \uplus \{p \mapsto (u, h, \{\}, \checkmark)\}, K'', N'', \{p \mapsto (e, \bot)\}, [\,]\rangle = \hat{P}}$$

To show that $P \sim_l \hat{P}$ we only need to focus on the cookie jars. Specifically, we want to show that for $\hat{K} = \mathit{sec\_upd\_ck}(K, u, ck)$ and $K'' = \mathit{sec\_upd\_ck}(K', u, ck)$ we have:

$$\mathit{cookie\_label}(d, f) \not\sqsubseteq l \Rightarrow (\hat{K}(d) = ck \wedge ck(k) = (n, f) \Leftrightarrow K''(d) = ck' \wedge ck'(k) = (n, f)),$$

which follows by Lemma 25;

Case (I-DOCREDIR): we have:

1. $C = \langle W, K, N \uplus \{n \mapsto (u, (), q)\}, \{\}, [\,]\rangle$;
2. $i = \mathsf{doc\_resp}(n, ck, u, u', h, e)$ with $l \vdash untainted(i)$;
3. $P = \langle W, \hat{K}, N \uplus \{n \mapsto (u', (), q')\}, \{\}, \mathsf{doc\_req}(ck'', u')\rangle$;
4. $q = \checkmark \Rightarrow \hat{K} = \mathit{sec\_upd\_ck}(K, u, ck)$;
5. $q = \times \Rightarrow \hat{K} = K$;
6. $q = \checkmark \wedge url\_label(u) = url\_label(u') \Rightarrow ck'' = \mathit{get\_http\_ck}(\hat{K}, u') \wedge q' = \checkmark$;
7. $q = \times \vee url\_label(u) \neq url\_label(u') \Rightarrow ck'' = \{\} \wedge q' = \times$.

Let $C' = \langle W', K', N', \{\}, [] \rangle$. Since $l \vdash \mathit{untainted}(i)$ holds true, we know that $\mathit{url\_label}(u) \not\sqsubseteq l$ and $\vdash_\diamond i$. If $q = \times$, then $P = \langle W, K, N \uplus \{n \mapsto (u', (), \times)\}, \{\}, \mathsf{doc\_req}(\{\}, u') \rangle$ and the redirect may succeed or not in C' by our hypothesis $C \sim_l C'$. If it succeeds, then $N' = N'' \uplus \{n \mapsto (u, (), \times)\}$ and we have:

(I-DocRedir)

$$C' = \langle W', K', N'' \uplus \{n \mapsto (u, (), \times)\}, \{\}, [] \rangle \xrightarrow{i} \langle W', K', N'' \uplus \{n \mapsto (u', (), \times)\}, \{\}, \mathsf{doc\_req}(\{\}, u') \rangle = P'$$

with $P \sim_l P'$. If instead the redirect fails, we have:

(I-Complete)

$$\frac{C' \not\xrightarrow{i}}{C' = \langle W', K', N', \{\}, [] \rangle \xrightarrow{i} \langle W', K', N', \{\}, \bullet \rangle = P''}$$

with $P \sim_l P''$, since $\vdash_\diamond i$ implies $\mathit{fn}(u') \subseteq \mathcal{N}_\perp$ and then $l \vdash \mathit{tainted}(\mathsf{doc\_req}(\{\}, u'))$ holds true. Finally, let $q = \checkmark$, then the redirect must succeed in C' by our hypothesis $C \sim_l C'$. In particular, we have $N' = N'' \uplus \{n \mapsto (u, (), \checkmark)\}$ and:

(I-DocRedir)

$$\mathit{url\_label}(u) = \mathit{url\_label}(u') \Rightarrow \hat{ck} = \mathit{get\_http\_ck}(K'', u') \wedge q'' = \checkmark$$
$$\mathit{url\_label}(u) \neq \mathit{url\_label}(u') \Rightarrow \hat{ck} = \{\} \wedge q'' = \times$$

$$C' = \langle W', K', N'' \uplus \{n \mapsto (u, (), \checkmark)\}, \{\}, [] \rangle \xrightarrow{i} \langle W', K'', N'' \uplus \{n \mapsto (u', (), q'')\}, \{\}, \mathsf{doc\_req}(\hat{ck}, u') \rangle = \hat{P}$$

We distinguish two cases. If $\mathit{url\_label}(u) = \mathit{url\_label}(u')$, then we know that $\mathit{url\_label}(u') \not\sqsubseteq l$ and:

$$P = \langle W, \hat{K}, N \uplus \{n \mapsto (u', (), \checkmark)\}, \{\}, \mathsf{doc\_req}(\mathit{get\_http\_ck}(\hat{K}, u'), u') \rangle$$
$$\hat{P} = \langle W', K'', N'' \uplus \{n \mapsto (u', (), \checkmark)\}, \{\}, \mathsf{doc\_req}(\mathit{get\_http\_ck}(K'', u'), u') \rangle$$

with $\hat{K} = \mathit{sec\_upd\_ck}(K, u, ck)$ and $K'' = \mathit{sec\_upd\_ck}(K', u, ck)$. The desired conclusion $P \sim_l \hat{P}$ follows by showing the invariant on the updated cookie jars (by Lemma 25) and the new document requests (by Lemma 24). If instead we have $\mathit{url\_label}(u) \neq \mathit{url\_label}(u')$, then:

$$P = \langle W, K, N \uplus \{n \mapsto (u', (), \checkmark)\}, \{\}, \mathsf{doc\_req}(\{\}, u') \rangle$$
$$\hat{P} = \langle W', K', N'' \uplus \{n \mapsto (u', (), \checkmark)\}, \{\}, \mathsf{doc\_req}(\{\}, u') \rangle$$

and the desired conclusion $P \sim_l \hat{P}$ easily follows;

Case (I-XhrResp): we have:

1. $C = \langle W \uplus \{p \mapsto (u_1, h_1, h_2 \uplus \{n \mapsto \lambda x.e\}, q)\}, K, N \uplus \{n \mapsto (u, p, q)\}, \{\}, [] \rangle$;
2. $i = \mathsf{xhr\_resp}(n, ck, u, \mathsf{blank}, v)$;
3. $P = \langle W \uplus \{p \mapsto (u_1, h_1, h_2, q)\}, \hat{K}, N, \{p \mapsto (e\{v/x\}, \hat{l})\}, [] \rangle$;
4. $q = \checkmark \Rightarrow \hat{K} = \mathit{sec\_upd\_ck}(K, u, ck)$;
5. $q = \times \Rightarrow \hat{K} = K$;
6. $\hat{l} = \mathit{url\_label}(u_1)$.

Let $C' = \langle W', K', N', \{\}, [] \rangle$. Since $l \vdash \mathit{untainted}(i)$ holds true, we know that $\mathit{url\_label}(u) \not\sqsubseteq l$ and $\vdash_\diamond i$. If $q = \times$, then $P = \langle W \uplus \{p \mapsto (u_1, h_1, h_2, \times)\}, K, N, \{p \mapsto (e\{v/x\}, \hat{l})\}, [] \rangle$ and the XHR response may succeed or not in C' by our hypothesis $C \sim_l C'$. If it succeeds, then $N' = N'' \uplus \{n \mapsto (u, p', \times)\}$ and $W' = \hat{W} \uplus \{p' \mapsto (u_1', h_1', h_2' \uplus \{n \mapsto \lambda x.e'\}, \times)\}$, thus we have:

(I-XhrResp)

$$\frac{W'' = \hat{W} \uplus \{p' \mapsto (u_1', h_1', h_2', \times)\} \qquad \hat{l}' = \mathit{url\_label}(u_1')}{C' = \langle W', K', N'' \uplus \{n \mapsto (u, p', \times)\}, \{\}, [] \rangle \xrightarrow{i} \langle W'', K', N'', \{p \mapsto (e'\{v/x\}, \hat{l}')\}, [] \rangle = P'}$$

with $P \sim_l P'$, since the two expressions $e\{v/x\}$ and $e'\{v/x\}$ are executed on tainted pages. If instead the XHR response fails, we have:

(I-Complete)

$$\frac{C' \not\xrightarrow{i}}{C' = \langle W', K', N', \{\}, [] \rangle \xrightarrow{i} \langle W', K', N', \{\}, \bullet \rangle = P''}$$

with $P \sim_l P''$, since we know that the expression $e\{v/x\}$ is executed on a tainted page. Finally, let $q = \checkmark$, then we have $P = \langle W \uplus \{p \mapsto (u_1, h_1, h_2, \checkmark)\}, \hat{K}, N, \{p \mapsto (e\{v/x\}, \hat{l})\}, [\,]\rangle$ with $\hat{K} = \textit{sec\_upd\_ck}(K, u, ck)$ and the XHR response must succeed in C' by our hypothesis $C \sim_l C'$. In particular, we have $N' = N'' \uplus \{n \mapsto (u, p, \checkmark)\}$ and $W' = W'' \uplus \{p \mapsto (u_1, h_1, h_2 \uplus \{n \mapsto \lambda x.e\}, \checkmark)\}$, thus we have:

(I-XHRRESP)
$$\frac{\hat{W} = W'' \uplus \{p \mapsto (u_1, h_1, h_2, \checkmark)\} \qquad K'' = \textit{sec\_upd\_ck}(K', u, ck)}{C' = \langle W', K', N'' \uplus \{n \mapsto (u, p, \checkmark)\}, \{\}, [\,]\rangle \overset{i}{\mapsto} \langle \hat{W}, K'', N'', \{p \mapsto (e\{v/x\}, \hat{l})\}, [\,]\rangle = \hat{P}}$$

The desired conclusion $P \sim_l \hat{P}$ follows by Lemma 25, which ensures that the invariant is preserved on the updated cookie jars;

Case (I-XHRREDIR): analogous to case (I-DOCREDIR);

3) by a case analysis on the rule applied to prove $C \xrightarrow{i} P$. If the applied rule is (I-COMPLETE), let $C = \langle W, K, N, \{\}, [\,]\rangle$ and $C' = \langle W', K', N', \{\}, [\,]\rangle$. We have $C \xrightarrow{i} \langle W, K, N, \{\}, \bullet\rangle = P$ and $P \sim_l C'$, since $l \vdash \textit{tainted}(\bullet)$ holds true. If the applied reduction rule is (I-MIRROR), we prove a similar statement where $\xrightarrow{i}$ has been replaced by $\overset{i}{\mapsto}$. The proof is by a case analysis on the rule applied to prove $C \overset{i}{\mapsto} P$:

Case (I-DOCRESP): we have:

1. $C = \langle W, K, N \uplus \{n \mapsto (u, (), q)\}, \{\}, [\,]\rangle$;
2. $i = \mathsf{doc\_resp}(n, ck, u, \mathsf{blank}, h, e)$ with $l \vdash \textit{tainted}(i)$;
3. $P = \langle W \uplus \{p \mapsto (u, h, \{\}, q)\}, \hat{K}, N, \{p \mapsto (e, \bot)\}, [\,]\rangle$;
4. $q = \checkmark \Rightarrow \hat{K} = \textit{sec\_upd\_ck}(K, u, ck)$;
5. $q = \times \Rightarrow \hat{K} = K$.

Let $C' = \langle W', K', N', \{\}, [\,]\rangle$ with $C \sim_l C'$, we want to show that $P \sim_l C'$. Since $l \vdash \textit{tainted}(i)$ holds true, we know that $\textit{url\_label}(u) \sqsubseteq l$, hence the addition of the new page p, the removal of the network connection n and the presence of the new running expression e cannot break the invariant. The most interesting point to show is related to the cookie jars, i.e., we have to show:

$$\textit{cookie\_label}(d, f) \not\sqsubseteq l \Rightarrow (\hat{K}(d) = ck \wedge ck(k) = (n, f) \Leftrightarrow K'(d) = ck' \wedge ck'(k) = (n, f))$$

If $q = \times$, then $\hat{K} = K$ and we conclude by the hypothesis $C \sim_l C'$; if $q = \checkmark$, then $\hat{K} = \textit{sec\_upd\_ck}(K, u, ck)$, hence the desired property follows by Lemma 26;

Case (I-DOCREDIR): we have:

1. $C = \langle W, K, N \uplus \{n \mapsto (u, (), q)\}, \{\}, [\,]\rangle$;
2. $i = \mathsf{doc\_resp}(n, ck, u, u', h, e)$ with $l \vdash \textit{tainted}(i)$;
3. $P = \langle W, \hat{K}, N \uplus \{n \mapsto (u', (), q')\}, \{\}, \mathsf{doc\_req}(ck'', u')\rangle$;
4. $q = \checkmark \Rightarrow \hat{K} = \textit{sec\_upd\_ck}(K, u, ck)$;
5. $q = \times \Rightarrow \hat{K} = K$;
6. $q = \checkmark \wedge \textit{url\_label}(u) = \textit{url\_label}(u') \Rightarrow ck'' = \textit{get\_http\_ck}(\hat{K}, u') \wedge q' = \checkmark$;
7. $q = \times \vee \textit{url\_label}(u) \neq \textit{url\_label}(u') \Rightarrow ck'' = \{\} \wedge q' = \times$.

Let $C' = \langle W', K', N', \{\}, [\,]\rangle$ with $C \sim_l C'$, we want to show that $P \sim_l C'$. Since $l \vdash \textit{tainted}(i)$ holds true, we know that $\textit{url\_label}(u) \sqsubseteq l$ and $\forall n \in \textit{fn}(i) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$. We distinguish two cases:

- if $q = \times$ or $\textit{url\_label}(u) \neq \textit{url\_label}(u')$, then $P = \langle W, K, N \uplus \{n \mapsto (u', (), \times)\}, \{\}, \mathsf{doc\_req}(\{\}, u')\rangle$, hence to conclude we just need to show that $l \vdash \textit{tainted}(\mathsf{doc\_req}(\{\}, u'))$ holds true. Since we know that $\forall n \in \textit{fn}(u') \subseteq \textit{fn}(i) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$, the conclusion follows;
- if $q = \checkmark$ and $\textit{url\_label}(u) = \textit{url\_label}(u')$, then:

$$P = \langle W, \hat{K}, N \uplus \{n \mapsto (u', (), \checkmark)\}, \{\}, \mathsf{doc\_req}(\textit{get\_http\_ck}(\hat{K}, u'), u')\rangle,$$

with $\hat{K} = \textit{sec\_upd\_ck}(K, u, ck)$. Since $\textit{url\_label}(u') = \textit{url\_label}(u) \sqsubseteq l$, the new network connection cannot break the invariant. As to the cookie store, we want to show that:

$$\textit{cookie\_label}(d, f) \not\sqsubseteq l \Rightarrow (\hat{K}(d) = ck \wedge ck(k) = (n, f) \Leftrightarrow K'(d) = ck' \wedge ck'(k) = (n, f)),$$

which follows by Lemma 26. Last, we need to show that $l \vdash \mathit{tainted}(\mathsf{doc\_req}(\mathit{get\_http\_ck}(\hat{K}, u'), u'))$ holds true. By Lemma 3 we know that $\forall n \in \mathit{fn}(\mathit{get\_http\_ck}(\hat{K}, u')) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$. We already know that $\forall n \in \mathit{fn}(u') \subseteq \mathit{fn}(i) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$ by the hypothesis $l \vdash \mathit{tainted}(i)$, hence the conclusion follows;

Case (I-XHRRESP): we have:

1. $C = \langle W \uplus \{p \mapsto (u', h_1, h_2 \uplus \{n \mapsto \lambda x.e\}, q)\}, K, N \uplus \{n \mapsto (u, p, q)\}, \{\}, [\,]\rangle$;
2. $i = \mathsf{xhr\_resp}(n, ck, u, \mathsf{blank}, v)$;
3. $P = \langle W \uplus \{p \mapsto (u_1, h_1, h_2, q)\}, \hat{K}, N, \{p \mapsto (e\{v/x\}, \hat{l})\}, [\,]\rangle$;
4. $q = \checkmark \Rightarrow \hat{K} = \mathit{sec\_upd\_ck}(K, u, ck)$;
5. $q = \times \Rightarrow \hat{K} = K$;
6. $\hat{l} = \mathit{url\_label}(u')$.

Let $C' = \langle W', K', N', \{\}, [\,]\rangle$. Since $l \vdash \mathit{tainted}(i)$ holds true, we know that $\mathit{url\_label}(u) \sqsubseteq l$ and $\forall n \in \mathit{fn}(i) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$. We distinguish two cases:

- if $q = \checkmark$, then $\mathit{url\_label}(u') = \mathit{url\_label}(u)$ by the typing assumption. Hence, $\mathit{url\_label}(u') \sqsubseteq l$ and the change in the page p and the introduction of the new expression $e\{v/x\}$ cannot break the invariant. To show the invariant on the new cookie jar $\hat{K}$ we appeal to Lemma 26;
- if $q = \times$, then the change in the page p and the introduction of the new expression $e\{v/x\}$ cannot break the invariant. Since $\hat{K} = K$, we can conclude;

Case (I-XHRREDIR): analogous to case (I-DOCREDIR);

4) first we prove that $P \sim_l C$ and $P \xrightarrow{o} Q$ implies $l \vdash \mathit{tainted}(o)$. If $P \xrightarrow{o} Q$ was derived by (O-COMPLETE), then $o = \bullet$ and $l \vdash \mathit{tainted}(\bullet)$ holds true. Otherwise, let $P \xrightarrow{o} Q$ be derived by (O-MIRROR), i.e., by the assumption $P \overset{o}{\mapsto} Q$. If $P \overset{o}{\mapsto} Q$ was derived by (O-FLUSH), we know that $P = \langle W, K, N, T, o\rangle$: since we assume $P \sim_l C = \langle W', K', N', \{\}, [\,]\rangle$, we must have $l \vdash \mathit{tainted}(o)$. Otherwise, assume that $P \overset{o}{\mapsto} Q$ was derived by any other rule, then we know that $P = \langle W, K, N, \{p \mapsto (e, \hat{l})\}, [\,]\rangle$. Assume by contradiction that $l \vdash \mathit{untainted}(o)$ holds true, then by Lemma 27 we know that $W(p) = (u, h, h', \checkmark)$ with $\mathit{url\_label}(u) \not\sqsubseteq l$. Since we assume $P = \langle W, K, N, \{p \mapsto (e, \hat{l})\}, [\,]\rangle \sim_l C = \langle W', K', N', \{\}, [\,]\rangle$, we get a contradiction, since the running expression e should occur also in C;
Now let $P \sim_l C$ and $P \xrightarrow{o} Q$, we show that $Q \sim_l C$. If $P \xrightarrow{o} Q$ was derived by (O-COMPLETE), we have:

$$P = \langle W, K, N, \{p \mapsto (e, \hat{l})\}, [\,]\rangle \xrightarrow{\bullet} \langle W, K, N, \{\}, [\,]\rangle = Q,$$

with $Q \sim_l C$. Indeed, since $P = \langle W, K, N, \{p \mapsto (e, \hat{l})\}, [\,]\rangle \sim_l C = \langle W', K', N', \{\}, [\,]\rangle$, we must have $W(p) = (u, h, h', q)$ with either $q = \times$ or $\mathit{url\_label}(u) \sqsubseteq l$, hence discarding the stuck expression e cannot break the invariant. If instead $P \xrightarrow{o} Q$ was derived by (O-MIRROR), we need to prove a similar statement where $P \xrightarrow{o} Q$ has been replaced by $P \overset{o}{\mapsto} Q$. The proof is by induction on the derivation of $P \overset{o}{\mapsto} Q$. Again, notice that, whenever $P = \langle W, K, N, \{p \mapsto (e, \hat{l})\}, O\rangle \sim_l C = \langle W', K', N', \{\}, [\,]\rangle$, we must have $W(p) = (u, h, h', q)$ with either $q = \times$ or $\mathit{url\_label}(u) \sqsubseteq l$. Hence, to show that the invariant is preserved we can disregard the structure of the running expression e as long as it does not change the internal data structures of the browser W, K, N, O. We focus on the remaining cases:

Case (O-LETCTX): we have:

1. $P = \langle W, K, N, \{p \mapsto (\mathsf{let}\ x = e'\ \mathsf{in}\ e, \hat{l})\}, [\,]\rangle$;
2. $Q = \langle W', K, N', \{p \mapsto (\mathsf{let}\ x = e''\ \mathsf{in}\ e, \hat{l})\}, [\,]\rangle$;
3. $P' = \langle W, K, N, \{p \mapsto (e', \hat{l})\}, [\,]\rangle \overset{o}{\mapsto} \langle W', K, N', \{p \mapsto (e'', \hat{l})\}, [\,]\rangle = Q'$.

Since $W(p) = (u, h, h', q)$ with either $q = \times$ or $\mathit{url\_label}(u) \sqsubseteq l$, we have $P' \sim_l C$. By induction hypothesis we get $Q' \sim_l C$, hence the conclusion $Q \sim_l C$ follows by the observation that $W'(p) = (u, h, h'', q)$ by Lemma 28, and we know that either $q = \times$ or $\mathit{url\_label}(u) \sqsubseteq l$;

Case (O-SET): the update of the cookie jar can only introduce a new cookie with a cookie label equal to $\bot$. Notice also that no cookie with a cookie label greater than $\bot$ can be overwritten in this step, hence the invariant must be preserved;

Case (O-XHR): we have:

1. $P = \langle W \uplus \{p \mapsto (u', h, h', q)\}, K, N, \{p \mapsto (\mathsf{xhr}(u, \lambda x.e), \hat{l})\}, [\,]\rangle$;

2. $o = \mathsf{xhr\_req}(ck, u)$;
3. $Q = \langle W \uplus \{p \mapsto (u', h, h' \uplus \{n \mapsto \lambda x.e\}, q)\}, K, N \uplus \{n \mapsto (u, p, q')\}, \{p \mapsto ((), l)\}, [\,]\rangle$;
4. $\hat{l} \neq \bot \Rightarrow \hat{l} = \mathit{url\_label}(u) = \mathit{url\_label}(u') \wedge q = \checkmark$;
5. $q = \checkmark \wedge \mathit{url\_label}(u) = \mathit{url\_label}(u') \Rightarrow ck = \mathit{get\_http\_ck}(K, u) \wedge q' = \checkmark$;
6. $q = \times \vee \mathit{url\_label}(u) \neq \mathit{url\_label}(u') \Rightarrow ck = \{\} \wedge q' = \times$.

We know that either $q = \times$ or $\mathit{url\_label}(u') \sqsubseteq l$, so we distinguish two cases:

- if $q = \times$, we have:

$$Q = \langle W \uplus \{p \mapsto (u', h, h' \uplus \{n \mapsto \lambda x.e\}, \times)\}, K, N \uplus \{n \mapsto (u, p, \times)\}, \{p \mapsto ((), l)\}, [\,]\rangle,$$

  hence the changes in the page p and in the network connection store cannot break the invariant and we have $Q \sim_l C$;
- if $\mathit{url\_label}(u') \sqsubseteq l$, without loss of generality assume that $q = \checkmark$ (otherwise we conclude as in the previous case). If $\mathit{url\_label}(u) \neq \mathit{url\_label}(u')$, we have:

$$Q = \langle W \uplus \{p \mapsto (u', h, h' \uplus \{n \mapsto \lambda x.e\}, \checkmark)\}, K, N \uplus \{n \mapsto (u, p, \times)\}, \{p \mapsto ((), l)\}, [\,]\rangle,$$

  hence the changes in the page p and in the network connection store cannot break the invariant and we have $Q \sim_l C$. Otherwise, assume $\mathit{url\_label}(u) = \mathit{url\_label}(u')$, then we have:

$$Q = \langle W \uplus \{p \mapsto (u', h, h' \uplus \{n \mapsto \lambda x.e\}, \checkmark)\}, K, N \uplus \{n \mapsto (u, p, \checkmark)\}, \{p \mapsto ((), l)\}, [\,]\rangle$$

  hence the only change which could potentially break our invariant is the introduction of the new network connection $\{n \mapsto (u, p, \checkmark)\}$. However, we know that $\mathit{url\_label}(u) = \mathit{url\_label}(u') \sqsubseteq l$, hence we can conclude $Q \sim_l C$;

Case (O-LOGIN): in this case we know that:

1. $P = \langle W, K, N, \{p \mapsto (\mathsf{auth}(u, c), \hat{l})\}, [\,]\rangle$;
2. $o = \mathsf{login}(ck, u, c)$ with $ck = \mathit{get\_http\_ck}(K, u')$;
3. $Q = \langle W, K, N \uplus \{n \mapsto (u, (), \checkmark)\}, \{p \mapsto ((), \hat{l})\}, [\,]\rangle$;
4. $W(p) = (u', h, h', \checkmark)$;
5. $\rho(c) = \mathit{url\_label}(u)$;
6. $\hat{l} = \mathit{url\_label}(u) = \mathit{url\_label}(u')$.

Since we know that $W(p) = (u, h, h', q)$ with either $q = \times$ or $\mathit{url\_label}(u) \sqsubseteq l$, we must have $\mathit{url\_label}(u) \sqsubseteq l$. Hence, the introduction of the new network connection $\{n \mapsto (u, (), \checkmark)\}$ cannot break the invariant and we have $Q \sim_l C$;

Case (O-FLUSH): let $P = \langle W, K, N, T, o\rangle \xrightarrow{o} \langle W, K, N, T, [\,]\rangle = Q$. Since $P \sim_l C = \langle W', K', N', \{\}, [\,]\rangle$, we know that $l \vdash \mathit{tainted}(o)$ holds true, hence the conclusion $Q \sim_l C$ follows;

5) let $P \xrightarrow{o} Q$ with $l \vdash \mathit{untainted}(o)$. We observe that $P \xrightarrow{o} Q$ can only be proved by rule (O-MIRROR), hence we show a similar statement where $P \xrightarrow{o} Q$ has been replaced by $P \overset{o}{\mapsto} Q$. The proof is by induction on the derivation of $P \overset{o}{\mapsto} Q$:

Case (O-LETCTX): we have:

1. $P = \langle W_1, K_1, N_1, \{p \mapsto (\mathsf{let}\ x = e'\ \mathsf{in}\ e, \hat{l})\}, [\,]\rangle$;
2. $Q = \langle W_1', K_1', N_1', \{p \mapsto (\mathsf{let}\ x = e''\ \mathsf{in}\ e, \hat{l})\}, [\,]\rangle$;
3. $\hat{P} = \langle W_1, K_1, N_1, \{p \mapsto (e', \hat{l})\}, [\,]\rangle \overset{o}{\mapsto} \langle W_1', K_1', N_1', \{p \mapsto (e'', \hat{l})\}, [\,]\rangle = \hat{Q}$.

By Lemma 27 we know that $W_1(p) = (u, h, h', \checkmark)$ with $\mathit{url\_label}(u) \not\sqsubseteq l$, hence we must have:

$$P' = \langle W_2, K_2, N_2, \{p \mapsto (\mathsf{let}\ x = e'\ \mathsf{in}\ e, \hat{l})\}, [\,]\rangle,$$

with $W_2(p) = (u, h, h', \checkmark)$ by the assumption $P \sim_l P'$. Let $\hat{P}' = \langle W_2, K_2, N_2, \{p \mapsto (e', \hat{l})\}, [\,]\rangle$, we have $\hat{P} \sim_l \hat{P}'$, hence by induction hypothesis we get $\hat{P}' \overset{o}{\mapsto} \langle W_2', K_2', N_2', \{p \mapsto (e'', \hat{l})\}, [\,]\rangle = \hat{Q}'$ with $\hat{Q} \sim_l \hat{Q}'$. Let then $Q' = \langle W_2', K_2', N_2', \{p \mapsto (\mathsf{let}\ x = e''\ \mathsf{in}\ e, \hat{l})\}, [\,]\rangle$, we can conclude $Q \sim_l Q'$;

Case (O-XHR): we have:

1. $P = \langle W_1 \uplus \{p \mapsto (u', h, h', q)\}, K_1, N_1, \{p \mapsto (\mathsf{xhr}(u, \lambda x.e), \hat{l})\}, [\,]\rangle$;

2. $o = \mathsf{xhr\_req}(ck, u)$;
3. $Q = \langle W_1 \uplus \{p \mapsto (u', h, h' \uplus \{n \mapsto \lambda x.e\}, q)\}, K_1, N_1 \uplus \{n \mapsto (u, p, q')\}, \{p \mapsto ((), l)\}, [\,]\rangle$;
4. $\hat{l} \neq \bot \Rightarrow \hat{l} = \mathit{url\_label}(u) = \mathit{url\_label}(u') \wedge q = \checkmark$;
5. $q = \checkmark \wedge \mathit{url\_label}(u) = \mathit{url\_label}(u') \Rightarrow ck = \mathit{get\_http\_ck}(K_1, u) \wedge q' = \checkmark$;
6. $q = \times \vee \mathit{url\_label}(u) \neq \mathit{url\_label}(u') \Rightarrow ck = \{\} \wedge q' = \times$.

By Lemma 27 we know that $q = \checkmark$ and $\mathit{url\_label}(u') \not\sqsubseteq l$, hence we must have:

$$P' = \langle W_2 \uplus \{p \mapsto (u', h, h', \checkmark)\}, K_2, N_2, \{p \mapsto (\mathsf{xhr}(u, \lambda x.e), \hat{l})\}, [\,]\rangle,$$

by the assumption $P \sim_l P'$. We distinguish two cases:

- if $\mathit{url\_label}(u) = \mathit{url\_label}(u')$, then:

  $$Q = \langle W_1 \uplus \{p \mapsto (u', h, h' \uplus \{n \mapsto \lambda x.e\}, \checkmark)\}, K_1, N_1 \uplus \{n \mapsto (u, p, \checkmark)\}, \{p \mapsto ((), l)\}, [\,]\rangle,$$

  and $o = \mathsf{xhr\_req}(ck, u)$ with $ck = \mathit{get\_http\_ck}(K_1, u)$. We then have:

  $$P' \xrightarrow{o'} \langle W_2 \uplus \{p \mapsto (u', h, h' \uplus \{n \mapsto \lambda x.e\}, \checkmark)\}, K_2, N_2 \uplus \{n \mapsto (u, p, \checkmark)\}, \{p \mapsto ((), l)\}, [\,]\rangle = Q',$$

  with $Q \sim_l Q'$ and $o' = \mathsf{xhr\_req}(ck', u)$ with $ck' = \mathit{get\_http\_ck}(K_2, u)$. To conclude we just need to show that $o = o'$, i.e., that $ck = ck'$: this follows by Lemma 24;
- if $\mathit{url\_label}(u) \neq \mathit{url\_label}(u')$, then:

  $$Q = \langle W_1 \uplus \{p \mapsto (u', h, h' \uplus \{n \mapsto \lambda x.e\}, \checkmark)\}, K_1, N_1 \uplus \{n \mapsto (u, p, \times)\}, \{p \mapsto ((), l)\}, [\,]\rangle,$$

  and $o = \mathsf{xhr\_req}(\{\}, u)$. We then have:

  $$P' \xrightarrow{o'} \langle W_2 \uplus \{p \mapsto (u', h, h' \uplus \{n \mapsto \lambda x.e\}, \checkmark)\}, K_2, N_2 \uplus \{n \mapsto (u, p, \times)\}, \{p \mapsto ((), l)\}, [\,]\rangle = Q',$$

  with $Q \sim_l Q'$ and $o' = o = \mathsf{xhr\_req}(\{\}, u)$;

Case (O-LOGIN): in this case we know that:

1. $P = \langle W_1, K_1, N_1, \{p \mapsto (\mathsf{auth}(u, c), \hat{l})\}, [\,]\rangle$;
2. $o = \mathsf{login}(ck, u, c)$ with $ck = \mathit{get\_http\_ck}(K_1, u)$;
3. $Q = \langle W_1, K_1, N_1 \uplus \{n \mapsto (u, (), \checkmark)\}, \{p \mapsto ((), \hat{l})\}, [\,]\rangle$;
4. $W(p) = (u', h, h', \checkmark)$;
5. $\rho(c) = \mathit{url\_label}(u)$;
6. $\hat{l} = \mathit{url\_label}(u) = \mathit{url\_label}(u')$.

By Lemma 27 we know that $W_1(p) = (u, h, h', \checkmark)$ with $\mathit{url\_label}(u) \not\sqsubseteq l$, hence we must have:

$$P' = \langle W_2, K_2, N_2, \{p \mapsto (\mathsf{auth}(u, c), \hat{l})\}, [\,]\rangle,$$

with $W_2(p) = \{p \mapsto (u, h, h', \checkmark)\}$ by the assumption $P \sim_l P'$. We then have:

$$P' \xrightarrow{o'} \langle W_2, K_2, N_2 \uplus \{n \mapsto (u, (), \checkmark)\}, \{p \mapsto ((), \hat{l})\}, [\,]\rangle = Q',$$

with $Q \sim_l Q'$ and $o' = \mathsf{login}(ck', u, c)$ with $ck' = \mathit{get\_http\_ck}(K_2, u)$. To conclude we just need to show that $o = o'$, i.e., that $ck = ck'$: this follows by Lemma 24;

Case (O-FLUSH): in this case we know that:

1. $P = \langle W_1, K_1, N_1, T_1, o\rangle$;
2. $Q = \langle W_1, K_1, N_1, T_1, [\,]\rangle$.

Let $P \sim_l P'$. Since $l \vdash \mathit{untainted}(o)$ holds true, we know that $P' = \langle W_2, K_2, N_2, T_2, o\rangle$, hence we have $P' \xrightarrow{o} Q' = \langle W_2, K_2, N_2, T_2, [\,]\rangle$ with $Q \sim_l Q'$.

□

Definition 13 (Corresponding States)**.** *We say that $\sigma = \langle Q, I, \tau, M\rangle$ and $\xi = \langle Q', I', \tau'\rangle$ are* corresponding *for a security label l, written $\sigma \approx_l \xi$, if and only if:*

1) $l \models \sigma$;

2) $Q \sim_l Q'$*;*
3) $[i \mid i \in I \wedge l \vdash \mathit{untainted}(i)] = [i \mid i \in I' \wedge l \vdash \mathit{untainted}(i)]$*;*
4) $\forall o \in \mathcal{O} : \forall l' \not\sqsubseteq l : \tau(o) = l' \Leftrightarrow \tau'(o) = l'$.

Lemma 30 (Untainted Login)**.** *Let* τ_1 *and* τ_2 *be two trust functions such that:*

$$\forall o \in \mathcal{O} : \forall l' \not\sqsubseteq l : \tau_1(o) = l' \Leftrightarrow \tau_2(o) = l'.$$

If $\tau_1 \xrightarrow{o} \tau_1'$*, then there exists* τ_2' *such that* $\tau_2 \xrightarrow{o} \tau_2'$ *and:*

$$\forall o \in \mathcal{O} : \forall l' \not\sqsubseteq l : \tau_1'(o) = l' \Leftrightarrow \tau_2'(o) = l'.$$

Proof. By a case analysis on the rule applied to prove $\tau_1 \xrightarrow{o} \tau_1'$. If the reduction rule is (A-NIL), we have $\tau_1' = \tau_1$ and the conclusion follows by letting $\tau_2 \xrightarrow{o} \tau_2$ by (A-NIL). Otherwise, let $o = \mathsf{login}(ck, u, c)$ and assume that either (A-FIX) or (A-SRV) was the applied reduction rule. We then let $\tau_2 \xrightarrow{o} \tau_2'$ by applying the same rule and we notice that our property of interest must be preserved after the reduction step. Indeed, we can distinguish two cases: if $\rho(c) \sqsubseteq l$, then the conclusion is immediate, since any increase of trust is bounded above by l. Otherwise, let $\rho(c) = l' \not\sqsubseteq l$, then we can create a mismatch upon reduction only if there exists o' such that $\tau_1(o') \sqsubseteq \rho(c)$ and $\tau_2(o') \not\sqsubseteq \rho(c)$ or vice-versa, but this is excluded by our hypothesis on the two trust functions. □

Definition 14 (Weak Move)**.** *We write* $l \vdash \xi \overset{\alpha}{\Rightarrow} \xi'$ *for either of the following:*

1) $\xi = \xi'$*;*
2) $\exists \beta_1, \ldots, \exists \beta_n : \forall i \in [1, n] : l \vdash \mathit{tainted}(\beta_i) \wedge \xi \xrightarrow{\beta_1} \xi_1 \xrightarrow{\beta_2} \ldots \xrightarrow{\beta_n} \xi'$*;*
3) $\exists \beta_1, \ldots, \exists \beta_n : \forall i \in [1, n] : l \vdash \mathit{tainted}(\beta_i) \wedge \xi \xrightarrow{\beta_1} \xi_1 \xrightarrow{\beta_2} \ldots \xrightarrow{\beta_j} \xi_j \xrightarrow{\alpha} \xi_{j+1} \xrightarrow{\beta_{j+1}} \ldots \xrightarrow{\beta_n} \xi'$.

Theorem 3 (Simulation)**.** *If* $\sigma \approx_l \xi$ *and* $l \vdash \sigma \xrightarrow{\alpha} \sigma'$*, then:*

1) if $\alpha = o$ *and* $l \vdash \mathit{untainted}(o)$*, then* $\xi \xrightarrow{\alpha} \xi'$ *and* $\sigma' \approx_l \xi'$*;*
2) otherwise, we have $l \vdash \xi \overset{\alpha}{\Rightarrow} \xi'$ *and* $\sigma' \approx_l \xi'$.

Proof. By a case analysis on the rule applied to prove $l \vdash \sigma \xrightarrow{\alpha} \sigma'$:

Case (AS-IN): we have $\sigma = \langle C, i :: I, \tau, M \rangle$ and $\sigma' = \langle P, I, \tau, M \rangle$ with $\alpha = i$ and $C \xrightarrow{i} P$. Let $\xi = \langle Q, I', \tau' \rangle$, then by the assumption $\sigma \approx_l \xi$ we know that:

1) $l \models \sigma$;
2) $C \sim_l Q$;
3) $[i' \mid i' \in i :: I \wedge l \vdash \mathit{untainted}(i')] = [i' \mid i' \in I' \wedge l \vdash \mathit{untainted}(i')]$;
4) $\forall o \in \mathcal{O} : \forall l' \not\sqsubseteq l : \tau(o) = l' \Leftrightarrow \tau'(o) = l'$.

We distinguish two cases, based on Q being a producer or a consumer state. First, let Q be a producer state P'. We first prove that there exists ξ' such that $\xi \xrightarrow{o} \xi'$ with $l \vdash \mathit{tainted}(o)$ and $\sigma \approx_l \xi'$. We observe that, since $C \sim_l P'$ by condition (2), we must have $l \vdash \mathit{tainted}(o)$ for any o such that $P' \xrightarrow{o} Q'$ and $C \sim_l Q'$ by Lemma 29. We have:

$$\text{(S-OUT)}\quad \frac{P' \xrightarrow{o} Q' \qquad \tau' \xrightarrow{o} \tau''}{\xi = \langle P', I', \tau' \rangle \xrightarrow{o} \langle Q', I', \tau'' \rangle = \xi'}$$

We now need to prove the following four conditions:

a. $l \models \sigma$;
b. $C \sim_l Q'$;
c. $[i' \mid i' \in i :: I \wedge l \vdash \mathit{untainted}(i')] = [i' \mid i' \in I' \wedge l \vdash \mathit{untainted}(i')]$;
d. $\forall o \in \mathcal{O} : \forall l' \not\sqsubseteq l : \tau(o) = l' \Leftrightarrow \tau''(o) = l'$.

Conditions (a) and (c) are exactly conditions (1) and (3). Condition (b) was proved above. Since $l \models \sigma$ by condition (1), we know that $\tau, l \vdash_\diamond M$ holds true. Moreover, we have that $\forall n \in \mathit{fn}(o) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$, which implies $\tau, l, M \Vdash o$ by (IS-OUT). Hence, condition (d) follows by condition (4), using Lemma 21. We iterate this reasoning until we reach a consumer state C' from P': this is always possible, since FF^+ can never loop. We

complete the case by performing the same reasoning we carry out below, where we assume that Q is a consumer state.

Let then Q be a consumer state C'. We further distinguish two cases, based on $ev\_label(i)$. First, assume that $ev\_label(i) \sqsubseteq l$: since $l \models \sigma$ holds true, we know that $\vdash_\diamond i$ holds also true. By Lemma 14 we know that $\forall n \in fn(i) : \exists l' \sqsubseteq ev\_label(i) \sqsubseteq l : n \in \mathcal{N}_{l'}$, hence we have $l \vdash tainted(i)$. We show that $\sigma' \approx_l \xi$, i.e., we prove the following four conditions:

a. $l \models \sigma'$;
b. $P \sim_l C'$;
c. $[i' \mid i' \in I \wedge l \vdash untainted(i')] = [i' \mid i' \in I' \wedge l \vdash untainted(i')]$;
d. $\forall o \in \mathcal{O} : \forall l' \not\sqsubseteq l : \tau(o) = l' \Leftrightarrow \tau'(o) = l'$.

Condition (a) follows by condition (1), using Theorem 2. Since $C \sim_l C'$ by condition (2), we have $P \sim_l C'$ by Lemma 29 and we proved (b). Condition (c) immediately follows by condition (3), since we proved that $l \vdash tainted(i)$ holds true. Condition (d) is exactly condition (4).

Let now $ev\_label(i) \not\sqsubseteq l$: since $l \models \sigma$ holds true, we know that $\vdash_\diamond i$ holds also true, hence we have $l \vdash untainted(i)$. Thus, by condition (3) we have $\exists i_1, \ldots, \exists i_n, \exists I'' : I' = i_1 :: \ldots :: i_n :: i :: I''$, where $\forall j \in [1, n] : l \vdash tainted(i_j)$ and $[i' \mid i' \in I \wedge l \vdash untainted(i')] = [i' \mid i' \in I'' \wedge l \vdash untainted(i')]$. We further distinguish two sub-cases:

- first, let $n = 0$, i.e., assume there is no tainted input event preceding i in I'. We prove that there exists ξ' such that $\xi \xrightarrow{i} \xi'$ and $\sigma' \approx_l \xi'$. Since $C \sim_l C'$ by condition (2) and $l \vdash untainted(i)$ holds true, we know that $C' \xrightarrow{i} P'$ for some P' such that $P \sim_l P'$ by Lemma 29. Hence, we have:

$$\text{(S-IN)}\ \frac{C' \xrightarrow{i} P'}{\xi = \langle C', i :: I'', \tau \rangle \xrightarrow{i} \langle P', I'', \tau \rangle = \xi'}$$

  We now need to show that $\sigma' \approx_l \xi'$ holds true, i.e., we have to prove the following four conditions:

  a. $l \models \sigma'$;
  b. $P \sim_l P'$;
  c. $[i' \mid i' \in I \wedge l \vdash untainted(i')] = [i' \mid i' \in I'' \wedge l \vdash untainted(i')]$;
  d. $\forall o \in \mathcal{O} : \forall l' \not\sqsubseteq l : \tau(o) = l' \Leftrightarrow \tau'(o) = l'$.

  Condition (a) follows by condition (1), using Theorem 2. Conditions (b) and (c) have been proved above, while condition (d) is exactly condition (4).

- now let $n > 0$, i.e., we have some tainted input events $i_1, \ldots, i_n$ preceding i in I'. We first prove that there exists ξ' such that $\xi \xrightarrow{i_1} \xi'$ with $l \vdash tainted(i_1)$ and $\sigma \approx_l \xi'$. We already showed that $l \vdash tainted(i_1)$ holds true and we know $C' \xrightarrow{i_1} P'$ for some P' by definition of reactive system, hence we have:

$$\text{(S-IN)}\ \frac{C' \xrightarrow{i_1} P'}{\xi = \langle C', I', \tau \rangle \xrightarrow{i_1} \langle P', i_2 :: \ldots :: i_n :: i :: I'', \tau \rangle}$$

  We need to show that $\sigma \approx_l \xi'$, i.e., we prove the following four conditions:

  a. $l \models \sigma$;
  b. $C \sim_l P'$;
  c. $[i' \mid i' \in i :: I \wedge l \vdash untainted(i')] = [i' \mid i' \in i_2 :: \ldots :: i_n :: i :: I'' \wedge l \vdash untainted(i')]$;
  d. $\forall o \in \mathcal{O} : \forall l' \not\sqsubseteq l : \tau(o) = l' \Leftrightarrow \tau''(o) = l'$.

  Condition (a) is exactly condition (1). Condition (b) follows by the assumption $C \sim_l C'$, using Lemma 29. Condition (c) follows by condition (3), since we proved that $l \vdash tainted(i_1)$ holds true. Condition (d) is exactly condition (4). We iterate this reasoning until we can consume the untainted input i and we conclude like in the previous sub-case;

Case (AS-OUT): we have $\sigma = \langle P, I, \tau, M \rangle$ and $\sigma' = \langle Q, I, \tau'', M \rangle$ with $\alpha = o$ and $P \xrightarrow{o} Q$ and $\tau \xrightarrow{o} \tau''$. Let $\xi = \langle Q', I', \tau' \rangle$, then by the assumption $\sigma \approx_l \xi$ we know that:

1) $l \models \sigma$;

2) $P \sim_l Q'$;
3) $[i \mid i \in I \wedge l \vdash \mathit{untainted}(i)] = [i \mid i \in I' \wedge l \vdash \mathit{untainted}(i)]$;
4) $\forall o \in \mathcal{O} : \forall l' \not\sqsubseteq l : \tau(o) = l' \Leftrightarrow \tau'(o) = l'$.

We distinguish two cases, based on o being tainted or not. If $l \vdash \mathit{tainted}(o)$ holds true, we show that $\sigma' \approx_l \xi$, i.e., we prove the following four conditions:

a. $l \models \sigma'$;
b. $Q \sim_l Q'$;
c. $[i \mid i \in I \wedge l \vdash \mathit{untainted}(i')] = [i \mid i \in I' \wedge l \vdash \mathit{untainted}(i')]$;
d. $\forall o \in \mathcal{O} : \forall l' \not\sqsubseteq l : \tau''(o) = l' \Leftrightarrow \tau'(o) = l'$.

Condition (a) follows by condition (1), using Theorem 2. Since $P \sim_l Q'$ by condition (2), we have $Q \sim_l Q'$ by Lemma 29 and we proved (b). Condition (c) is exactly condition (3) above. Since $l \models \sigma$ by condition (1), we know that $\tau, l \vdash_\diamond M$ holds true. Moreover, we have that $\forall n \in \mathit{fn}(o) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$, which implies $\tau, l, M \Vdash o$ by (IS-OUT). Hence, condition (d) follows by condition (4), using Lemma 21.

Let now $l \vdash \mathit{untainted}(o)$ hold true, we show that there exists ξ' such that $\xi \xrightarrow{o} \xi'$ and $\sigma' \approx_l \xi'$. Since $P \sim_l Q'$ by condition (2), we know that Q' must be a producer state P' and $P' \xrightarrow{o} Q''$ for some Q'' such that $Q \sim_l Q''$ by Lemma 29. Since $l \models P$ by condition (1), by Lemma 30 we know that there exists $\hat{\tau}$ such that $\tau' \xrightarrow{o} \hat{\tau}$ and $\forall o \in \mathcal{O} : \forall l' \not\sqsubseteq l : \tau''(o) = l' \Leftrightarrow \hat{\tau}(o) = l'$. Hence, we have:

$$\text{(S-OUT)}\quad \frac{P' \xrightarrow{o} Q'' \qquad \tau' \xrightarrow{o} \hat{\tau}}{\xi = \langle P', I', \tau' \rangle \xrightarrow{o} \langle Q'', I', \hat{\tau} \rangle = \xi'}$$

We now need to show that $\sigma' \approx_l \xi'$ holds true, i.e., we have to prove the following four conditions:

a. $l \models \sigma'$;
b. $Q \sim_l Q''$;
c. $[i \mid i \in I \wedge l \vdash \mathit{untainted}(i)] = [i \mid i \in I' \wedge l \vdash \mathit{untainted}(i)]$;
d. $\forall o \in \mathcal{O} : \forall l' \not\sqsubseteq l : \tau''(o) = l' \Leftrightarrow \hat{\tau}(o) = l'$.

Condition (a) follows by condition (1), using Theorem 2. Conditions (b) and (d) have been proved above. Condition (c) is exactly condition (3);

Case (AS-GETIN): we have $\sigma = \langle Q, i :: I, \tau, M \rangle$ and $\sigma' = \langle Q, I, \tau, M \cup \{i\} \rangle$ with $\alpha = \bullet$ and $\tau, l \dagger i$ being true by the premise of the rule. Let $\xi = \langle Q', I', \tau' \rangle$, then by the assumption $\sigma \approx_l \xi$ we know that:

1) $l \models \sigma$;
2) $Q \sim_l Q'$;
3) $[i \mid i \in i :: I \wedge l \vdash \mathit{untainted}(i)] = [i \mid i \in I' \wedge l \vdash \mathit{untainted}(i)]$;
4) $\forall o \in \mathcal{O} : \forall l' \not\sqsubseteq l : \tau(o) = l' \Leftrightarrow \tau'(o) = l'$.

We want to show that $\sigma' \approx_l \xi$, i.e., we have to prove the following four conditions:

1) $l \models \sigma'$;
2) $Q \sim_l Q'$;
3) $[i \mid i \in I \wedge l \vdash \mathit{untainted}(i)] = [i \mid i \in I' \wedge l \vdash \mathit{untainted}(i)]$;
4) $\forall o \in \mathcal{O} : \forall l' \not\sqsubseteq l : \tau(o) = l' \Leftrightarrow \tau'(o) = l'$.

Condition (a) follows by condition (1), using Theorem 2. Condition (b) is exactly condition (2) above. Since $\tau, l \dagger i$ holds true, we know that $\mathit{ev\_label}(i) \sqsubseteq l$, hence we know that $l \vdash \mathit{untainted}(i)$ does not hold: this is enough to prove condition (c) from condition (3). Condition (d) is exactly condition (4) above;

Case (AS-GETOUT): we have $\sigma = \langle P, I, \tau, M \rangle$ and $\sigma' = \langle Q, I, \tau, M \cup \{o\} \rangle$ with $\alpha = \bullet$ and $P \xrightarrow{o} Q$. By the premises of the reduction rule, we also know that $\tau, l \dagger o$. Let $\xi = \langle Q', I', \tau' \rangle$, then by the assumption $\sigma \approx_l \xi$ we know that:

1) $l \models \sigma$;
2) $P \sim_l Q'$;
3) $[i \mid i \in I \wedge l \vdash \mathit{untainted}(i)] = [i \mid i \in I' \wedge l \vdash \mathit{untainted}(i)]$;
4) $\forall o \in \mathcal{O} : \forall l' \not\sqsubseteq l : \tau(o) = l' \Leftrightarrow \tau'(o) = l'$.

We want to show that $\sigma' \approx_l \xi$, i.e., we have to prove the following four conditions:

1) $l \models \sigma'$;
2) $Q \sim_l Q'$;
3) $[i \mid i \in I \wedge l \vdash \mathit{untainted}(i)] = [i \mid i \in I' \wedge l \vdash \mathit{untainted}(i)]$;
4) $\forall o \in \mathcal{O} : \forall l' \not\sqsubseteq l : \tau(o) = l' \Leftrightarrow \tau'(o) = l'$.

Condition (a) follows by condition (1), using Theorem 2. Since $\tau, l \dagger o$ holds, we know that $\mathit{ev\_label}(o) \sqsubseteq l$: this implies that $l \vdash \mathit{tainted}(o)$ holds true by Lemma 4. Given that $P \sim_l Q'$ by condition (2), we have $Q \sim_l Q'$ by Lemma 29 and we proved condition (b). Conditions (c) and (d) are exactly (3) and (4);

Case (AS-HEARIN): we have $\sigma = \langle Q, i :: I, \tau, M \rangle$ and $\sigma' = \langle Q, i :: I, \tau, M \cup \{i\} \rangle$ with $\alpha = \bullet$ and $\tau, l \,?\, i$ being true by the premise of the rule. Let $\xi = \langle Q', I', \tau' \rangle$, then by the assumption $\sigma \approx_l \xi$ we know that:

1) $l \models \sigma$;
2) $Q \sim_l Q'$;
3) $[i' \mid i' \in i :: I \wedge l \vdash \mathit{untainted}(i')] = [i' \mid i' \in I' \wedge l \vdash \mathit{untainted}(i')]$;
4) $\forall o \in \mathcal{O} : \forall l' \not\sqsubseteq l : \tau(o) = l' \Leftrightarrow \tau'(o) = l'$.

We want to show that $\sigma' \approx_l \xi$, i.e., we have to prove the following four conditions:

1) $l \models \sigma'$;
2) $Q \sim_l Q'$;
3) $[i' \mid i \in i' :: I \wedge l \vdash \mathit{untainted}(i')] = [i' \mid i' \in I' \wedge l \vdash \mathit{untainted}(i')]$;
4) $\forall o \in \mathcal{O} : \forall l' \not\sqsubseteq l : \tau(o) = l' \Leftrightarrow \tau'(o) = l'$.

Condition (a) follows by condition (1), using Theorem 2. Conditions (b), (c) and (d) are exactly conditions (2), (3) and (4) above;

Case (AS-HEAROUT): analogous to case (AS-OUT);

Case (AS-SYNIN): we have $\sigma = \langle C, I, \tau, M \rangle$ and $\sigma' = \langle P, I, \tau, M \rangle$ with $\alpha = i$ and $C \xrightarrow{i} P$. By the premises of the reduction rule, we also know that $\tau, l, M \Vdash i$. Let $\xi = \langle Q', I', \tau' \rangle$, then by the assumption $\sigma \approx_l \xi$ we know that:

1) $l \models \sigma$;
2) $C \sim_l Q'$;
3) $[i \mid i \in I \wedge l \vdash \mathit{untainted}(i)] = [i \mid i \in I' \wedge l \vdash \mathit{untainted}(i)]$;
4) $\forall o \in \mathcal{O} : \forall l' \not\sqsubseteq l : \tau(o) = l' \Leftrightarrow \tau'(o) = l'$.

We distinguish two cases, based on Q' being a producer or a consumer state. First, let Q' be a producer state P'. We first prove that there exists ξ' such that $\xi \xrightarrow{o} \xi'$ with $l \vdash \mathit{tainted}(o)$ and $\sigma \approx_l \xi'$. We observe that, since $C \sim_l P'$ by condition (2), we must have $l \vdash \mathit{tainted}(o)$ for any o such that $P' \xrightarrow{o} Q'$ by Lemma 29. We have:

$$\frac{P' \xrightarrow{o} Q' \qquad \tau' \xrightarrow{o} \tau''}{\xi = \langle P', I', \tau' \rangle \xrightarrow{o} \langle Q', I', \tau'' \rangle = \xi'} \text{(S-OUT)}$$

We now need to prove the following four conditions:

a. $l \models \sigma$;
b. $C \sim_l Q'$;
c. $[i' \mid i' \in i :: I \wedge l \vdash \mathit{untainted}(i')] = [i' \mid i' \in I' \wedge l \vdash \mathit{untainted}(i')]$;
d. $\forall o \in \mathcal{O} : \forall l' \not\sqsubseteq l : \tau(o) = l' \Leftrightarrow \tau''(o) = l'$.

Conditions (a) and (c) are exactly conditions (1) and (3). Condition (b) follows by the assumption $C \sim_l P'$, using Lemma 29. Since $l \models \sigma$ by condition (1), we know that $\tau, l \vdash_\diamond M$ holds true. Moreover, we have that $\forall n \in \mathit{fn}(o) : \exists l' \sqsubseteq l : n \in \mathcal{N}_{l'}$, which implies $\tau, l, M \Vdash o$ by (IS-OUT). Hence, condition (d) follows by condition (4), using Lemma 21. We iterate this reasoning until we reach a consumer state C' from P': this is always possible, since FF^+ can never loop. We complete the case by performing the same reasoning we carry out below, where we assume that Q' is a consumer state.

Let then Q' be a consumer state C'. Since $l \models \sigma$ by condition (1), we know that $\tau, l \vdash_\diamond M$ holds true. Given that $\tau, l, M \Vdash i$ holds true, we have $l \vdash \mathit{tainted}(i)$ by Lemma 17. We want to show that $\sigma' \approx_l \xi$, i.e., we have to prove the following four conditions:

1) $l \models \sigma'$;

2) $Q \sim_l C'$;
3) $[i \mid i \in I \wedge l \vdash untainted(i)] = [i \mid i \in I' \wedge l \vdash untainted(i)]$;
4) $\forall o \in \mathcal{O} : \forall l' \not\sqsubseteq l : \tau(o) = l' \Leftrightarrow \tau'(o) = l'$.

Condition (a) follows by condition (1), using Theorem 2. Condition (b) follows by condition (2), using Lemma 29. Condition (c) and (d) are exactly conditions (3) and (4);

Case (AS-SYNOUT): we have $\sigma = \langle Q, I, \tau, M\rangle$ and $\sigma' = \langle Q, I, \tau'', M\rangle$ with $\alpha = o$ and $\tau \xrightarrow{o} \tau''$. By the premises of the reduction rule, we also know that $\tau, l, M \Vdash o$. Let $\xi = \langle Q', I', \tau'\rangle$, then by the assumption $\sigma \approx_l \xi$ we know that:

1) $l \models \sigma$;
2) $Q \sim_l Q'$;
3) $[i \mid i \in I \wedge l \vdash untainted(i)] = [i \mid i \in I' \wedge l \vdash untainted(i)]$;
4) $\forall o \in \mathcal{O} : \forall l' \not\sqsubseteq l : \tau(o) = l' \Leftrightarrow \tau'(o) = l'$.

Since $l \models \sigma$ by condition (1), we know that $\tau, l \vdash_\diamond M$ holds true. Given that $\tau, l, M \Vdash o$ holds true, we have $l \vdash tainted(o)$ by Lemma 18. Hence, it is enough to show that $\sigma' \approx_l \xi$, i.e., we have to prove the following four conditions:

1) $l \models \sigma'$;
2) $Q \sim_l Q'$;
3) $[i \mid i \in I \wedge l \vdash untainted(i)] = [i \mid i \in I' \wedge l \vdash untainted(i)]$;
4) $\forall o \in \mathcal{O} : \forall l' \not\sqsubseteq l : \tau''(o) = l' \Leftrightarrow \tau'(o) = l'$.

Condition (a) follows by condition (1), using Theorem 2. Conditions (b) and (c) are exactly conditions (2) and (3). Since $\tau, l \vdash_\diamond M$ and $\tau, l, M \Vdash o$ hold true, condition (d) follows by condition (4), using Lemma 21.

□

In the following results, let $C_0^+ = \langle\{\}, \{\}, \{\}, \{\}, [\,]\rangle$ be the initial state of FF^+.

Lemma 31 (Type-checking)**.** *$l \models \langle C_0^+, I, \tau_\perp, \emptyset\rangle$ for any security label l and any well-formed input stream I.*

Proof. Notice that $l \models C_0^+$ holds true, since all the conditions dictated by Definition 7 are trivially met. To show that $l \models \langle C_0^+, I, \tau_\perp, \emptyset\rangle$ holds true, we simply have to prove the remaining conditions in Definition 9: condition 2 holds true, since we are assuming that I is well-formed, while all the other conditions are trivial. □

Lemma 32 (Simulation for Initial State)**.** *$C_0^+ \sim_l C_0^+$ for any security label l.*

Proof. Notice that $l \models C_0^+$ holds true, since all the conditions dictated by Definition 7 are trivially met. All the other conditions in Definition 12 are trivially true. □

Lemma 33 (Simulation for FF^+)**.** *$\langle C_0^+, I, \tau_\perp, \emptyset\rangle \approx_l \langle C_0^+, I, \tau_\perp\rangle$ for any security label l and any well-formed input stream I.*

Proof. By Lemma 31 we know that $l \models \langle C_0^+, I, \tau_\perp, \emptyset\rangle$. By Lemma 32 we also know that $C_0^+ \sim_l C_0^+$ holds true. The remaining conditions of Definition 13 are trivially met. □

Lemma 34 (Determinism for FF^+)**.** *If $\xi \xrightarrow{\alpha} \xi'$ and $\xi \xrightarrow{\beta} \xi''$, then $\alpha = \beta$ and $\xi' = \xi''$.*

Proof. By a case analysis on the rule applied to prove $\xi \xrightarrow{\alpha} \xi'$.

If the applied rule is (S-IN), then $\xi = \langle C, i :: I, \tau\rangle$ and $\xi' = \langle P, I, \tau\rangle$ with $\alpha = i$ and $C \xrightarrow{i} P$. Hence, $\xi \xrightarrow{\beta} \xi''$ can only be proved by (S-IN) and we must have $\alpha = \beta$. Let then $\xi'' = \langle P', I, \tau\rangle$ for some P' such that $C \xrightarrow{i} P'$, we can show that $P = P'$ by a case analysis on the rule applied to show $C \xrightarrow{i} P$.

If the applied rule is (S-OUT), then $\xi = \langle P, I, \tau\rangle$ and $\xi' = \langle Q, I, \tau'\rangle$ with $P \xrightarrow{o} Q$, $\tau \xrightarrow{o} \tau'$ and $\alpha = o$. Hence, $\xi \xrightarrow{\beta} \xi''$ can only be proved by (S-OUT). Let $\xi'' = \langle Q', I, \tau''\rangle$ for some Q' such that $P \xrightarrow{o'} Q'$ and some τ'' such that $\tau \xrightarrow{o'} \tau''$, we can show that $o = o'$ and $Q = Q'$ by a case analysis on the rule applied to show $P \xrightarrow{o} Q$. Similarly, we can show that $\tau' = \tau''$ by a case analysis on the rule applied to show $\tau \xrightarrow{o} \tau'$. □

Theorem 4 (Session Integrity)**.** *FF^+ enforces session integrity for any well-formed trace.*

Proof. By definition of trace we know that $\tau_\perp \vdash C_0^+(I) \rightsquigarrow O$. Now let l be an arbitrary opponent and assume we have $\tau_\perp, l \vdash C_0^+(I) \rightsquigarrow O'$ for some trace O'. By Lemma 1 we know that $\tau_\perp \vdash C_0^+(I) \rightsquigarrow O$ implies that:

$$\underbrace{\langle Q_0, I_0, \tau_0\rangle}_{\xi_0} \xrightarrow{\alpha_1} \underbrace{\langle Q_1, I_1, \tau_1\rangle}_{\xi_1} \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_n} \underbrace{\langle Q_n, I_n, \tau_n\rangle}_{\xi_n},$$

where $Q_0 = C_0^+$, $I_0 = I$, $\tau_0 = \tau_\perp$ and $O = [(o_i, \tau_{i-1}(o_i))_{1 \leq i \leq n} \mid o_i = \alpha_i]$. Similarly, by Lemma 2 we know that $\tau_\perp, l \vdash C_0^+(I) \rightsquigarrow O'$ implies that:

$$l \vdash \underbrace{\langle Q'_0, I'_0, \tau'_0, M_0\rangle}_{\sigma_0} \xrightarrow{\beta_1} \underbrace{\langle Q'_1, I'_1, \tau'_1, M_1\rangle}_{\sigma_1} \xrightarrow{\beta_2} \dots \xrightarrow{\beta_m} \underbrace{\langle Q'_m, I'_m, \tau'_m, M_m\rangle}_{\sigma_m},$$

where $Q'_0 = C_0^+$, $I'_0 = I$, $\tau'_0 = \tau_\perp$, $M_0 = \emptyset$ and $O' = [(o_i, \tau'_{i-1}(o_i))_{1 \leq i \leq m} \mid o_i = \beta_i]$.

We observe that $l \models \sigma_0$ by Lemma 31, hence $l \models \sigma_j$ holds true for any $j \in [0, m]$ by Theorem 2. In particular, this implies that:

$$\forall j \in [0,m] : \forall o \in \mathcal{O} : \tau'_j(o) = l' \not\sqsubseteq l \Rightarrow o \in \{*\_\mathsf{req}(ck, u) \mid \mathit{url\_label}(u) = l' \wedge \mathit{ck\_vals}(ck) \cap \mathcal{N}_{l'} \neq \emptyset\}.$$

For this reason, we know that for any output event $o = \beta_j \in \{\beta_1, \dots, \beta_m\}$ such that $l \vdash \mathit{tainted}(o)$ we must have $\tau'_{j-1}(o) \sqsubseteq l$; conversely, whenever $\tau'_{j-1}(o) \not\sqsubseteq l$, we must have $l \vdash \mathit{untainted}(o)$.

If $\forall j \in [1, m] : \tau'_{j-1}(\beta_j) \sqsubseteq l$, then we conclude, since we have $O' \downarrow l' = [\,]$ for any $l' \not\sqsubseteq l$. Let then β_j be the first output event such that $\tau'_{j-1}(\beta_j) = l' \not\sqsubseteq l$, we already showed that $l \vdash \mathit{untainted}(\beta_j)$ must hold true. Since $\sigma_0 \approx_l \xi_0$ by Lemma 33 and the unattacked semantics of FF^+ is deterministic (Lemma 34), we know by Theorem 3 that there exists $k \in [0, n]$ such that:

1) $\sigma_{j-1} \approx_l \xi_k$;
2) $\forall o \in \{\alpha_1, \dots, \alpha_k\} : l \vdash \mathit{tainted}(o)$;
3) $\alpha_{k+1} = \beta_j$.

Again, we know that $\forall \alpha_i \in \{\alpha_1, \dots, \alpha_k\} : \tau_{i-1}(\alpha_i) \sqsubseteq l$ and we have $\tau_k(\alpha_{k+1}) = \tau'_{j-1}(\beta_j)$, since we are assuming $\tau'_{j-1}(\beta_j) = l' \not\sqsubseteq l$ and $\sigma_{j-1} \approx_l \xi_k$ holds true. We conclude the proof by iterating this reasoning for an appropriate number of times. □