

Better Abstractions for Secure Server-Side Scripting

Dachuan Yu Ajay Chander Hiroshi Inamura Igor Serikov

DoCoMo Communications Laboratories USA

{yu,chander,inamura,iserikov}@docomolabs-usa.com

Abstract

It is notoriously difficult to program a solid web application. Besides addressing web interactions, state maintenance, and whimsical user navigation behaviors, programmers must also avoid a minefield of security vulnerabilities. The problem is fundamentally twofold. On the one hand, we lack a clear understanding of the essence of the new computation model underlying web applications. On the other hand, we lack proper abstractions for hiding common and subtle coding details that are orthogonal to the business functionalities of specific web applications.

This report addresses both issues. First, we present a high-level language BASS for declarative server-side scripting. BASS allows programmers to work in an ideal world, using some new abstractions to tackle common but problematic aspects of web programming. The meta properties of BASS provide useful security guarantees. Second, we present a low-level language MOSS reflecting realistic web programming concepts and scenarios, thus articulating the computation model behind web programming. Finally, we present a rigorous translation from BASS to MOSS, demonstrating how the ideal programming model and security guarantees of BASS can be implemented and enforced in practice.

1. Introduction

The Web of today is commonly used as a platform for hosting sophisticated applications and services. A key component of such web-based applications is the programs running on the web servers. These server-side programs (referred to as “server programs” from now on) take client input in the form of HTTP requests, perform computation and enforce the business logics of the applications, and produce output for the client in the form of HTML pages.

With the growing popularity of web applications, especially those which manage sensitive data and perform critical functionalities (*e.g.*, online banking), the security of the server programs has become a major concern. Indeed, web applications face more security threats than conventional desktop applications [10, 26]. Some representative ones include command injection [27], cross-site scripting (XSS) [22], cross-site request forgery (CSRF) [5], and session fixation [19]. Any of these could cause serious consequences: sensitive user information could be stolen, data and associated belongings could be damaged, or service availability could be compromised.

In response to such security threats, existing languages and frameworks for server-side scripting (*e.g.*, JSP [11], PHP [1], Ruby on Rails [17]), as well as the web application security community (*e.g.*, CERT [7], OWASP [25], WASC [30]), largely promote secure coding practices (*e.g.*, input validation) and provide useful libraries in support (*e.g.*, filtering functions). However, there is no guarantee that programmers will follow the recommendations correctly, if they followed them at all. Furthermore, even if all programs are written with the security practices strictly enforced, the extra care that programmers spend on preventing vulnerabilities much dis-

tracts from the business functionalities. Take the implementation of online payments as an example. To securely code a web interaction of “obtaining payment details,” one must correctly perform input validation, maintain program states across the interaction, and prevent CSRF.

Observing that many of the security issues are orthogonal to the business logic of specific web applications, we propose some new abstractions for writing secure server programs. These abstractions provide an ideal view of some key aspects of web programming (*e.g.*, “to obtain a web input”), and hide the common security handling (*e.g.*, input validation, state maintenance, CSRF prevention). Using these abstractions, a language for server-side scripting can be given a high-level syntax and semantics that reflect useful and secure web operations, with the correct enforcement of the semantics taken care of by the underlying language implementation following established security practices once and for all. As a result, all programs written in the language, although not directly dealing with low-level security constructs, are guaranteed to be free of certain common vulnerabilities. In addition, now thinking in terms of the high-level semantics, programmers can focus more on the business logic, which results in better security and higher productivity. To some extent, the new abstractions hide security details in the same way as object creation primitives in object-oriented languages hide low-level memory management details.

On the technical side, this report presents two formal models (languages): one is an ideal model that we propose for future web programs to be written in; the other is the real-world model that underlies existing web programming. By comparing these two models side-by-side, and by formalizing a translation between them, we argue that web programming (secure server-side scripting, in particular) could benefit significantly through the use of domain-specific language abstractions. Specifically, this report makes the following contributions.

First, we identify some common aspects of web programming that are important for security, and propose new abstractions to help implement the corresponding tasks (Section 2). The new abstractions are rigorously presented in a self-contained formal language BASS (“better abstractions for server-side scripting”) (Section 3). The new BASS primitives and operational semantics shield programmers from certain malicious exploits, thus serving as an ideal model of web programming. The security guarantees of BASS are articulated in its formal semantics and meta properties.

Next, we formalize a second language MOSS (“a model of server-side scripting”) that reflects the programming model underlying existing server-side scripting (Section 4). This directly demonstrates the complexity of web programming and the benefit of the proposed BASS abstractions. In addition, MOSS is of independent value—it articulates server program execution, attacker exploits, and client-server interactions all within the same model, thus serving as a foundation for formal studies of web programming. MOSS can be used to illustrate both secure programs and vulnerable ones.




Figure 1. Idealistic work flow of online banking




Figure 2. Actual work flow of online banking

We give formal encoding of some common exploit patterns and typical attacks as an example use of MOSS.

Finally, the BASS approach to security is that the operational semantics dictates “good” behaviors only, thus leaving no room for certain attacks. A BASS compiler must carry out the operational semantics correctly. We present a formal translation from BASS to MOSS (Section 5), where the ideal model of BASS is enforced using proper MOSS primitives for manipulating security concepts. We explain how the security guarantees of BASS are maintained.

2. Server-Side Scripting and Security

2.1 An Ideal View

We use an online-banking example to explain what is involved in secure server-side scripting, and how proper abstractions can help. This application provides two services: showing the account balance (the “balance” service) and setting up a payment (the “payment” service). A user must be logged in to access the services.

Although serving multiple users, this web application logically deals with one client at a time. In an ideal view, there are multiple orthogonal instances of the server program running, each taking care of a single client. Every single instance of the program can be viewed as a sequential program of a conventional application. A work flow graph following this ideal view is shown in Figure 1.

2.2 A Limited Mechanism

The above ideal view cannot be directly implemented, because of some limitations of the underlying HTTP mechanism for web interactions. In particular, there is no persistent channel for a server program to obtain input from a client. Instead, HTTP supports a one-shot request-response model where a client requests resource identified by a URL, and a server responds with the resource if the request is accepted. Using HTTP, web interactions are typically carried out as a sequence of requests (form submissions) and responses (web pages), where the requests provide user input, and the responses present information to the user:

HTTP 0 → HTML 0 → HTTP 1 → HTML 1 → HTTP 2 → HTML 2 . . .

Using this model, a server program is often split into multiple fragments, each taking an HTTP request and producing an HTML response. In the response, there can be a web form with an embedded URL pointing to the next fragment, so that the next request is targeted correctly. Therefore, the server work flow of our banking

application is more accurately described as in Figure 2. There are 4 program fragments connected with URL embeddings, as indicated by the dashed lines. In particular, because the payment service requires user input, the structure of the service loop in the ideal view can no longer be coded as an explicit loop. Instead, a goto-style structure is exhibited through URL embeddings. Such fragmentation and low-level control structures obscure the control flow of the server program.

Besides obscurity, there is a bigger issue: since HTTP is stateless, server programs must maintain program states on their own. In the example, the user name obtained from the login input must be saved and restored explicitly across the later web interactions. In addition, one must somehow correlate incoming requests with specific clients, since multiple clients may be interacting with the server at the same time, although in logically separate transactions.

In general, a web application needs to encode states at a level above HTTP. There are several alternatives applicable to different situations. Typically, a virtual concept of “a session” is used to refer to a logical transaction. Every session is associated with a unique ID, called SID. Saved program states and incoming client requests are both identified by the SID. As a result, much code in server programs is dedicated to managing sessions. Before generating a response, a server program must save state and embed the SID in the response. Upon receiving a request, the server program must obtain an SID from the request and load state. Based on the application, some parts of the state should be saved on the server, whereas others should be on the client via the cookie or URL embedding. These routine manipulations increase the complexity of server programs, reduce productivity, and extend the chances of programming errors.

2.3 A Dangerous World

Assuming a programmer has taken care of all the above issues correctly, the result program may still not be ready for deployment. The problem there is security: clients in the real world may be malicious, or attackers may trick innocent clients into making mistakes.

Indeed, there have been many common web security vulnerabilities identified. Secure programming solutions have been proposed accordingly. Nonetheless, a programmer must be aware of all the issues involved and implement the related defenses. Most of the defenses are orthogonal to the business logic of specific web applications. Their handling further complicates server-side scripting.

In what follows, we briefly overview some representative security issues. The goal is to demonstrate the complexity of web programming, rather than to give a comprehensive list [10, 26].

CSRF: An attacker may forge a request as if it were intended by another user. This is mainly applicable when SIDs are stored in cookies. Suppose a banking program is not coded securely. Such an attack can be launched if a user, while logged in, opens a malicious email containing a crafted image link. Trying to load the image, the user’s browser may follow the link and send a request to the server asking for a payment to be set up to the attacker.

XSS: An attacker may inject malicious code into the web page that a server program sends to a victim user. In one scenario, an attacker sends to a victim a crafted link with malicious JavaScript code embedded; when the victim loads the link, a vulnerable server program may propagate the embedded code into the HTML response. The malicious code, now coming from the server, gains the privilege of the server domain. For example, it may read the cookie set by the server and send it to the attacker. This scenario is somewhat related to CSRF. However, there are also second-order attacks [24] that do not require the use of forged requests.

Session fixation: An attacker may fix an SID beforehand and trick a user into communicating with the server using it. This is mainly applicable if SIDs are embedded in URLs. An innocent user may click on a link in a malicious email which claims to be from our

banking site. The link takes the user to our banking site, but using an SID fixed by an attacker. If the server programs use the same SID for later web interactions after the user logs in, the knowledge of the SID will grant the attacker the user's privileges.

Others: Many other aspects affect security [26]. Consider the fact that a web application is implemented as multiple program fragments. Each fragment is open as a service interface. An attacker could make up requests to these interfaces without strictly following the links provided in the server responses. Using the crafted requests, they could poison program states (*e.g.*, by modifying naïve implementations of client-side states), inject malformed input (*e.g.*, by exploiting insufficient input validation), or circumvent intended work flows (*e.g.*, by unexpected use of the “back” button).

Programmers need to be aware of all these issues and follow the relevant security practices. In the result code, business logic is intertwined with security manipulations. Consequently, secure web programming is difficult, and web programs are hard to maintain.

2.4 A Declarative Layer

We propose to program web applications using a high-level language that handles some common security aspects behind the scene. This language benefits from new abstractions designed for web programming. A program in this language more directly reflects the business logic of the target application; therefore, it is easier to write, to reason about, and to maintain. The language implementation (the compiler) will generate secure target code following established security practices.

Abstracting web interactions From earlier discussions, it is easy to see that much of the complication is due to the need of obtaining user input. Therefore, supporting web interactions is a key. We introduce a dedicated construct `form` for this purpose:

```
form(p : “username”, q : “password”);
```

The intention of the `form` construct is to present an HTML page to the client and obtain some input. In a realistic deployment, such a construct would take as argument a full-fledged HTML document, possibly generated dynamically and/or containing multiple component forms. In this report, we simply let it take arguments to describe input parameters of a single web form. In the above example, the `form` construct presents to the client a simple form comprised of two fields: `username` and `password`. After the client fills in the form and submits it, the `form` construct assigns the values provided in the fields to the two variables `p` and `q`.

There are a few issues handled transparently by the underlying implementation. First, the server program is split upon a `form` construct, and proper URL embedding is used in the web page to connect the control flow. Second, the input values of the variables are parsed from the form submission, and input validation is performed according to the declared variable types. Third, standard security practices are followed to manage sessions, maintain states, and defend against common exploits. The details of the implementation will be discussed in Section 5. For now, it suffices to understand this construct from a programmer's point of view as an abstract and secure web input operation which does not break the control flow.

Supporting user navigation Such a `form` construct implicitly opens up a service interface for receiving user requests. There would be vulnerabilities if it were not handled properly, or if its handling were not sufficiently understood by the programmer. Previous work (*sans* security) [4, 6, 9] on web interaction abstractions requires that the interface be “open” only once—a second request to the interface would not be accepted. This significantly restricts user navigation [15, 16, 28]. In practice, it is common for a user to return to a previous navigation stage using the “back” button. In general, the user could revisit any item in the browser history.

The validity of such an operation should be determined by the application, rather than be dismissed all together. For our banking example, it is reasonable to allow a user to duplicate the service selection page and proceed with both “balance” and “payment” in parallel.

For this purpose, we allow two modes of web interactions: a *single-use* mode (`forms`) and a *multi-use* mode (`formM`). In the single-use mode, the interface is open for request only once; revisiting the interface (*e.g.*, by activating a browser bookmark) results in an error. In the multi-use mode, the interface remains open for future requests. The semantics of BASS articulates the program behavior in both cases; therefore, the programmer can choose the suitable one based on the application. In either case, a request is accepted only if it follows the intended control flow of the server program to the interface. Consider our banking example. It is okay if the user reached the service selection page, bookmarked it, and reused it before logging out. However, it is not okay if the user forged a payment request without first going through the login and service selection pages.

Maintaining program states Multi-use forms are sufficient to accommodate all possible client navigation behaviors, because any behavior can be viewed as revisiting a point in the browser history. From a programmer's point of view, the program is virtually forked into multiple parallel “threads” at a multi-use `form` statement, together with all appropriate program state. The handling of the program state there is nontrivial. Some parts of the state could be local to the individual threads, whereas others could be global to all threads. The “local” state of a thread will not be affected by other threads, whereas the “global” state may be affected by others. Careless treatment of the state may result in logical errors [16].

The exact partitioning of the state should be determined by the application. We let programmers declare mutable variables as either *volatile* or *nonvolatile*. In the BASS implementation, volatile state can be stored in a database on the server across web interactions, thus all forked threads refer to the same volatile state. In contrast, nonvolatile state (after proper protection against client modifications) can be embedded in the URLs of web forms upon web interactions, thus every forked thread has its own nonvolatile state.

Manipulating client history Sometimes it is useful to disable some previously forked threads. Suppose the user tries to reload the service selection page after logging out of our banking application. The server program will receive a request that should not be processed. In general, we need a mechanism to disable some of the entries in the client history. In existing web applications, this is sometimes handled by embedding special tokens into web forms and checking them upon requests. While logging out, the server program expires the corresponding token, thus further requests to the service selection page will no longer be processed.

We do not wish to expose the details of token embedding to the programmer. Instead, we introduce a simple `clear` command for a similar purpose. From the programmer's point of view, `clear` simply resets the client history, so that all previously forked threads are discarded. This corresponds to the “session timeout” behaviors of many web applications. However, instead of thinking in terms of disabling the token of SID, BASS encourages programmers to think in terms of the client history. There are certainly more general ways to manipulate the client history. One possibility is to attach a privilege level to every web form, and to manipulate the forked threads by levels. Note that such level management can easily be implemented by the programmer using our provided abstractions of web input, volatile state, and nonvolatile state. More specifically, the current level can be stored in a volatile variable on the server as the global level. Upon a web interaction, the current level can be attached to the current thread using a nonvolatile variable as the local level. Comparison between the global and local levels and

```

string  user, pass, payee;
int     sel, amnt;

forms(user : "username", pass : "password");
LoginCheck(user, pass);
while (true) do {
  forms(sel : "1 : balance; 2 : payment; others : logout");
  if (sel == 1)
  then ShowBalance(user)
  else if (sel == 2)
  then {forms(payee : "payee"; amnt : "amount");
        DoPayment(user, payee, amnt)}
  else {clear; break}
}

```

Figure 3. Simple banking in BASS

proper handling can be done upon any user requests. Although not increasing the expressiveness of the language, abstractions for such level-based privilege management may prove convenient for certain applications. However, since its handling is straightforward, we omit it and use the simple `clear` command for demonstration.

Example revisited Figure 3 demonstrates the appeal of these abstractions by revisiting our banking example. The new abstractions provide an ideal world where there is only one client and the client is well behaved. In the code, we obtain login information from the client, perform login check (`LoginCheck` terminates the program if the check fails), and proceed with a service loop. In the loop, based on the service selection of the client, we carry out the balance service (`ShowBalance`) or the payment service (`DoPayment`), or log the user out. The service selection input is coded using a multi-use form, therefore the user may duplicate the corresponding web page and proceed with the two services in parallel. In addition, `clear` is used to disable all service threads when the user logs out. In this example, only the `user` variable is live across web interactions. Its value is obtained from a single-use form, and will not be updated after the login process. Therefore, it can be declared as either volatile or nonvolatile.

This code corresponds well to the work flow of Figure 1, and is much cleaner than a version directly written in an existing language. More importantly, it does not sacrifice security, because the implementation of the new abstractions will take care of the “plumbings” transparently—it will split the program into multiple fragments, maintain program states across web interactions, filter input based on variable types, and carry out relevant security manipulations such as the embedding of secret tokens.

Of course, these abstractions may not be suitable for all programming tasks. For applications which require little dynamic manipulation and state maintenance, it may not be as big a burden to directly code the security details. However, observing the growing popularity and complexity of web applications, especially in critical areas, we believe such abstractions will bring many benefits.

3. BASS

We formalize the above ideas in BASS, which provides a declarative view for programmers to write secure programs without worrying about the low-level details described in Section 2. The meta properties of BASS provide useful security guarantees. Assuming the semantics is enforced by an underlying implementation, all well-formed BASS programs will be secure in a certain sense. For example, values obtained from web input will have the expected types, forged requests from attackers will be rejected, and the program control flow will be enforced. These properties will be articulated in Section 3.4.




Figure 4. A virtual view of the execution environment

$$\begin{aligned}
(\text{World}) \quad W &::= (\Sigma, \kappa, \vec{\kappa}) \\
(\text{Global Env}) \quad \Sigma &::= \{\vec{a} : \vec{\tau} = \vec{v}\} \\
(\text{Closure}) \quad \kappa &::= (\sigma, C) \mid \sigma \\
(\text{Local Env}) \quad \sigma &::= \{\vec{x} : \vec{\tau} = \vec{v}\} \\
(\text{Type}) \quad \tau &::= \text{int} \mid \text{str} \\
(\text{Command}) \quad C &::= \text{skip} \mid a := E \mid x := E \mid C; C \\
&\quad \mid \text{if } E \text{ then } C \text{ else } C \mid \text{while } E \text{ do } C \\
&\quad \mid \text{forms}(\vec{a} : \vec{s}) \mid \text{form}_M(\vec{a} : \vec{s}) \mid \text{clear} \\
(E \text{xp}) \quad E &::= a \mid x \mid i \text{ (integers)} \mid s \text{ (strings)} \mid \text{op}(\vec{E}) \\
(V \text{alue}) \quad v &::= i \mid s
\end{aligned}$$

Figure 5. BASS syntax

3.1 Syntax

Figure 4 gives a graphical depiction of the virtual execution environment of BASS. We emphasize the word “virtual” because the view does not dictate how BASS is implemented; it simply provides a tractable world for the programmer to deal with.

In this world, there is a single server and a single client. The server consists of a global environment for volatile variables and a closure of the current execution. The closure is made up of a local environment for nonvolatile variables and some code. During execution, the server executes the code with respect to the two variable environments. Upon a web interaction, the server sends a proper closure to the client, and the client is then activated.

The client involves a list of closures reflecting the browsing history. When the client is active, it chooses an arbitrary closure in the history, fills in the corresponding web form, and sends the result to the server. We point out that although many code pieces appear in the closures on the client, they would actually be realized in an implementation as program pointers embedded in the forms. Similarly, the local environments can also be realized in forms.

Figure 5 gives the syntax. We use the vector $\vec{\kappa}$ to refer to the list $[\kappa_1 \dots \kappa_n]$. Similar notations apply to other meta variables, too.

A world W is a 3-tuple $(\Sigma, \kappa, \vec{\kappa})$. The first element Σ is the global variable environment, which is a mapping from volatile variables to their types and values. The second element κ is the closure currently active on the server. The third element $\vec{\kappa}$ is the client browsing history.

A closure is usually a pair of a local environment and a command. Sometimes, an environment without a command also makes a closure, which either contains useful information or signifies termination. A local environment maps nonvolatile variables to types and values. We assume that the names of volatile and nonvolatile variables are disjoint. By convention, we use a, a', a_1, \dots for volatile variables, and x, x', x_1, \dots for nonvolatile variables. We use two types: integers (`int`) and strings (`str`). Booleans are simulated with integers.

Commands are common except for web interactions (`forms`, `formM`) and history clearing (`clear`) as described in Section 2.4. For conciseness, web forms only input volatile variables; this does not affect expressiveness, because the input values can be passed on to nonvolatile variables with assignments. Expressions and values are as expected. We use `op` to abstract over all operations free from

side-effects, such as string concatenation and comparison. We omit function calls for conciseness, because they introduce no new issues.

3.2 Operational Semantics

We use a big-step semantics for expressions. Details are standard, thus omitted. We present in Figure 6 the small-step execution of a world. The key concept is a “world step” relation $W \rightsquigarrow W'$. A multi-step relation $W \rightsquigarrow^* W'$ is defined as the reflexive and transitive closure of the world step relation.

In world step, some rules reflect server-side computation, whereas others reflect client-side input. Rule (1) describes a server computation step by referring to a single-thread step relation $(\Sigma, \kappa) \rightsquigarrow (\Sigma', \kappa')$. This is applicable only when the current command starts with skip, assignment, conditional, or loop (the single-thread step relation is undefined on other cases). If the current command is `clear`, Rules (2) and (3) apply to remove all closures from the client. In these rules and the remainder of this report, we use ϵ to denote the empty vector. Upon any `form` commands, Rule (4) applies to transfer the closure to the client. The notation $[\kappa] \cup \vec{\kappa}$ means to combine the singleton vector $[\kappa]$ with $\vec{\kappa}$.

Client input through a single-use form is captured in Rules (5) and (6). Note that there is no further command to execute on the server (the closure on the server is simply σ_0), therefore the client takes control. In Rule (5), the client selects the closure κ_i to proceed with next. This corresponds to picking an arbitrary point in the browsing history. The client may input arbitrary values to the form, as long as they are of the right types ($\vec{v} \in \vec{\tau}$). The global environment is updated using the macro $update(\Sigma, \vec{a}, \vec{v})$, which updates Σ so that the variables \vec{a} will have the values \vec{v} . As a result of the step, the global environment is updated, and the chosen client closure is moved to the server for execution. The notation $\vec{\kappa} - [\kappa_i]$ means to remove κ_i from $\vec{\kappa}$. Rule (6) is a variant of Rule (5) where there is no further command in the chosen closure to execute after the form input. In this case, the chosen closure is simply removed.

Rules (7) and (8) are for client input through multi-use forms. The difference in comparison with the single-use ones lies in the result browsing history. For single-use forms, the selected closure is removed from the history so that it cannot be revisited. For multi-use forms, the history remains the same. All client request rules are non-deterministic on closure selection. A client may “abandon” a closure by never selecting it.

The single-thread step relation is straightforwardly described in Rules (9)–(16). Here the macro $update$ is overloaded to update both the global and the local environments.

3.3 Typing Rules

We give some non-standard typing rules in Figure 7. In Rule (17), a world $(\Sigma, \kappa, \vec{\kappa})$ is well-formed if all closures (κ and $\vec{\kappa}$) are well-formed with respect to the global environment (Σ), and every closure (κ_i) on the client side contains code which starts with a web form. The well-formedness of closures is defined in Rules (18) and (19). A trivial closure σ is always well-formed. A nontrivial closure (σ, C) is well-formed with respect to Σ if C is well-formed with respect to Σ and σ . We omit the standard handling on commands and expressions, where types of variables are pulled from the corresponding environments.

3.4 Meta Properties

BASS enjoys standard type soundness via preservation and progress.

Lemma 1 (Preservation) If $\vdash W$ and $W \rightsquigarrow W'$, then $\vdash W'$.

Lemma 2 (Progress) If $\vdash W$ then either there exists W' such that $W \rightsquigarrow W'$, or W is of the form $(\Sigma, \sigma, \epsilon)$.

$\Sigma; \sigma \vdash E \Downarrow v$	Omitted.
$W \rightsquigarrow^* W'$	Omitted.
$W \rightsquigarrow W'$	

Server computation:

$$\frac{(\Sigma, \kappa) \rightsquigarrow (\Sigma', \kappa')}{(\Sigma, \kappa, \vec{\kappa}) \rightsquigarrow (\Sigma', \kappa', \vec{\kappa})} \quad (1)$$

History clearing:

$$(\Sigma, (\sigma, \text{clear}; C), \vec{\kappa}) \rightsquigarrow (\Sigma, (\sigma, C), \epsilon) \quad (2)$$

$$(\Sigma, (\sigma, \text{clear}), \vec{\kappa}) \rightsquigarrow (\Sigma, \sigma, \epsilon) \quad (3)$$

Server response:

$$\frac{C \in \{(\text{form}_S(\vec{a}; \vec{s}); C'), \text{form}_S(\vec{a}; \vec{s}), (\text{form}_M(\vec{a}; \vec{s}); C'), \text{form}_M(\vec{a}; \vec{s})\}}{(\Sigma, (\sigma, C), \vec{\kappa}) \rightsquigarrow (\Sigma, \sigma, [(\sigma, C)] \cup \vec{\kappa})} \quad (4)$$

Client request:

$$\frac{\kappa_i \in \vec{\kappa} \quad \kappa_i = (\sigma, (\text{form}_S(\vec{a}; \vec{s}); C)) \quad (\vec{a} : \vec{\tau} = \dots) \in \Sigma \quad \vec{v} \in \vec{\tau} \quad update(\Sigma, \vec{a}, \vec{v}) = \Sigma'}{(\Sigma, \sigma_0, \vec{\kappa}) \rightsquigarrow (\Sigma', (\sigma, C), \vec{\kappa} - [\kappa_i])} \quad (5)$$

$$\frac{\kappa_i \in \vec{\kappa} \quad \kappa_i = (\sigma, (\text{form}_S(\vec{a}; \vec{s})))}{(\Sigma, \sigma_0, \vec{\kappa}) \rightsquigarrow (\Sigma, \sigma_0, \vec{\kappa} - [\kappa_i])} \quad (6)$$

$$\frac{\kappa_i \in \vec{\kappa} \quad \kappa_i = (\sigma, (\text{form}_M(\vec{a}; \vec{s}); C)) \quad (\vec{a} : \vec{\tau} = \dots) \in \Sigma \quad \vec{v} \in \vec{\tau} \quad update(\Sigma, \vec{a}, \vec{v}) = \Sigma'}{(\Sigma, \sigma_0, \vec{\kappa}) \rightsquigarrow (\Sigma', (\sigma, C), \vec{\kappa})} \quad (7)$$

$$\frac{\kappa_i \in \vec{\kappa} \quad \kappa_i = (\sigma, (\text{form}_M(\vec{a}; \vec{s})))}{(\Sigma, \sigma_0, \vec{\kappa}) \rightsquigarrow (\Sigma, \sigma_0, \vec{\kappa})} \quad (8)$$

$(\Sigma, \kappa) \rightsquigarrow (\Sigma', \kappa')$	
--	--

$$(\Sigma, (\sigma, \text{skip})) \rightsquigarrow (\Sigma, \sigma) \quad (9)$$

$$\frac{\Sigma; \sigma \vdash E \Downarrow v \quad update(\Sigma, a, v) = \Sigma'}{(\Sigma, (\sigma, a := E)) \rightsquigarrow (\Sigma', \sigma)} \quad (10)$$

$$\frac{\Sigma; \sigma \vdash E \Downarrow v \quad update(\sigma, x, v) = \sigma'}{(\Sigma, (\sigma, x := E)) \rightsquigarrow (\Sigma, \sigma')} \quad (11)$$

$$\frac{(\Sigma, (\sigma, C)) \rightsquigarrow (\Sigma', \sigma')}{(\Sigma, (\sigma, C; C')) \rightsquigarrow (\Sigma', (\sigma', C'))} \quad (12)$$

$$\frac{\Sigma; \sigma \vdash E \Downarrow i \quad i \neq 0}{(\Sigma, (\sigma, \text{if } E \text{ then } C \text{ else } C')) \rightsquigarrow (\Sigma, (\sigma, C))} \quad (13)$$

$$\frac{\Sigma; \sigma \vdash E \Downarrow 0}{(\Sigma, (\sigma, \text{if } E \text{ then } C \text{ else } C')) \rightsquigarrow (\Sigma, (\sigma, C'))} \quad (14)$$

$$\frac{\Sigma; \sigma \vdash E \Downarrow i \quad i \neq 0}{(\Sigma, (\sigma, \text{while } E \text{ do } C)) \rightsquigarrow (\Sigma, (\sigma, C; \text{while } E \text{ do } C))} \quad (15)$$

$$\frac{\Sigma; \sigma \vdash E \Downarrow 0}{(\Sigma, (\sigma, \text{while } E \text{ do } C)) \rightsquigarrow (\Sigma, \sigma)} \quad (16)$$

Figure 6. BASS operational semantics

$\vdash W$	$\Sigma \vdash \kappa \quad \forall \kappa_i \in \vec{\kappa}. \begin{cases} \kappa_i = (\sigma_i, C_i) \\ C_i \text{ starts with } \text{form}_S \text{ or } \text{form}_M \end{cases} \quad \Sigma \vdash \kappa_i$	(17)
$\Sigma \vdash \kappa$	$\Sigma \vdash \sigma$	(18)
$\Sigma; \sigma \vdash C$	$\Sigma \vdash (\sigma, C)$	(19)
$\Sigma; \sigma \vdash C$	Omitted.	
$\Sigma; \sigma \vdash E$	Omitted.	

Figure 7. BASS typing rules

As usual, these imply that well-formed programs will execute till termination without getting stuck. These seemingly simple theorems provide useful guarantees for server-side scripting. They essentially guarantee that the program execution will be confined by the operational semantics—only the transitions defined by the operational semantics may occur. Therefore, all behaviors reflected by the operational semantics are guaranteed, including:

1. The control flow will follow the explicit high-level structures;
2. Values obtained from web input will be well typed;
3. Program states will be kept intact across web interactions;
4. Entries in the client history will be disabled upon `clear`;
5. There will be only one client interacting with the program instance.

These free programmers from some low-level maneuvers, such as:

1. Embedding and enforcing the control flow;
2. Validating the names and types of web input parameters;
3. Recovering program states across web interactions;
4. Addressing unexpected user navigation behaviors;
5. Protecting against malicious exploits that involve multiple clients (*e.g.*, CSRF, first-order XSS, and session fixation).

We stress the point that the BASS approach to security is to directly preclude questionable behaviors from the operational semantics. BASS programmers enjoy a clean view of web programming, with common security handling transparently taken care of by the BASS implementation (the compiler). This enables programmers to focus on the business logic of specific web applications, thus improving security and productivity. Specifically on security, the BASS semantics leaves no room for CSRF, first-order XSS, session fixation, session poisoning, malformed input, and work-flow circumvention. Of course, the semantics must be enforced by an implementation. This is the topic of the next two sections, where we formalize a realistic web programming model MOSS, and translate BASS into MOSS so that the BASS semantics is faithfully carried out.

4. MOSS

MOSS articulates server program execution and client behaviors. It reflects realistic web interactions, and can be used to illustrate both well-behaved client activities and malicious exploits. We will explain how concepts in MOSS correspond to real-world entities throughout the section.

4.1 Syntax

A graphical depiction of the model is given in Figure 8. On the server, there are a group of programs collaborating to implement the




Figure 8. A model of web programming

	(World) $W ::= (\dot{\Sigma}, \dot{\kappa}, \vec{\Phi})$	
(Global Env)	$\dot{\Sigma} ::= \{[\vec{l} : \vec{c}], [f = (\vec{a}) \triangleright C]^*\}$	
(Session Env)	$\varsigma ::= \{\vec{a} = \vec{v}\}$	
(Closure)	$\dot{\kappa} ::= (i, v, \dot{\sigma}, C) \mid (v, \dot{\sigma})$	
(Local Env)	$\dot{\sigma} ::= \{\vec{x} = \vec{v}\}$	
(Client)	$\Phi ::= (v, \vec{\phi})$	
(Form)	$\phi ::= \text{form } f(\vec{s} \vec{a} = \vec{v}) \text{ with } \psi$	
(Sealed)	$\psi ::= \langle \vec{x} = \vec{v} \rangle$	
(Command)	$C ::= \text{sk} \mid a \leftarrow E \mid C; C \mid \text{cond}(E, C, C) \mid \text{loop}(E, C) \mid f(\vec{E}) \mid \text{input}(E) \mid \text{setCk}(E) \mid \text{unpack}(E) \mid \text{verif}(\vec{E}, \vec{\tau}) \mid \text{err}$	
(Exp)	$E ::= a \mid x \mid i \text{ (integers)} \mid s \text{ (strings)} \mid \text{op}(\vec{E}) \mid \{\vec{E}\} \mid \text{in}(E, E) \mid \text{ins}(E, E) \mid \text{del}(E, E) \mid \text{form } f(\vec{s} \vec{a} = \vec{v}) \text{ with } E \mid \text{getCk}() \mid \psi \mid \text{pack}() \mid \iota \text{ (tokens)} \mid \text{newTk}()$	
(Value)	$v, w ::= i \mid s \mid \{\vec{v}\} \mid \phi \mid \psi \mid \iota$	
(BASS Type)	$\tau ::= \text{int} \mid \text{str}$	

Figure 9. MOSS syntax

desired application logic. They access and update data that belong to multiple clients. The server interacts with multiple clients simultaneously, although in separate logical transactions. Every client has its own cookie, and maintains a list of forms as the browsing history. Although forms are expected to be obtained from the server through web interactions, a client may also forge a form on its own (as illustrated using Form_x and Form_y , which appear to come from nowhere). This reflects real-world exploits where a client makes up an HTTP request without going through the intended application logic. In addition, a malicious client B may trick an innocent client A by inserting a form from B's history (part of which may be forged) into A's history (as illustrated using Form_z). In the real-world counterpart, an attacker may send crafted links to a victim in an email; sometimes, these links are image links that may be automatically loaded by a browser.

Whereas the virtual execution environment in Figure 4 provides an ideal view for the programmer to work with, the model in Figure 8 reflects real-world web application scenarios, where server programs interact with and maintain data for multiple clients, who are no longer always well-behaved. To implement secure web applications in this model, the server programs need special primitives for manipulating security concepts and constructs. For example, special tokens (*e.g.*, SID) can be used to identify users and logical transactions, and those tokens can be embedded in the cookies and forms on the clients.

We present the details of this model as a low-level language MOSS. The syntax is given in Figure 9. MOSS is untyped for conciseness, but it is straightforward to give it a standard typing. Vector notations (*e.g.*, $\vec{\Phi}$) are sometimes used when referring to a

list of items (e.g., $[\Phi_1 \dots \Phi_n]$). The symbols for certain language constructs are borrowed from BASS to show correspondence of the underlying concepts, but minor modifications are applied to avoid confusion. In particular, we use a different font for reused English letters (\mathbb{W} , \mathbb{C} , \mathbb{E} and \mathbb{v}), and add a dot on top for reused Greek letters ($\vec{\Sigma}$, $\dot{\kappa}$, and $\dot{\sigma}$).

A world \mathbb{W} in this language consists of three elements. The first is a global environment $\vec{\Sigma}$ on the server. The second is a closure $\dot{\kappa}$ currently active on the server. The third is a list of browsing histories $\vec{\Phi}$, each belonging to a different client.

The global environment $\vec{\Sigma}$ contains both data and code. The data part is essentially a list of session environments ς organized by some tokens ι for identifying sessions (i.e., SIDs). Every session environment ς captures the values of the volatile variables (e.g., a, a', a_1) for a session. The code part is simply a collection of functions, with every function consisting of a name f , a list of input parameters \vec{a} , and a body command \mathbb{C} . Here, the symbols $\llbracket \cdot \rrbracket$ are used as parentheses of the meta language, rather than as part of MOSS.

The closure $\dot{\kappa}$ is usually a tuple consisting of a client identifier i (indicating that the current computation is for the i th client; this corresponds to the transient but dedicated connection established between the client and the server for an HTTP request-response), a value \mathbb{v} as the content of the cookie (transferred from the client to the server in a request and passed back to the client in a response), a local environment $\dot{\sigma}$, and a command \mathbb{C} . Sometimes a pair of a cookie value and a local environment also makes a closure; such a closure cannot be executed, but the information therein may be propagated to other closures. The local environment $\dot{\sigma}$ collects the values of non-volatile variables (e.g., x, x', x_1). We assume that $\dot{\sigma}$ contains a special variable x_{sid} dedicated for storing the SID of the underlying session. Alternatively, one may wish to store the SID in the cookie. Since the cookie is shared by all forms on a client, that would prevent the same client from accessing multiple sessions simultaneously. Therefore, we choose to store the SID in the local environments of the individual “threads.”

The clients vector $\vec{\Phi}$ deserves attention. In BASS, the single client is idealized as a list of closures. In MOSS, a world involves multiple clients. Each client Φ consists of a list of forms $\vec{\phi}$ and a cookie with the value \mathbb{v} . A form ϕ is made up of input parameters \vec{a} and their textual explanations \vec{s} and default values \vec{v} (i.e., input fields in a web form), a target f (i.e., the web link behind the “submit” button), and some sealed client-side state ψ . Note that ψ is not meant to be modified by the client. This restriction will be enforced in the operational semantics. In the syntax, we use the notation $\langle \dots \rangle$ to indicate that it is a “package” that cannot be taken apart by the client. In a real language, such sealed state can be signed by the server to prevent client modifications. We also point out that the forms in a client may belong to different sessions, because the client may initiate different sessions in different browser windows. Although reserving a dedicated variable x_{sid} in the local environments for session identification, MOSS does not automatically maintain it (or any other non-volatile variables) across web interactions. As in the case of real web programming, the programmer may maintain the relevant information in the sealed state.

Commands \mathbb{C} involve the common ones and some new primitives adapted from real-world entities for secure server-side scripting. For distinction from BASS, we use a different syntax for the common commands of skip (sk), assignment ($a \leftarrow E$), sequential composition ($\mathbb{C}; \mathbb{C}$), conditional ($\text{cond}(E, \mathbb{C}, \mathbb{C})$), and while-loop ($\text{loop}(E, \mathbb{C})$). $f(\vec{E})$ is a function call to f with the arguments \vec{E} . When the server program needs web input, $\text{input}(E)$ is used to send a form computed from E to the client. $\text{setCk}(E)$ sets a new value for the cookie in the current closure. $\text{unpack}(E)$ unpacks the sealed state from E . $\text{verify}(\vec{E}, \vec{\tau})$ validates \vec{E} against BASS types $\vec{\tau}$. Finally, err exits the program upon an error; we use this abrupt pro-

$$\boxed{\varsigma; \mathbb{v}_{ck}; \dot{\sigma} \vdash E \Downarrow \mathbb{v}} \quad \varsigma; \mathbb{v}_{ck}; \dot{\sigma} \vdash a \Downarrow \varsigma(a) \quad (20)$$

$$\varsigma; \mathbb{v}_{ck}; \dot{\sigma} \vdash x \Downarrow \dot{\sigma}(x) \quad (21)$$

$$\varsigma; \mathbb{v}_{ck}; \dot{\sigma} \vdash \mathbb{v} \Downarrow \mathbb{v} \quad (22)$$

$$\frac{\varsigma; \mathbb{v}_{ck}; \dot{\sigma} \vdash \vec{E} \Downarrow \vec{v}}{\varsigma; \mathbb{v}_{ck}; \dot{\sigma} \vdash \text{op}(\vec{E}) \Downarrow \hat{op}(\vec{v})} \quad (23)$$

$$\frac{\varsigma; \mathbb{v}_{ck}; \dot{\sigma} \vdash \vec{E} \Downarrow \vec{v}}{\varsigma; \mathbb{v}_{ck}; \dot{\sigma} \vdash \{\vec{E}\} \Downarrow \{\vec{v}\}} \quad (24)$$

$$\frac{\varsigma; \mathbb{v}_{ck}; \dot{\sigma} \vdash E \Downarrow \mathbb{v} \quad \varsigma; \mathbb{v}_{ck}; \dot{\sigma} \vdash E' \Downarrow \{\vec{v}\}}{\varsigma; \mathbb{v}_{ck}; \dot{\sigma} \vdash \text{in}(E, E') \Downarrow i} \quad i = \begin{cases} 0 & \text{if } \mathbb{v} \notin \vec{v} \\ 1 & \text{if } \mathbb{v} \in \vec{v} \end{cases} \quad (25)$$

$$\frac{\varsigma; \mathbb{v}_{ck}; \dot{\sigma} \vdash E \Downarrow \mathbb{v} \quad \varsigma; \mathbb{v}_{ck}; \dot{\sigma} \vdash E' \Downarrow \{\vec{v}\}}{\varsigma; \mathbb{v}_{ck}; \dot{\sigma} \vdash \text{ins}(E, E') \Downarrow \langle \vec{v}, \mathbb{v} \rangle} \quad (26)$$

$$\frac{\varsigma; \mathbb{v}_{ck}; \dot{\sigma} \vdash E \Downarrow \mathbb{v} \quad \varsigma; \mathbb{v}_{ck}; \dot{\sigma} \vdash E' \Downarrow \{\vec{v}\} \quad \vec{w} = \vec{v} - \{\mathbb{v}\}}{\varsigma; \mathbb{v}_{ck}; \dot{\sigma} \vdash \text{del}(E, E') \Downarrow \langle \vec{w} \rangle} \quad (27)$$

$$\frac{\varsigma; \mathbb{v}_{ck}; \dot{\sigma} \vdash E \Downarrow \mathbb{v}}{\varsigma; \mathbb{v}_{ck}; \dot{\sigma} \vdash \text{form}(f(\vec{s} \vec{a} = \vec{\tau})) \text{ with } E \Downarrow \text{form } f(\vec{s} \vec{a} = \vec{\tau}) \text{ with } \mathbb{v}} \quad (28)$$

$$\varsigma; \mathbb{v}_{ck}; \dot{\sigma} \vdash \text{getCk}() \Downarrow \mathbb{v}_{ck} \quad (29)$$

$$\varsigma; \mathbb{v}_{ck}; \langle \vec{x} = \vec{v} \rangle \vdash \text{pack}() \Downarrow \langle \vec{x} = \vec{v} \rangle \quad (30)$$

$$\frac{\text{fresh}(\iota)}{\varsigma; \mathbb{v}_{ck}; \dot{\sigma} \vdash \text{newTk}() \Downarrow \iota} \quad (31)$$

Figure 10. MOSS expression evaluation

gram termination for simplicity, although real web applications typically perform more sophisticated error handling (e.g., giving error messages, redirecting to other pages). The exact meanings of these commands will be defined by the operational semantics.

MOSS reuses the variables a and x , integers i , strings s , and abstract operations $\text{op}(\vec{E})$ of BASS for ease of exposition of the translation in Section 5. It also introduces some set expressions: $\{\vec{E}\}$ is a set with elements \vec{E} , $\text{in}(E, E')$ checks if E belongs to the set E' , $\text{ins}(E, E')$ inserts E into the set E' , and $\text{del}(E, E')$ removes E from the set E' . form constructs a form by putting together some input parameters, a target function, and an expression carrying the sealed state. $\text{getCk}()$ obtains the cookie value. ψ is a sealed state, and it is typically created using $\text{pack}()$. Finally, tokens ι are unique and unforgeable entities created using $\text{newTk}()$; $\text{newTk}()$ will never create two identical tokens, and the tokens created cannot be guessed or forged by clients. The exact form of tokens is abstract in MOSS. We have already mentioned the use of tokens as SIDs. Later when translating BASS, we will also use tokens to identify clients, enforce control flows, and prevent request forgeries.

4.2 Operational Semantics

Figure 10 gives the semantics of expressions. The evaluation is carried out with respect to three entities: a session environment ς , a cookie value \mathbb{v}_{ck} , and a local environment $\dot{\sigma}$. The values for volatile and nonvolatile variables are pulled from their environments respectively in Rules (20) and (21). Values require no further evaluation (Rule (22)). In Rule (23), we use \hat{op} to refer to the corresponding meta-level computation of op . The set expressions are evaluated as expected in Rules (24)–(27), where integers are used to simulate booleans. Rule (28) evaluates the expression component of a form, which is expected to be a sealed state, and composes a form value. Rule (29) simply provides the cookie as the result of evalu-

<div style="border: 1px solid black; padding: 2px; display: inline-block; margin-bottom: 5px;">$W \rightsquigarrow^* W'$</div> <p style="margin-left: 40px;">Omitted.</p> <div style="border: 1px solid black; padding: 2px; display: inline-block; margin-bottom: 5px;">$W \rightsquigarrow W'$</div> <p>Server computation:</p> $\frac{(\dot{\Sigma}, \dot{\kappa}) \rightsquigarrow (\dot{\Sigma}', \dot{\kappa}')}{(\dot{\Sigma}, \dot{\kappa}, \vec{\Phi}) \rightsquigarrow (\dot{\Sigma}', \dot{\kappa}', \vec{\Phi})} \quad (32)$ <p>Server response:</p> $\frac{\begin{array}{l} \dot{\kappa} = (i, \mathbf{v}_{ck}, \dot{\sigma}, \mathbf{input}(\mathbf{E})) \text{ or } \dot{\kappa} = (i, \mathbf{v}_{ck}, \dot{\sigma}, \mathbf{input}(\mathbf{E})) ; \mathbf{C} \\ \dot{\sigma}(x_{sid}) = \iota \quad \dot{\Sigma}(\iota); \mathbf{v}_{ck}; \dot{\sigma} \vdash \mathbf{E} \Downarrow \phi \\ \Phi_i = (\mathbf{v}, \vec{\phi}) \quad \Phi'_i = (\mathbf{v}_{ck}, [\phi] \cup \vec{\phi}) \end{array}}{(\dot{\Sigma}, \dot{\kappa}, \vec{\Phi}) \rightsquigarrow (\dot{\Sigma}, (\mathbf{v}_{ck}, \dot{\sigma}), (\vec{\Phi} - [\Phi_i]) \cup [\Phi'_i])} \quad (33)$ <p>Client forging (form filling):</p> $\frac{\begin{array}{l} \Phi_i \in \vec{\Phi} \quad \Phi_i = (\mathbf{v}_{ck}, \vec{\phi}) \\ \phi \in \vec{\phi} \quad \phi = \mathbf{form} f(\vec{s} \vec{a} = \vec{v}) \text{ with } \psi \\ \phi' = \mathbf{form} f'(\vec{s}' \vec{a}' = \vec{w}) \text{ with } \psi \quad \Phi'_i = (\mathbf{v}', [\phi'] \cup \vec{\phi}) \end{array}}{(\dot{\Sigma}, (\mathbf{v}, \dot{\sigma}), \vec{\Phi}) \rightsquigarrow (\dot{\Sigma}, (\mathbf{v}, \dot{\sigma}), (\vec{\Phi} - [\Phi_i]) \cup [\Phi'_i])} \quad (34)$ <p>Client request:</p> $\frac{\begin{array}{l} \Phi_i \in \vec{\Phi} \quad \Phi_i = (\mathbf{v}_{ck}, \vec{\phi}) \quad \mathbf{form} f(\vec{s} \vec{a} = \vec{v}) \text{ with } \psi \in \vec{\phi} \\ (f = (a_{srl}, \vec{a}) \triangleright \mathbf{C}) \in \dot{\Sigma} \end{array}}{(\dot{\Sigma}, (\mathbf{v}, \dot{\sigma}), \vec{\Phi}) \rightsquigarrow (\dot{\Sigma}, (i, \mathbf{v}_{ck}, \{\}, f(\psi, \vec{v})), \vec{\Phi})} \quad (35)$ <p>Attacker tricking:</p> $\frac{\begin{array}{l} \Phi_i \in \vec{\Phi} \quad \Phi_i = (\mathbf{v}', \vec{\phi}) \quad \phi \in \vec{\phi} \\ \Phi_j \in \vec{\Phi} \quad \Phi_j = (\mathbf{v}_{ck}, \vec{\phi}') \quad \Phi'_j = (\mathbf{v}_{ck}, [\phi] \cup \vec{\phi}') \end{array}}{(\dot{\Sigma}, (\mathbf{v}, \dot{\sigma}), \vec{\Phi}) \rightsquigarrow (\dot{\Sigma}, (\mathbf{v}, \dot{\sigma}), (\vec{\Phi} - [\Phi_j]) \cup [\Phi'_j])} \quad (36)$ <p>New session:</p> $\frac{\begin{array}{l} \Phi_i \in \vec{\Phi} \quad \Phi_i = (\mathbf{v}_{ck}, \vec{\phi}) \quad \mathbf{fresh}(\iota) \\ \dot{\Sigma}' = \dot{\Sigma} \cup \{\iota : \{\}\} \quad \dot{\kappa} = (i, \mathbf{v}_{ck}, \{x_{sid} = \iota\}, f_{start}()) \end{array}}{(\dot{\Sigma}, (\mathbf{v}, \dot{\sigma}), \vec{\Phi}) \rightsquigarrow (\dot{\Sigma}', \dot{\kappa}, \vec{\Phi})} \quad (37)$ <p>New client:</p> $\frac{\begin{array}{l} \mathbf{fresh}(\iota) \quad \dot{\Sigma}' = \dot{\Sigma} \cup \{\iota : \{\}\} \quad \vec{\Phi} = n \\ \mathbf{fresh}(\iota') \quad \dot{\kappa} = (n+1, \iota', \{x_{sid} = \iota\}, f_{start}()) \end{array}}{(\dot{\Sigma}, (\mathbf{v}, \dot{\sigma}), \vec{\Phi}) \rightsquigarrow (\dot{\Sigma}', \dot{\kappa}, \vec{\Phi} \cup [(\bullet, \epsilon)])} \quad (38)$	<div style="border: 1px solid black; padding: 2px; display: inline-block; margin-bottom: 5px;">$(\Sigma, \kappa) \rightsquigarrow (\Sigma', \kappa')$</div> $(\dot{\Sigma}, (i, \mathbf{v}_{ck}, \dot{\sigma}, \mathbf{sk})) \rightsquigarrow (\dot{\Sigma}, (\mathbf{v}_{ck}, \dot{\sigma})) \quad (39)$ $\frac{\begin{array}{l} \dot{\sigma}(x_{sid}) = \iota \quad \dot{\Sigma}(\iota); \mathbf{v}_{ck}; \dot{\sigma} \vdash \mathbf{E} \Downarrow \mathbf{v} \\ \mathbf{update}(\dot{\Sigma}(\iota), a, \mathbf{v}) = \dot{\Sigma}' \end{array}}{(\dot{\Sigma}, (i, \mathbf{v}_{ck}, \dot{\sigma}, a \leftarrow \mathbf{E})) \rightsquigarrow (\dot{\Sigma}', (\mathbf{v}_{ck}, \dot{\sigma}))} \quad (40)$ $\frac{\begin{array}{l} \dot{\sigma}(x_{sid}) = \iota \quad \dot{\Sigma}(\iota); \mathbf{v}_{ck}; \dot{\sigma} \vdash \mathbf{E} \Downarrow \mathbf{v} \\ \mathbf{update}(\dot{\sigma}, x, \mathbf{v}) = \dot{\sigma}' \end{array}}{(\dot{\Sigma}, (i, \mathbf{v}_{ck}, \dot{\sigma}, x \leftarrow \mathbf{E})) \rightsquigarrow (\dot{\Sigma}, (\mathbf{v}_{ck}, \dot{\sigma}'))} \quad (41)$ $\frac{(\dot{\Sigma}, (i, \mathbf{v}_{ck}, \dot{\sigma}, \mathbf{C})) \rightsquigarrow (\dot{\Sigma}', (\mathbf{v}, \dot{\sigma}'))}{(\dot{\Sigma}, (i, \mathbf{v}_{ck}, \dot{\sigma}, \mathbf{C};; \mathbf{C}')) \rightsquigarrow (\dot{\Sigma}', (i, \mathbf{v}, \dot{\sigma}', \mathbf{C}'))} \quad (42)$ $\frac{\begin{array}{l} \dot{\sigma}(x_{sid}) = \iota \quad \dot{\Sigma}(\iota); \mathbf{v}_{ck}; \dot{\sigma} \vdash \mathbf{E} \Downarrow i \quad i \neq 0 \end{array}}{(\dot{\Sigma}, (i, \mathbf{v}_{ck}, \dot{\sigma}, \mathbf{cond}(\mathbf{E}, \mathbf{C}, \mathbf{C}')) \rightsquigarrow (\dot{\Sigma}, (i, \mathbf{v}_{ck}, \dot{\sigma}, \mathbf{C}))} \quad (43)$ $\frac{\begin{array}{l} \dot{\sigma}(x_{sid}) = \iota \quad \dot{\Sigma}(\iota); \mathbf{v}_{ck}; \dot{\sigma} \vdash \mathbf{E} \Downarrow 0 \end{array}}{(\dot{\Sigma}, (i, \mathbf{v}_{ck}, \dot{\sigma}, \mathbf{cond}(\mathbf{E}, \mathbf{C}, \mathbf{C}')) \rightsquigarrow (\dot{\Sigma}, (i, \mathbf{v}_{ck}, \dot{\sigma}, \mathbf{C}'))} \quad (44)$ $\frac{\begin{array}{l} \dot{\sigma}(x_{sid}) = \iota \quad \dot{\Sigma}(\iota); \mathbf{v}_{ck}; \dot{\sigma} \vdash \mathbf{E} \Downarrow i \quad i \neq 0 \end{array}}{(\dot{\Sigma}, (i, \mathbf{v}_{ck}, \dot{\sigma}, \mathbf{loop}(\mathbf{E}, \mathbf{C})) \rightsquigarrow (\dot{\Sigma}, (i, \mathbf{v}_{ck}, \dot{\sigma}, \mathbf{C};; \mathbf{loop}(\mathbf{E}, \mathbf{C})))} \quad (45)$ $\frac{\begin{array}{l} \dot{\sigma}(x_{sid}) = \iota \quad \dot{\Sigma}(\iota); \mathbf{v}_{ck}; \dot{\sigma} \vdash \mathbf{E} \Downarrow 0 \end{array}}{(\dot{\Sigma}, (i, \mathbf{v}_{ck}, \dot{\sigma}, \mathbf{loop}(\mathbf{E}, \mathbf{C})) \rightsquigarrow (\dot{\Sigma}, (\mathbf{v}_{ck}, \dot{\sigma}))} \quad (46)$ $\frac{\begin{array}{l} (f = (\vec{a}) \triangleright \mathbf{C}) \in \dot{\Sigma} \quad \vec{\mathbf{E}} = \vec{a} \\ \dot{\sigma}(x_{sid}) = \iota \quad \dot{\Sigma}(\iota); \mathbf{v}_{ck}; \dot{\sigma} \vdash \vec{\mathbf{E}} \Downarrow \vec{v} \end{array}}{(\dot{\Sigma}, (i, \mathbf{v}_{ck}, \dot{\sigma}, f(\vec{\mathbf{E}})) \rightsquigarrow (\dot{\Sigma}, (i, \mathbf{v}_{ck}, \dot{\sigma}, \mathbf{C}[\vec{v}/\vec{a}]))} \quad (47)$ $\frac{\begin{array}{l} \dot{\sigma}(x_{sid}) = \iota \quad \dot{\Sigma}(\iota); \mathbf{v}_{ck}; \dot{\sigma} \vdash \mathbf{E} \Downarrow \mathbf{v} \end{array}}{(\dot{\Sigma}, (i, \mathbf{v}_{ck}, \dot{\sigma}, \mathbf{setCk}(\mathbf{E})) \rightsquigarrow (\dot{\Sigma}, (\mathbf{v}, \dot{\sigma}))} \quad (48)$ $\frac{\begin{array}{l} \dot{\sigma}(x_{sid}) = \iota \quad \dot{\Sigma}(\iota); \mathbf{v}_{ck}; \dot{\sigma} \vdash \mathbf{E} \Downarrow \langle \vec{x} = \vec{v} \rangle \\ \mathbf{update}(\dot{\sigma}, \vec{x}, \vec{v}) = \dot{\sigma}' \end{array}}{(\dot{\Sigma}, (i, \mathbf{v}_{ck}, \dot{\sigma}, \mathbf{unpack}(\mathbf{E})) \rightsquigarrow (\dot{\Sigma}, (\mathbf{v}_{ck}, \dot{\sigma}'))} \quad (49)$ $\frac{\begin{array}{l} \dot{\sigma}(x_{sid}) = \iota \quad \dot{\Sigma}(\iota); \mathbf{v}_{ck}; \dot{\sigma} \vdash \vec{\mathbf{E}} \Downarrow \vec{v} \\ \forall k. (\exists i. \mathbf{v}_k = i \text{ and } \tau_k = \mathbf{int}) \text{ or } (\exists s. \mathbf{v}_k = s \text{ and } \tau_k = \mathbf{str}) \end{array}}{(\dot{\Sigma}, (i, \mathbf{v}_{ck}, \dot{\sigma}, \mathbf{verif}(\vec{\mathbf{E}}, \vec{\tau})) \rightsquigarrow (\dot{\Sigma}, (\mathbf{v}_{ck}, \dot{\sigma}))} \quad (50)$
--	---

Figure 11. MOSS world execution

ating `getCk()`. Rule (30) creates a sealed state based on the local environment. Finally, Rule (31) produces a fresh token.

The remainder of the operational semantics is given in Figure 11. Similar to BASS, the key concept is a “world step” relation $W \rightsquigarrow W'$. A multi-step relation $W \rightsquigarrow^* W'$ is defined as the reflexive and transitive closure of the world step relation.

Most of the server-side computation is captured in Rule (32), which delegates the task to a single-thread step relation $(\dot{\Sigma}, \dot{\kappa}) \rightsquigarrow (\dot{\Sigma}', \dot{\kappa}')$. Upon an `input(E)` command (Rule (33)), the expression \mathbf{E} is evaluated to a form ϕ , and the result is transferred to the client. Note that the special variable x_{sid} is used to obtain the SID ι , and ι is used to locate the session environment ($\dot{\Sigma}(\iota)$ refers to the session environment identified by ι in $\dot{\Sigma}$). In addition, the client identifier i is used to locate the client history Φ_i . Φ_i may contain a different cookie value \mathbf{v} , because the server may have updated the cookie to \mathbf{v}_{ck} in the current closure $\dot{\kappa}$. We update the client history from Φ_i to Φ'_i using the new cookie \mathbf{v}_{ck} and form ϕ . In the result of the step, the clients vector is updated by removing the old Φ_i and inserting

the updated Φ'_i . Any command \mathbf{C} following `input(E)` is discarded. This reflects the case that the server execution stops after sending a response to the client, and resumes after receiving the next request; the control flow is connected by a link embedded in the web form \mathbf{E} , not by sequencing with another command \mathbf{C} . This Rule (33) is the only one that describes the execution upon an `input(E)` command. Therefore, the execution of a MOSS program gets stuck if \mathbf{E} is not a form, if ι is not recognized by $\dot{\Sigma}$, or if x_{sid} does not contain a valid token at all. Similar observations apply to all other rules.

The remainder rules of the world step relation concern things initiated by a client. An ideal case is that a client fills in a form with values. However, in the real-world handling of a web form, the target function name and input fields are subject to malicious modifications, because they are directly embedded behind the “submit” button for the browser to recognize and produce an HTTP request. Such a scenario is captured in Rule (34), where a client Φ_i forges a new form ϕ' (the ideal form filling is a special case of this). Here, the client may make up arbitrary things on the function name and

input fields, but not the sealed state ψ (e.g., the client cannot forge the server’s signature); ψ must come from some form ϕ in the history of Φ_i . In addition, the cookie value may also be modified by the client from v_{ck} to v' . As a result of the step, the new client Φ'_i includes the forged form ϕ' and cookie v' , and the clients vector $\vec{\Phi}$ is updated accordingly. The server side remains unchanged.

In Rule (35), the i th client picks a form from the history and submits it. On the server side, the function f must be recognized by the global environment, and the function must expect the proper number of arguments (a_{srl} is a special argument for the sealed state). In the result of the step, a new closure is composed for execution, where the cookie is obtained from the client, and the function f is applied to the sealed state and the input values. No input validation on the argument types is done in this rule. Therefore, a program which does not perform proper input validation on its own may go wrong in later computation.

Sometimes an attacker tricks a victim into sending a request composed by the attacker. In Rule (36), the attacker client i picks an arbitrary form ϕ from Φ_i , and inserts it into the victim client j ’s history. The server side remains unchanged. Now that the form is injected into the victim client, a later request following Rule (35) may submit it to the server, thus completing the attack.

An existing client may initiate a new session using Rule (37). On the server side, a fresh token ι is created as the SID, an empty session environment is added into the global environment $\vec{\Sigma}$, and a new closure κ is composed. In the closure κ , the local environment is initialized so that x_{sid} is bound to ι , and the code starts from f_{start} , which is a special function reserved as the entry point of MOSS programs. The session environment and the local environment do not contain entries for programmer-defined variables yet. Instead, those variables will be declared implicitly upon their first use, which is articulated in later rules for executing assignments.

The last client-initiated step, Rule (38), applies when a new client joins. Similar to Rule (37), a new session with SID ι is created. In addition, another fresh token ι' is created as an ID for uniquely identifying the client (we will refer to this as a CID). ι' is stored in the cookie of the new closure κ ,¹ and will be propagated to the client upon a web interaction using Rule (33). As a result of the step, the global environment is updated, the new closure is activated for execution, and the clients vector is updated to include the new client, which has the transient client identifier $n+1$, an empty cookie indicated by the special value \bullet , and an empty history vector ϵ .

The right-hand side of Figure 11 defines the single-thread step relation $(\vec{\Sigma}, \kappa) \rightsquigarrow (\vec{\Sigma}', \kappa')$. The cases for skip (Rule (39)), sequencing (Rule (42)), conditional (Rules (43) and (44)), and while-loop (Rules (45) and (46)) are standard. Assignments are defined in Rules (40) and (41), where a macro *update* is overloaded to update both session environments ς and local environments $\dot{\sigma}$. If the variable to be updated is not in the environment, the macro will create a new entry, thus the variable is implicitly declared upon first use.

Rule (47) carries out a function call using capture-avoiding substitution, after checking the number of and evaluating the arguments. Rule (48) updates the cookie of the current closure. Rule (49) evaluates E to a sealed state, and updates the local environment $\dot{\sigma}$ accordingly. Finally, Rule (50) validates \vec{E} against BASS types $\vec{\tau}$. There is no rule for the case where the validation fails—the execution gets stuck upon a failed validation.

¹ We store CID in cookie automatically for a simpler exposition of the translation in Section 5. This does not restrict the expressiveness of MOSS, because the semantics does not inspect the cookie for CID, and a programmer may always choose to ignore the CID and use the cookie for other purposes.

4.3 Example Use: Patterns of Exploits

MOSS is designed to be very flexible, and some syntactically correct programs could result in stucks during execution. Some causes of stucks include unrecognized SIDs or function names, unmatched number of function parameters, illegal usage of values (e.g., unpacking an integer), and failed input validations. Indeed, MOSS is meant to reflect real-world scenarios, rather than to confine web program behaviors using a restricted semantics or a type system. Similar to existing web programming languages, MOSS can be used to write both secure programs and vulnerable ones.

Besides serving as the target of our BASS translation, MOSS is of independent value to understanding web programming, because it helps to articulate the essence of common exploits and attacks. We demonstrate this by encoding two exploit patterns, which cover all the example attacks discussed in Section 2.3.

The first is an *attacker-victim* pattern, where an attacker takes advantage by tricking a victim into submitting an unintended request, which in turn triggers a vulnerability in a server program. We assume the application stores the SID in the client cookie.

$$\begin{aligned} & (\vec{\Sigma}, (v, \dot{\sigma}), [(\iota_A, \vec{\phi}_A), (\iota_B, \vec{\phi}_B)]) \\ \rightsquigarrow & (\vec{\Sigma}, (v, \dot{\sigma}), [(\iota_A, \vec{\phi}_A), (\iota_B, \vec{\phi}_B \cup [\phi])]) && \text{(forge by Rule (34))} \\ \rightsquigarrow & (\vec{\Sigma}, (v, \dot{\sigma}), [(\iota_A, \vec{\phi}_A \cup [\phi]), (\iota_B, \vec{\phi}_B \cup [\phi])]) && \text{(trick by Rule (36))} \\ \rightsquigarrow & (\vec{\Sigma}, \kappa, [(\iota_A, \vec{\phi}_A \cup [\phi]), (\iota_B, \vec{\phi}_B \cup [\phi])]) && \text{(request by Rule (35))} \end{aligned}$$

Here the attacker B forges a form ϕ in the first step, and injects it into A (e.g., by emailing a link) in the second step. When A submits the form (e.g., by following the link) in the third step, the server program composes a closure κ to process the request. Whereas a careful server program may inspect ϕ to identify injected requests (Section 5), a vulnerable one might simply identify the session based on the SID ι_A transferred from A ’s cookie. Therefore, the request is processed as if it were intended by A .

This pattern captures different attacks when ϕ and κ are instantiated with different entities based on different program vulnerabilities. For instance, suppose a banking program has a CSRF vulnerability—it accepts payment requests solely based on the SID stored in the cookie, while a client is logged in. The attacker B may instantiate the above pattern with $\phi = \text{form } f_{pay}(\text{“payee” } p = \text{“}B\text{”, “amt” } a = 100) \text{ with } \langle \rangle$. Upon receiving the request, the server program would compose a closure to set up a payment from A . Similarly, a program with an XSS vulnerability would propagate a crafted string (e.g., cookie-stealing JavaScript code) from B ’s forged form to an HTML page displayed in A ’s browser, and a program with a session fixation vulnerability would set A up to use an SID fixed by B .

The second is a *malicious client* pattern, where a malicious client crafts a request to exploit server program vulnerabilities.

$$\begin{aligned} & (\vec{\Sigma}, (v, \dot{\sigma}), [(\iota_A, \vec{\phi}_A)]) \\ \rightsquigarrow & (\vec{\Sigma}, (v, \dot{\sigma}), [(\iota_A, \vec{\phi}_A \cup [\phi])]) && \text{(forge by Rule (34))} \\ \rightsquigarrow & (\vec{\Sigma}, \kappa, [(\iota_A, \vec{\phi}_A \cup [\phi])]) && \text{(request by Rule (35))} \end{aligned}$$

Since a malicious client A may use arbitrary target function names, parameter names and values, and sealed states (e.g., ones obtained from other forms in the history) to compose the request ϕ , a server program which does not carefully code against these possibilities would be subject to attacks including session poisoning, malformed input, and work-flow circumvention.

5. Translation

By comparing BASS (our proposed ideal programming model) and MOSS (the actual programming model in use today), the advantage of BASS is clear—none of the above attacks can happen to a BASS program, simply because BASS has a “well-behaved” semantics. We now discuss how this semantics can be implemented. Specif-

ically, we will point out how the two patterns of attacks are prevented. Although focusing on a particular translation, we note that other pertinent secure coding practices could be applied instead. The bottom line is, with the BASS abstractions implemented following secure coding practices once and for all, programmers enjoy an ideal model that is free from common exploits.

In particular, BASS has an ideal model where the control flow is not disrupted by web input, there is only a single session, and the client is well-behaved. These no longer hold in the MOSS model of real-world scenarios. To connect the two models, the translation takes care of program splitting, form making, session management, state maintenance, forgery prevention, and input validation. Recall from Figures 8 and 9 that a client in MOSS consists of a cookie and a list of forms. In our translation, we maintain the invariant that the cookie stores the CID of the client, and the forms are of the following shape:

```
form  $f(\vec{s} \vec{a} = \vec{w})$ 
  with  $\langle x_{cid} = l_{cid}, x_{sid} = l_{sid}, x_{tok} = l_{tok}, x_{fun} = s_f, \vec{x} = \vec{v} \rangle$ 
```

A key to the translation, this is in essence an encoding of BASS client closures in MOSS forms. Besides the target function f and input parameters \vec{a} , we also encode information about the local environment (nonvolatile variables $\vec{x} = \vec{v}$). In addition, some special variables are used for session management and forgery prevention.

The first special variable, x_{cid} , stores the CID that the server assigned to the client in Rule (38). In a translated program, before sending a response to the client, we obtain the CID from the cookie in the current closure, and store it in variable x_{cid} . On the other hand, when processing a request from a client, we inspect x_{cid} and check it against the client cookie. This ensures that the request indeed comes from the client, because: (1) a sealed state cannot be modified by a client; (2) the cookie of a client cannot be updated by others; (3) although a client may modify her cookie, she cannot change it to the value of another client's unforgeable CID. These together prevents attacker-crafted forms, inserted into a victim's history using Rule (36), from being accepted by the server. As a result, the attacker-victim pattern will not succeed.

The variable x_{sid} stores the SID. The variable x_{tok} stores a server generated token for history control. The validity of this token will be checked when the server program processes the request. For single-use forms, the token will be invalidated after the first use, thus preventing future resubmissions. In addition, the token can also be invalidated if we choose to expire all existing forms (cf. a BASS `clear` operation). We point out that values in x_{sid} and x_{tok} only need to be unique; they do not have to be unforgeable on their own, because they are stored in the sealed state and thus cannot be forged. We reuse the same token construct for CID, SID and history control in MOSS, but a realization may relax this.

Besides the three tokens, there is a special variable x_{fun} that stores the target function name. This is to prevent certain request forgeries where the client modifies the target function of the form. Upon a request, x_{fun} will be inspected and the request will be processed only if it matches the actual target function. This effectively binds the target function name and sealed state together in a form. Together with some type-based input validation code generated by the translation, this prevents attacks of the malicious client pattern.

The exact translation is given in Figure 12. Since all BASS expressions/values are MOSS expressions/values, the expression translation is trivial. In the translation rules, we simply adjust the font, updating E to \mathbf{E} to indicate the trivial translation from a BASS expression E to the corresponding MOSS expression \mathbf{E} .

Command translation has the form $|C|(\Sigma, \dot{\Sigma}, f) = (\dot{\Sigma}', \mathbf{C})$. Here we translate the BASS command C into the MOSS command \mathbf{C} . The translation takes as input a BASS global environment Σ , a MOSS global environment $\dot{\Sigma}$, and a function name f . Σ is propagated

$$\boxed{|C|(\Sigma, \dot{\Sigma}, f) = (\dot{\Sigma}', \mathbf{C})}$$

$$|\text{skip}|(\Sigma, \dot{\Sigma}, f) = (\dot{\Sigma}, f()) \quad (51)$$

$$\frac{|C|(\Sigma, \dot{\Sigma}, f) = (\dot{\Sigma}', \mathbf{C})}{|\text{skip}; C|(\Sigma, \dot{\Sigma}, f) = (\dot{\Sigma}', \text{sk};; \mathbf{C})} \quad (52)$$

$$\frac{|C|(\Sigma, \dot{\Sigma}, f) = (\dot{\Sigma}', \mathbf{C})}{|a := E; C|(\Sigma, \dot{\Sigma}, f) = (\dot{\Sigma}', a \leftarrow \mathbf{E};; \mathbf{C})} \quad (53)$$

$$\frac{|C|(\Sigma, \dot{\Sigma}, f) = (\dot{\Sigma}', \mathbf{C})}{|x := E; C|(\Sigma, \dot{\Sigma}, f) = (\dot{\Sigma}', x \leftarrow \mathbf{E};; \mathbf{C})} \quad (54)$$

$$\frac{\begin{array}{l} |C|(\Sigma, \dot{\Sigma}, f) = (\dot{\Sigma}_0, \mathbf{C}) \quad \text{fresh}(f') \\ \dot{\Sigma}_1 = \text{update}(\dot{\Sigma}_0, f', () \triangleright \mathbf{C}) \\ |C_1; \text{skip}|(\Sigma, \dot{\Sigma}_1, f') = (\dot{\Sigma}_2, \mathbf{C}_1) \\ |C_2; \text{skip}|(\Sigma, \dot{\Sigma}_2, f') = (\dot{\Sigma}_3, \mathbf{C}_2) \end{array}}{|\text{if } E \text{ then } C_1 \text{ else } C_2; C|(\Sigma, \dot{\Sigma}, f) = (\dot{\Sigma}_3, \text{cond}(\mathbf{E}, \mathbf{C}_1, \mathbf{C}_2))} \quad (55)$$

$$\frac{\begin{array}{l} |C_1|(\Sigma, \dot{\Sigma}, f) = (\dot{\Sigma}_0, \mathbf{C}_1) \quad \text{fresh}(f') \\ \dot{\Sigma}_1 = \text{update}(\dot{\Sigma}_0, f', () \triangleright \text{sk}) \\ |C; \text{skip}|(\Sigma, \dot{\Sigma}_1, f') = (\dot{\Sigma}_2, \mathbf{C}) \\ \dot{\Sigma}_3 = \text{update}(\dot{\Sigma}_2, f', () \triangleright \text{cond}(\mathbf{E}, \mathbf{C}, \mathbf{C}_1)) \end{array}}{|\text{while } E \text{ do } C; C_1|(\Sigma, \dot{\Sigma}, f) = (\dot{\Sigma}_3, f'())} \quad (56)$$

$$\frac{\begin{array}{l} (\vec{a} : \vec{\tau} = \vec{v}) \in \Sigma \quad |C|(\Sigma, \dot{\Sigma}, f) = (\dot{\Sigma}_0, \mathbf{C}) \\ \text{fresh}(f') \quad s = \text{toString}(f') \\ \mathbf{C}_1 = x_{cid} \leftarrow \text{getCk}();; x_{fun} \leftarrow s;; \\ \quad x_{tok} \leftarrow \text{newTk}();; a_{stok} \leftarrow \text{ins}(x_{tok}, a_{stok});; \\ \quad \text{input}(\text{form } f'(\vec{s} \vec{a} = \vec{v}) \text{ with pack}()) \\ \mathbf{C}_2 = \text{verif}(\vec{a}, \vec{\tau});; \text{unpack}(a_{srl});; \text{cond}(\mathbf{E}, \mathbf{C}_3, \text{err}) \\ \text{where } \mathbf{E} = (x_{cid} == \text{getCk}()) \wedge \text{in}(x_{tok}, a_{stok}) \wedge (x_{fun} == s) \\ \quad \mathbf{C}_3 = a_{stok} \leftarrow \text{del}(x_{tok}, a_{stok});; \mathbf{C} \\ \dot{\Sigma}_1 = \text{update}(\dot{\Sigma}_0, f', (a_{srl}, \vec{a}) \triangleright \mathbf{C}_2) \end{array}}{|\text{form}_S(\vec{a} : \vec{s}); C|(\Sigma, \dot{\Sigma}, f) = (\dot{\Sigma}_1, \mathbf{C}_1)} \quad (57)$$

$$\frac{\begin{array}{l} (\vec{a} : \vec{\tau} = \vec{v}) \in \Sigma \quad |C|(\Sigma, \dot{\Sigma}, f) = (\dot{\Sigma}_0, \mathbf{C}) \\ \text{fresh}(f') \quad s = \text{toString}(f') \\ \mathbf{C}_1 = x_{cid} \leftarrow \text{getCk}();; x_{fun} \leftarrow s;; x_{tok} \leftarrow a_{mtok};; \\ \quad \text{input}(\text{form } f'(\vec{s} \vec{a} = \vec{v}) \text{ with pack}()) \\ \mathbf{C}_2 = \text{verif}(\vec{a}, \vec{\tau});; \text{unpack}(a_{srl});; \text{cond}(\mathbf{E}, \mathbf{C}, \text{err}) \\ \text{where } \mathbf{E} = (x_{cid} == \text{getCk}()) \wedge (x_{tok} == a_{mtok}) \wedge (x_{fun} == s) \\ \dot{\Sigma}_1 = \text{update}(\dot{\Sigma}_0, f', (a_{srl}, \vec{a}) \triangleright \mathbf{C}_2) \end{array}}{|\text{form}_M(\vec{a} : \vec{s}); C|(\Sigma, \dot{\Sigma}, f) = (\dot{\Sigma}_1, \mathbf{C}_1)} \quad (58)$$

$$\frac{\begin{array}{l} |C|(\Sigma, \dot{\Sigma}, f) = (\dot{\Sigma}', \mathbf{C}) \\ \mathbf{C}' = a_{stok} \leftarrow \{\};; a_{mtok} \leftarrow \text{newTk}();; \mathbf{C} \end{array}}{|\text{clear}; C|(\Sigma, \dot{\Sigma}, f) = (\dot{\Sigma}', \mathbf{C}')} \quad (59)$$

$$\boxed{|W|_{\Sigma} = \mathbf{W}}$$

$$\frac{\begin{array}{l} \Sigma = \{\vec{a} : \vec{\tau} = \vec{v}\} \quad \sigma = \{\vec{x} : \vec{\tau} = \vec{w}\} \\ \dot{\Sigma}_0 = \{f_{exit} = () \triangleright \dots\} \\ |C; \text{skip}|(\Sigma, \dot{\Sigma}_0, f_{exit}) = (\dot{\Sigma}_1, \mathbf{C}) \\ \mathbf{C}' = \vec{a} \leftarrow \vec{v};; \vec{x} \leftarrow \vec{w};; a_{stok} \leftarrow \{\};; a_{mtok} \leftarrow \text{newTk}();; \mathbf{C} \\ \dot{\Sigma}_2 = \text{update}(\dot{\Sigma}_1, f_{start}, () \triangleright \mathbf{C}') \end{array}}{|(\Sigma, (\sigma, \mathbf{C}), \epsilon)|_{\Sigma} = (\dot{\Sigma}_2, \circ, \epsilon)} \quad (60)$$

Figure 12. Translation

throughout the translation without change; it is used to obtain the types of input variables when producing code for input validation. Σ is an environment to be updated to Σ' , so that new code blocks introduced during program splitting can be accumulated. f is provided as a continuation to the current command.

Rule (51) describes the base case of translating a `skip`. The translation simply produces the same environment Σ and a call to the continuation f . All other rules translate compound commands.

Rules (52), (53) and (54) translate compound commands which start with a `skip` or assignment. These commands are simply translated into their counterparts in MOSS. Note that the environment Σ may be updated to Σ' when translating the tail command C .

Rule (55) translates a compound command starting with a conditional. First, we translate the tail C into C , producing the environment Σ_0 . Next, we create a fresh label f' and add it into Σ_0 to refer to C . The updated environment Σ_1 is then used to translate the branches C_1 and C_2 (`skip` is inserted so that command translation needs only one base case—Rule (51)), with the new label f' as the continuation for both. Finally, we produce the translation result using the latest environment Σ_3 and a MOSS conditional. In this rule, the new label f' is created to restructure the control flow in case either branch involves web input. Suppose C_1 is of the form $a := 1; \text{form}_s(\dots); a := 2$. Rule (55) would produce split target code— C_1 would be $a \leftarrow 1;; \text{input}(\phi)$, and the link embedded in ϕ would point to a function with the code $a \leftarrow 2;; f'()$. If neither branch involves web input, an optimized version can be used to produce simpler target code, which is straightforward and omitted.

Rule (56) translates a while-loop. A fresh label f' is created as the entry point of the loop. We insert into Σ_1 some dummy code for f' , use Σ_1 to translate C with f' as the continuation, and update the code of f' with a conditional in the final environment Σ_3 . In the expression E evaluates to true, C will be executed; a loop is formed because C uses f' as the continuation. If E evaluates to false, C_1 (translated from C_1) will be executed, thus exiting the loop. Again, an optimized version can be used if C does not involve web input.

The translation of web input follows the intuition described at the beginning of this section. Rule (57) translates a single-use form `forms($\vec{a}; \vec{s}$)`. The program is split into two parts: the first part C_1 will be explicit in the translation result; the second C_2 will be stored in the result environment Σ_1 under a fresh label f' . In C_1 , we put the cookie value (the CID) into x_{cid} , put the function name s into x_{fun} (`toString` converts function names to unique strings), put a fresh token into x_{tok} , update a volatile variable a_{stok} to include the new token in the set of valid tokens, and send out a form. The form contains the target f' , the default values \vec{v} (we obtain these from Σ , but other values can be used in practice), and the sealed state created using `pack()`. In C_2 , which will be executed when the form is submitted, we validate the input parameters by type, unpack the sealed state, and check the validity of the request using E . In E , we compare the CID in x_{cid} against the cookie to detect attacker injected forms, check the validity of the token x_{tok} to prevent the single-use form from being submitted a second time, and check the target name x_{fun} to prevent forged links. If all these checks succeed, we remove the used token x_{tok} from the list of valid tokens a_{stok} , and proceed with the command C translated from C . Otherwise, we terminate the program with `err`.

Rule (58) translates a multi-use form. The idea is similar to Rule (57), except that we reuse the same token stored in a_{mtok} for all multi-use forms. In command C_2 , after the validation and unpacking, we check whether the token x_{tok} still matches a_{mtok} .

The last command translation case is described in Rule (59). Upon a `clear`, we simply reset the volatile variables a_{stok} and a_{mtok} . This effectively disables all existing forms on the client,

because the form translations in Rules (57) and (58) check the validity of the embedded tokens against a_{stok} and a_{mtok} .

World translation is given in Rule (60). We only translate those that correspond to unexecuted programs, where the client history is empty. The translation uses two special labels: f_{start} as the entry point of the program, and f_{exit} as the exit point. We translate the program body C with f_{exit} as the continuation, and create an entry for f_{start} in the final environment Σ_2 . The body of f_{start} initializes the program variables based on BASS environments Σ and σ , and assigns initial values to a_{stok} and a_{mtok} .

This concludes our formal translation, where the high-level semantics of BASS is enforced using careful manipulations of tokens and sealed states in MOSS. Note that a BASS program corresponds to a single session in MOSS. With the translation (a compiler) taking care of the common manipulations on web interactions and security once and for all, programs can be written in BASS following a simple and elegant declarative view of web programming.

5.1 Correctness and Prototyping

The correctness of the translation has been discussed throughout the translation. We now summarize some key arguments on the preservation of all the properties discussed in Section 3.4, which can be easily established by inspecting the form encoding presented earlier. Since the sealed state is protected, the only form component subject to client modification is the input parameters. Modifications on the target function name will not succeed, because the function name is also stored in the sealed state. As a result, the control flow of a BASS program will be enforced, since clients cannot forge target links—they can only follow the server-provided forms to access the server programs. In addition, an attacker cannot masquerade as another client—although an attacker may modify the cookie, she cannot set it to the unforgeable CID of another client. More interestingly, CSRF, first-order XSS and session fixation, which require an attacker to inject a form into a client (Rule (36)), are prevented because the injected form will have a different CID than that stored in the client’s cookie. Finally, input validation, session state recovery, and history control are all properly carried out.

We have implemented a prototype compiler of BASS in OCaml. The compiler takes well-formed BASS programs and produces secure Perl code following Figure 12. SSL is used for web interactions to prevent eavesdropping and tampering. All constructs of MOSS and the translation have clear counterparts in the prototype, demonstrating that our modeling is faithful to real-world scenarios, and that the translation can be carried out effectively. The compiler consists of 1456 lines of OCaml code (about 49KB), and uses a runtime library of 342 lines of Perl code (about 7KB). The executable in FreeBSD ELF format is about 346KB. The compiler translated the banking example of Figure 3 into 99 lines of Perl code. Focusing on the formal BASS language, the scale of this prototype is small. We plan to experiment with a larger scale prototype in the future, where the ideas of this report will be implemented for an existing language as a library.

6. Related and Future Work

Declarative web programming MAWL [3, 4] and its descendants (`<bigwig>` [6], JWIG [9]) use domain-specific language constructs to program web applications. They view web applications as form-based services, and provide abstractions on some key aspects such as web input and state management. These abstractions hide implementation details (e.g., input validation, embedding of continuation and state in URL), thus preventing certain programming errors. In another line of work, Graunke *et al.* [14] propose the design and implementation of an I/O construct for web interactions. This construct helps to program web applications in a more

traditional model of interactions, and avoids the manual saving and restoring of control state between interactions.

Although similar in spirit to BASS on these aspects, the above work does not provide a formal semantics with the same security guarantees. However, security should not be overlooked for declarative web programming—now that the details of web interactions are hidden by new abstractions, programmers can no longer carry out the secure coding practices by themselves. As a result, a naïve application of new abstractions could suffer from security vulnerabilities such as CSRF. It is thus crucial that the proposed abstractions and their implementation provide related security guarantees.

On expressiveness, MAWL and descendants enforce a strict control flow where every form is, in the BASS terminology, single-use. For example, users will be redirected to the beginning of a session if they hit the back button. In contrast, BASS leaves the design decision to the programmer, rather than disabling “whimsical navigation” [16] altogether. This additional flexibility is important for many applications [15, 16, 28].

We emphasize that the goal of BASS is to facilitate secure web programming with language abstractions more suitable for the domain. Besides having declarative support on web interactions, single-/multi-use forms, state declarations, and history control, it is important that the features are all modelled within an original and self-contained semantic specification. BASS has an intuitive and formal programming model with an explicit notion of a client, and its meta properties are articulated. This allows programmers to fully grasp how BASS programs behave. The common task of following secure coding practices, which is orthogonal to the specific application logic, is carried out by a BASS implementation once and for all.

There has also been work developing domain-specific languages or frameworks for web programming as libraries of existing server-safe languages. Some notable examples are the Curry-based server-side scripting language by Hanus [18], Haskell-based WASH/CGI by Thiemann [28], and Smalltalk-based Seaside by Ducasse *et al.* [12]. These provide useful abstractions in the form of libraries to handle some common aspects of web programming, such as structured HTML generation, session management, and client-server communication. However, they do not provide the formal security guarantees of BASS. Furthermore, there is no stand-alone formal semantics for the new abstractions, although in principle the behaviors could be inferred from the implementations and the semantics of the host languages. Finally, they are tied to the host languages, thus the ideas are not easily applicable to other languages. In contrast, BASS is translated into a flexible model underlying web programming in general. The rigorous BASS and MOSS allow us to formally derive security guarantees.

Finally, some recent work uses unified languages and frameworks for web application development, automatically partitioning programs among the server and the client. Most of such work does not address security. A notable exception is Swift [8], which ensures confidentiality and integrity of web application information using type annotations which reflect information flow policies. The security guarantee of Swift is largely orthogonal to those of BASS. If programmers use proper annotations, the general information flow guarantees of Swift can help guarding against some common vulnerabilities such as SQL injection and XSS. However, it remains vulnerable to (it is up to the programmer to write secure code against) others such as CSRF. In addition, the formal semantics of BASS helps providing a rigorous foundation on web interaction and browser navigation in general.

Models of web programs Besides building new abstractions, it is also important to understand what the existing web programming model is and how to write secure programs within it. Although there are many security recommendations and coding practices available

on the topic, web programming has rarely been formally studied from the language principles. An exception is the work by Graunke *et al.* on modeling web interactions [16]. In particular, they model web interactions in presence of whimsical navigation behaviors of well-intended (non-malicious) users, identify two classes of errors (form field mismatch, client-server state mismatch), and propose to catch the errors with a static type system (typed web forms) and dynamic checks (time-stamped states).

In comparison with our work, Graunke *et al.* do not address the wider security questions, where the server interacts with multiple clients, some of whom may be malicious. Specifically, they do not have security-related primitives in their server-side computation, they model web programs in the presence of a single client, and their transitions lack the counterparts of the forging (Rule (34)) and tricking (Rule (36)) behaviors of MOSS. In contrast, we more accurately model web interactions and server computations. This allows us to encode common exploit patterns (Section 4.3), which are then amenable to formal reasoning. Instead of using a type system or dynamic checks to close up the two classes of errors, we use principled but flexible abstractions in BASS for web programming and present their realization in MOSS. This hides some common details of secure web programming and precludes some different classes of errors and vulnerabilities. Although we do not propose techniques preventing errors associated with whimsical navigations, the corresponding error scenarios can be rigorously illustrated in BASS and MOSS; therefore, programmers are fully aware of the issues, and the techniques of Graunke *et al.* can be applied.

Future work We believe that web programming will benefit significantly from the use of domain-specific abstractions, and much can be done in the area. For starters, the form constructs of BASS can be generalized to take an arbitrary HTML page as a parameter in practice. Sometimes an HTML page may contain multiple forms, each with a separate submit button linking to a different target. A more interesting aspect is to include client-side scripting (*e.g.*, JavaScript) into the model. This is orthogonal to the focus of server-side scripting in this report, and previous work on the formal aspects of JavaScript [29, 2, 32] and web application development frameworks and policies [13, 21] provide some good starting points.

Using some new abstractions, BASS provides several useful security guarantees as articulated in Section 3.4. However, the BASS abstractions are not meant to be “complete,” and there are other desirable program properties uncovered. It is thus useful to explore abstractions for other areas, such as dynamic HTML generation, privilege management, and dynamic SQL construction. HTML generation has been much studied [18, 9, 23, 28], where the well-formedness of the generated document can be guaranteed statically. Privilege management has been discussed in Section 2.4. Dynamic SQL construction is worth mentioning. Most web languages manipulate SQL commands with strings, naïve handling of which is subject to injection attacks. Previous work [27, 31] has articulated that injection attacks happen when client input changes the intended structure of the constructed commands. In response, it is useful to provide abstractions that statically enforce the structures of SQL commands before client input is plugged in at runtime. Prepared statements [20] provide some support on this aspect if used properly, but further study on more generally shaped queries and more friendly syntax is useful. Similar ideas may also help addressing second-order XSS [24].

Designed for web programming in general, BASS addresses only common security aspects, rather than issues on the specifics of an application. For example, directory traversal [10] (accessing the parent directory using the path “. . . \”) is not prevented by common type-based input validation, and programmers must perform additional filtering. Application-specific security and code analysis will still be necessary. However, with the new abstractions closing

up some common vulnerabilities and clarifying the control flow, such analysis should be easier. In general, the new abstractions should also help the analysis, reasoning, and testing of web programs, because they provide an ideal model (e.g., structured control flow, automatic state maintenance, single well-behaved client) that is amenable to established language techniques.

In this report, MOSS serves mainly as the target of the BASS translation. However, it is also of independent value to web programming studies. In MOSS, a client may take control only if the server contains no active closure. This is similar in spirit to non-preemptive multi-threading. A preemptive version of MOSS is interesting future work. In addition, type systems and program logics for confining and reasoning about MOSS programs are useful.

7. Conclusion

Web applications reflect a different computation model than conventional software, and the security issues therein deserve careful study from both language principles and practical implementations. This report serves as a useful step towards building a formal foundation for secure server-side scripting. In particular, we propose two self-contained formalizations on the topic, using familiar language concepts such as continuations, threads, and small-step semantics.

The first is a formal language BASS for server-side scripting. BASS provides an ideal programming model where the server interacts with a single client, using some dedicated constructs to obtain web input and manipulate client history. The meta properties and formal guarantees of BASS allow programmers to focus on the application logic without being distracted by common implementation details, thus improving productivity and security.

The second is a formal model MOSS characterizing realistic web programming concepts. As is the case of existing server-side scripting languages, MOSS can be used to write both secure programs and vulnerable ones. We present a formal translation from BASS to MOSS, demonstrating how the BASS abstractions and guarantees can be enforced using a few common primitives for manipulating security concepts. Our prototype shows much promise on the idea of better abstractions for secure server-side scripting, and we hope to experiment more on the topic with real-world web programming languages and frameworks.

Acknowledgments

We wish to thank Zhong Shao, Gary Wassermann, and the anonymous reviewers at WWW 2008 for helpful suggestions.

References

- [1] M. Achour *et al.*. PHP manual. <http://www.php.net>, 2007.
- [2] C. Anderson, P. Giannini, and S. Drossopoulou. Towards type inference for JavaScript. In *Proc. 19th European Conference on Object-Oriented Programming*, pages 429–452, July 2005.
- [3] D. Atkins, T. Ball, M. Benedikt, G. Bruns, K. Cox, P. Mataga, and K. Rehor. Experience with a domain specific language for form-based services. In *Proc. 1997 Conference on Domain-Specific Languages*, 1997.
- [4] D. L. Atkins, T. Ball, G. Bruns, and K. Cox. MAWL: A domain-specific language for form-based services. *IEEE Trans. on Software Engineering*, 25(3):334–346, 1999.
- [5] R. Auger. The Cross-Site Request Forgery FAQ. <http://www.cgisecurity.com/articles>, 2007.
- [6] C. Brabrand, A. Møller, and M. I. Schwartzbach. The <bigwig> project. *ACM Trans. on Internet Technology*, 2(2):79–114, 2002.
- [7] CERT. Computer emergency response team. <http://www.cert.org>, 2007.
- [8] S. Chong, J. Liu, A. C. Myers, X. Qi, K. Vikram, L. Zheng, and X. Zheng. Secure web applications via automatic partitioning. In *Proc. 21st Symposium on Operating Systems Principles*, pages 31–44, Oct. 2007.
- [9] A. S. Christensen, A. Møller, and M. I. Schwartzbach. Extending Java for high-level Web service construction. *ACM Trans. on Programming Languages and Systems*, 25(6):814–875, Nov. 2003.
- [10] S. Christey and R. A. Martin. Vulnerability type distributions in CVE. <http://cve.mitre.org/docs/vuln-trends>, 2007.
- [11] P. Delisle, J. Luehe, and M. Roth. JavaServer Pages specification. <http://java.sun.com/products/jsp>, 2006.
- [12] S. Ducasse, A. Lienhard, and L. Renggli. Seaside — a multiple control flow web application framework. In *Proc. 12th International Smalltalk Conference*, pages 231–257, Sept. 2004.
- [13] Ú. Erlingsson, B. Livshits, and Y. Xie. End-to-end web application security. In *Proc. 11th Workshop on Hot Topics in Operating Systems*, May 2007.
- [14] P. Graunke, R. B. Findler, S. Krishnamurthi, and M. Felleisen. Automatically restructuring programs for the Web. In *Proc. 16th International Conference on Automated Software Engineering*, pages 211–222, Nov. 2001.
- [15] P. Graunke and S. Krishnamurthi. Advanced control flows for flexible graphical user interfaces. In *Proc. 2002 International Conference on Software Engineering*, pages 277–296, 2002.
- [16] P. Graunke, S. Krishnamurthi, S. V. D. Hoeven, and M. Felleisen. Modeling web interactions. In *Proc. 2003 European Symposium on Programming*, pages 122–136, 2003.
- [17] D. H. Hansson. Ruby on Rails. <http://www.rubyonrails.org>.
- [18] M. Hanus. High-level server side web scripting in Curry. In *Proc. 3rd International Symposium on Practical Aspects of Declarative Languages*, pages 76–92, 2001.
- [19] J. Kolšek. Session fixation vulnerability in web-based applications. <http://www.acrossecurity.com/papers.htm>, 2002.
- [20] J. Lam. Prepared statements and SQL injections. <http://isc.sans.org/diary.html?storyid=2301&rss>, 2007.
- [21] B. Livshits and Ú. Erlingsson. Using web application construction frameworks to protect against code injection attacks. In *Proc. Programming Languages and Analysis for Security*, June 2007.
- [22] G. A. D. Lucca, A. R. Fasolino, M. Mastoianni, and P. Tramontana. Identifying XSS vulnerabilities in web applications. In *Proc. 6th International Workshop on Web Site Evolution*, pages 71–80, 2004.
- [23] K. Nørmark. Web programming in Scheme with LAML. *Journal of Functional Programming*, 15(1):53–65, 2005.
- [24] G. Ollmann. Second-order code injection attacks. <http://www.nextgenss.com/papers>, 2004.
- [25] OWASP. Open web application security project. <http://www.owasp.org>, 2007.
- [26] OWASP Foundation. The ten most critical web application security vulnerabilities. <http://www.owasp.org>, 2007.
- [27] Z. Su and G. Wassermann. The essence of command injection attacks in web applications. In *Proc. 33rd Symposium on Principles of Programming Languages*, pages 372–382, Jan. 2006.
- [28] P. Thiemann. An embedded domain-specific language for type-safe server-side web scripting. *ACM Trans. on Internet Technology*, 5(1):1–46, Feb. 2005.
- [29] P. Thiemann. Towards a type system for analyzing JavaScript programs. In *Proc. 2005 European Symposium on Programming*, pages 408–422, Apr. 2005.
- [30] WASC. Web application security consortium. <http://www.webappsec.org>, 2007.
- [31] G. Wassermann and Z. Su. Sound and precise analysis of web applications for injection vulnerabilities. In *Proc. 2007 Conference on Programming Language Design and Implementation*, pages 32–41, June 2007.
- [32] D. Yu, A. Chander, N. Islam, and I. Serikov. JavaScript instrumentation for browser security. In *Proc. 34th Symposium on Principles of Programming Languages*, pages 237–249, Jan. 2007.