

An Abstract Semantics for Inference of Types and Effects in a Multi-Tier Web Language

Letterio Galletta Giorgio Levi

Dipartimento di Informatica, Università di Pisa

{galletta, levi}@di.unipi.it

Types-and-effects are type systems, which allow one to express general semantic properties and to statically reason about program's execution. They have been widely exploited to specify static analyses, for example to track computational side effects, exceptions and communications in concurrent programs. In this paper we adopt abstract interpretation techniques to reconstruct (following the Cousot's methodology) a types-and-effects system developed to handle security problems of a multi-tier web language. Our reconstruction allows us to show that this types-and-effects system is not sound with respect to the semantics of the language. In addition, we correct the soundness issues in the analysis and systematically construct a correct analyser.

1 Introduction

Types-and-effects systems are a powerful extension of type systems which allows one to express general semantic properties and to statically reason about program's execution. The underlying idea is to refine the type information so as to express further intensional or extensional properties of the semantics of the program: in practice, they compute the type of each program's sentence and an approximate (but sound) description of its run-time behavior. Since they are defined over the well understood theory of type systems, they are an intuitive framework for specifying and for developing static analyses. Such systems were originally introduced in [12] to statically track side effects in languages that mix functional and imperative feature. However, they have been employed to control many other kinds of computational effects and analyses, e.g. exceptions [18], region inference [22] and communications in concurrent programs [21]. Recently, they have been used in [2] to handle security issues in LINKS [4].

LINKS is a strict, typed, functional language for web applications. Its main feature is to be multi-tier, that is, it enables the developer to mix client, server and database source code by delegating the charge of code and data partitioning to the compiler: from a single source file the compiler generates code for the database back-end, for the web server and the client front-end, ensuring that all data is stored either in client or in database. In [2] Baltopoulos and Gordon have shown that storing unencrypted application data on the client opens LINKS to attacks that may expose secrets and modify control-flow and application data. In order to overcome these problems they have proposed a compilation strategy based on authenticated encryption¹ and a types-and-effects system to enforce programs to satisfy a particular class of integrity constraints (event-based assertions). This types-and-effects system formalizes source level reasoning about LINKS programs and allows them to prove security properties by inspection of the source code. For the definition of this system they have followed a methodology characterized by translating each LINKS expression to an expression of a concurrent λ -calculus with refinement types [3]. This translation hides the properties of the analysis, and does not guarantee the soundness with respect to the

¹a combination of secrecy and integrity protection obtained by encrypting together data and its hash.

semantics of the language. Hence, we decided to study the properties of this analysis by reconstructing it by abstract interpretation [11].

Abstract interpretation [6, 7, 8, 9] is a general theory for approximating the semantics of dynamic systems. The key idea behind abstract interpretation is that the description of the behavior of a system (at various levels of abstraction) is an approximation of its formal semantics. In static analysis this means that every property of a program can be observed in its semantics and computed as an approximation: the intuition is that the analysis can be systematically derived by throwing away superfluous information from the semantics. In practice, the approximated semantics (abstract semantics) is obtained from the standard one (called concrete) by substituting the actual (concrete) domain of computation and its basic semantic operations with abstract domain and abstract semantic operations, respectively. The basic idea is that the abstract domain is a representation of some properties of interest about concrete domain's values, while abstract operations simulate, over the abstract domain, the behavior of their concrete counterparts. Hence the abstract semantics computes the properties of interest and the analysis algorithm corresponds to evaluating programs over the abstract domain. Since the abstract domain is a sound approximation of the concrete one, the analysis algorithm is correct with respect to the semantics by construction.

Type systems (and corresponding type inference algorithms) have been reconstructed as a hierarchy of abstract interpretations by Cousot [5]. In order to reconstruct the types-and-effects analysis of LINKS we extend Cousot's methodology by defining an abstract domain able to express types augmented by effects. In this paper we give the following contributions:

- we demonstrate that the analysis defined by Baltopoulos and Gordon is not sound: in fact, the expression `get(Text("Hello!"))` is type-checked but it results in a run-time type error (Section 3)
- we show how to fix this unsoundness issue (Section 3)
- we systematically derive an abstract semantics which represents a correct analyser (we have implemented it in OCaml [16]) (Sections 4 and 5)

In the next sections we first will sketch the type-and-effect system proposed for LINKS (Section 2), then we describe the ideas and the methodology underlying our reconstruction.

2 Secure Compilation of LINKS

Standard web applications have a multi-tier architecture: user interface, application logic and data management are implemented over three different tiers. Each tier runs on a different computational environment (web browser, web server and database respectively) characterized by its own language and its data representation. This heterogeneity gives rise to the problem of *impedance mismatch* [19]: because each language has its own data type, data exchanged between tiers of same application have a different representation. This problem complicates the development of web applications because programmers need to define routines to interchange and convert data. To solve this problem a new class of web languages (multi-tiers languages) have been developed. These languages allow programmers to blend server, client and database source code and provide automatic mechanisms for the partition of the application over tiers.

LINKS is a functional programming language for web applications that belongs to the class of multi-tiers languages. LINKS enables developers to mix client, server and database source code by delegating

the charge of code and data partitioning to the compiler: from a single source code the compiler generates code for the database back-end, for the server and for the client front-end.

In this way LINKS overcomes the problem of impedance mismatch by abstracting details of a single tier and by supporting an unified programming model similar to the one used for GUI applications. To realize this cross-tier programming model LINKS exploits the mechanism of the web continuation [23]. These continuations are implemented as closures (expression to be executed plus values of free variables) and are stored in HTML pages either as hidden fields of forms or as URL parameters. This approach gives rise to security risks since a malicious client may modify those closures to enforce unexpected computations on the server.

In particular, Baltopoulos and Gordon in [2] have demonstrated that the approach adopted by LINKS of storing unencrypted data on the client is not secure because an attacker may violate the data secrecy, the data integrity and the control-flow integrity of the application. To overtake these problems they have proposed a secure implementation of LINKS that includes a compilation strategy based on authenticated encryption to protect the closures held in the browser and a types-and-effects system to enable source level reasoning about security of web applications. This secure implementation has been formalized for TINYLINKS, a simple subset of LINKS.

TINYLINKS is a λ -calculus augmented with XML values for representing web pages and annotation expressions for expressing safety properties. Its syntax is shown in Figure 1. HTML pages are values created by applying the data constructors `Text` and `Elem`: the first one represents simple text in HTML document, the second one a generic tag element. To express links and forms exists two ad-hoc data constructors that contains suspended expressions ². `href(E)` is a link that, when clicked, evaluates the expression `E`. `form([l1, ..., ln], E)` is a HTML form with a suspended computation (the expression `E`) which requires user input. The input is represented by labels `[l1, ..., ln]` that will contain the values inserted in the input fields of the form. The evaluation of `href` and `form` can be accomplished by using the operators `get` and `post`, respectively ³. The annotations `eventL` and `assertL` have no computational meaning. They allow us to annotate TINYLINKS programs with event-based assertions expressing suitable safety properties. An expression is safe if whenever an assertion `assertL` occurs in the execution, there exists a previous occurrence of an event `eventL`.

Baltopoulos and Gordon have defined a dependent types-and-effects system to verify that each expression of a program is safe. This system is specified by a set of inductively defined typing judgments. These judgments are of the form $\Gamma; F \vdash E \xrightarrow{exp} \langle _ : T \rangle \{ F' \}$, where Γ is the typing environment, F is the set of events which have occurred and are needed to safe evaluation of the expression `E` (precondition); T and F' are, respectively, the type of value and the set of events (post-condition) yielded by the execution of `E`.

The typing rules for the operations `get` and `post`, for the annotations `event` and `assert` and for the function application are shown in Figure 2. Rule (T-Get) establishes that the type assigned to `get` is `xml` (that represent the type of a generic HTML tag) with empty effect, provided that `V` is another HTML tag. By (T-Post), the type of `post` expression is `xml` with empty effect, provided that the values associated with submission labels are strings and that `U` is a HTML tag. By (T-Event) `eventL` has type `unit` and effect `L`, provided that the values in the event `L` have a type. Rule (T-Assert) is similar to (T-Event) except that requires `L ∈ F`, that is the precondition of the judgment includes `L`. Rule (T-App) is typical for application and shows how the mechanism of the annotations works: the expression is type checked if only if the events in the precondition F_1 of the function have occurred in F with same values. The

²we can look at these values as special kinds of functional abstractions.

³we can look at these operations as special kinds of function application.

f, y, x	Variables
p	Predicates
$c ::= \text{Unit} \mid \text{Zero} \mid \text{Succ} \mid \text{String}$ $\mid \text{Nil} \mid \text{Cons} \mid \text{Tuple} \mid \text{Elem} \mid \text{Text}$	Data constructors
$g ::= + \mid - \mid * \mid /$	Primitive operators
$L ::= p(V_1, \dots, V_n)$	Events: a predicate and a list of values
$V, U ::= x \mid c(V_1, \dots, V_n) \mid \text{href}(E)$ $\mid \lambda x_1. \dots, x_n. E \mid \text{form}([l_1, \dots, l_n], E)$	Values
$E ::= V \mid \text{var } x = E_1; E_2 \mid g(E_1, E_2)$ $\mid V(U_1, \dots, U_n) \mid \text{post}([l_1 = V_1, \dots, l_n = V_n], U)$ $\mid \text{get}(V) \mid \text{event } L \mid \text{assert } L$ $\quad \text{switch}(V)\{$ $\quad \quad \text{case } c(x_1, \dots, x_n) \rightarrow E_1$ $\quad \quad - \rightarrow E_2$ $\quad \}$	Expressions

Figure 1: Syntax of TINYLINKS

events generated after application include the ones of the post-condition of U .

We say that a web application E is safe if and only if there is a derivation within the types-and-effects system of the judgment $\emptyset; \emptyset \vdash E \stackrel{exp}{\Leftarrow} \langle _ : \text{xml} \rangle \{ \}$, meaning that E is a closed expression which requires no precondition and which yields a web page without generating further events.

After the definition of typing rules, the standard methodology requires to state and prove the soundness theorem which guarantees the validity of the analysis with respect to the semantics of the language. Baltopoulos and Gordon adopt a different approach by translating each TINYLINKS expression to an expression of a concurrent λ -calculus with *refinement types*. This translation hides the details and the properties of the defined types-and-effects system, in particular the soundness. For instance, the expression $\text{get}(\text{Text}(\text{"Hello!"}))$ is safe because a derivation exists for the judgment $\emptyset; \emptyset \vdash \text{get}(\text{Text}(\text{"Hello!"})) \stackrel{exp}{\Leftarrow} \langle _ : \text{xml} \rangle \{ \}$. However, we will show in the next section that the proposed types-and-effects system is not sound because, even if this expression is type checked, its evaluation results in a run-time type error.

3 A Denotational Semantics for TINYLINKS

In this paper we adopt the approach described by Cousot in [5]. We define a denotational semantics for TINYLINKS, by considering it as an untyped λ -calculus. Furthermore, since we deal with effects, we explicitly consider assertions of events. To this purpose we introduce a special environment (*events environment*) which will store occurred events. The semantics of $\text{assert } q(V_1, \dots, V_n)$ will require checking that q is bound in this environment to values V_1, \dots, V_n . If this check succeeds, the evaluation yields a `Unit` value, otherwise a “sentinel” value indicating an error.

For the sake of simplicity, we restrict the values in an event to integers only. We will also assume that functions have a single argument and predicates in events are bound to a single value. Since we regard TINYLINKS an untyped λ -calculus, we define the semantics domain of values (*Eval*) as a recursive sum of cpos, by using the inverse limit construction described in [24]. Each element of this sum represents a specific class of values. For instance, \mathcal{Z} is the set of integers; U and S are singletons of the unit value

$$\begin{array}{c}
\text{(T-Get)} \frac{\Gamma; F \vdash V \stackrel{val}{\Leftarrow} \text{xml}}{\Gamma; F \vdash \text{get}(V) \stackrel{exp}{\Rightarrow} \langle _ : \text{xml} \rangle \{ \}} \\
\text{(T-Post)} \frac{\Gamma; F \vdash V_i \stackrel{val}{\Leftarrow} \text{string} \quad \forall i \in \{1, \dots, n\} \quad \Gamma; F \vdash U \stackrel{val}{\Leftarrow} \text{xml}}{\Gamma; F \vdash \text{post}([\mathbf{l}_1 = V_1, \dots, \mathbf{l}_n = V_n], U) \stackrel{exp}{\Rightarrow} \langle _ : \text{xml} \rangle \{ \}} \\
\text{(T-Event)} \frac{\Gamma \vdash \diamond \quad fV(F, L) \subseteq \text{dom}(\Gamma) \quad L = \mathbf{p}(V_1, \dots, V_n) \quad \Gamma; F \vdash V_i \stackrel{val}{\Rightarrow} T_i \quad \forall i \in \{1, \dots, n\}}{\Gamma; F \vdash \text{event } L \stackrel{exp}{\Rightarrow} \langle _ : \text{unit} \rangle \{ L \}} \\
\text{(T-Assert)} \frac{\Gamma \vdash \diamond \quad fV(F, L) \subseteq \text{dom}(\Gamma) \quad L \in F \quad L = \mathbf{p}(V_1, \dots, V_n) \quad \Gamma; F \vdash V_i \stackrel{val}{\Rightarrow} T_i \quad \forall i \in \{1, \dots, n\}}{\Gamma; F \vdash \text{assert } L \stackrel{exp}{\Rightarrow} \langle _ : \text{unit} \rangle \{ L \}} \\
\text{(T-App)} \frac{\Gamma; F \vdash U \stackrel{val}{\Rightarrow} T \quad T = \langle x_1 : T_1, \dots, x_n : T_n \rangle \{ F_1 \} \rightarrow T_2 \{ F_2 \} \quad fV(T) = \emptyset \quad \Gamma; F \vdash V_i \stackrel{val}{\Leftarrow} T_i \quad \forall i \in \{1, \dots, n\} \quad F_1[V_1/x_1] \dots [V_n/x_n] \subseteq F}{\Gamma; F \vdash U(V_1, \dots, V_n) \stackrel{exp}{\Rightarrow} T_2 \{ F_2[V_1/x_1] \dots [V_n/x_n] \}}
\end{array}$$

Figure 2: Some examples of rules specifying the type-and-effect system for the correspondences analysis.

and the error value ⁴; $EEnv \rightarrow Eval \rightarrow (Eval \times EEnv)$, $EEnv \rightarrow (Eval \times EEnv)$ and $EEnv \rightarrow [Eval] \rightarrow (Eval \times EEnv)$ are the sets of the denotations of functions, links and forms, respectively.

The environment (Env) is a function from identifier (Ide) to values ($Eval$). The events environment ($EEnv$) maps predicates ($Pred$) to pairs formed by an element of $Dval$ and an element of $Mark$. $Dval$ denotes values which can occur in an event ⁵. $Mark$ is the state of an event: E indicates that the event has occurred, EA that has occurred and has been asserted, A that has only been asserted.

We define two semantic functions $\mathcal{V}[_] : VAL \rightarrow Env \rightarrow EEnv \rightarrow Eval$ for values and $\llbracket _ \rrbracket : EXP \rightarrow Env \rightarrow EEnv \rightarrow (Eval \times EEnv)$ for expressions. The semantics of values is straightforward, because we only need to construct the corresponding denotation. Some examples of semantic equation are shown in Figure 3. In the definition, we use injections into $Eval$ (like $Unit$, $Href$, Fun), continuous semantic operators (like $bindList$) and a meta-language which includes:

- $if e_1 \text{ then } e_2 \text{ else } e_3$ (conditional);
- $let x = e_1 \text{ in } e_2$ as a cleaner notation for $((\lambda x.e_2) e_1)$;
- $let^* x = e_1 \text{ in } e_2$ for $((\lambda x.e_2)^* e_1)$;
- $case e_1 \text{ of } in_1(x_1) \rightarrow e_2 _ \rightarrow e_3$ for $[\lambda x_1.e_1, \lambda x_2.e_3, \dots, \lambda x_2.e_3]$;
- $let (x_1 x_2) = e_1 \text{ in } e_2$ for $let y = e_1 \text{ in } let x_1 = \pi_1(y) \text{ let } x_2 = \pi_2(y) \text{ in } e_2$;

where π_i , $[_]$, \star and $[_, \dots, _]$ are the standard operators for product, lifting and sum of cpos [26].

The semantics of expressions is similar to the one of the untyped λ -calculus. The most interesting cases of semantic equations are shown in Figure 4 and below we give some comments about them.

⁴this value is used to show a run-time type error

⁵in the following we will call them denotable values

$$\begin{aligned}
\mathcal{V} \llbracket \lambda x. E \rrbracket \rho \phi &= \lfloor \text{Fun}(\lambda \phi'. \lambda v. \llbracket E \rrbracket \rho [v/x] \phi') \rfloor \\
\mathcal{V} \llbracket \text{Href}(E) \rrbracket \rho \phi &= \lfloor \text{Href}(\lambda \phi'. \llbracket E \rrbracket \rho \phi') \rfloor \\
\mathcal{V} \llbracket \text{Form}(ll, E) \rrbracket \rho \phi &= \lfloor \text{Form}(\lambda \phi'. \lambda vl. \text{let}^* \rho' = \text{bindList}(\rho, ll, vl) \text{ in } \llbracket E \rrbracket \rho' \phi') \rfloor
\end{aligned}$$

Figure 3: Examples of semantic equations for values.

The semantics of $\text{get}(V)$ asks to evaluate V ; if the evaluation results into the denotation of a link ($\text{Href}(f)$), we evaluate the corresponding suspended expression (the closure f), otherwise we return an error value.

The semantics of $\text{post}(VL, V)$ is similar: if the evaluation of V is a form ($\text{Form}(f)$) and the evaluation of VL is a list of strings, we return the result of the application of the functional value f to the denotation of VL and to the current events environment ϕ .

The semantics of $\text{event } q(V)$ requires the evaluation of V ; if the produced value is an integer, we create a new binding for the predicate q in ϕ and return a unit value otherwise we raise an error.

The semantics of $\text{assert } q(V)$ is similar, but requires the evaluation of V to be equal to the value bound to the predicate q in ϕ . In this case we update the state of the event in ϕ and return a unit value.

By using the semantic equation of get we prove that the evaluation of the expression $\text{get}(\text{Text}(\text{"Hello!"}))$ results in a run-time type error (the value $\lfloor \text{WrongValue}() \rfloor$) because the denotation of $\text{Text}(\text{"Hello!"})$ is not a link. Although a link is an XML value, it is different from other XML values because it is a special kind of functional abstraction. Notice that the type-and-effect system proposed for TINYLINKS does not handle this special nature of links correctly, because it assigns the same type to the all XML values. Note that the same remark can be made for forms. Our above arguments demonstrate that the types-and-effects system of [2] is unsound because exists an expression which is type checked but its evaluation yields yet a run-time type error. We argue that the solution to this problem is to use a type system with subtypes. For the sake of simplicity, in our reconstruction we will not use subtypes, but we will instead define two ad-hoc types for forms and links which will handled so as have a sound analysis.

4 An Abstract Semantics for Inference of Types and Effects

Following the classical methodology of abstract interpretation, once we have defined a concrete semantics, we need to define a collecting semantics by extending $\mathcal{V} \llbracket - \rrbracket$ and $\llbracket - \rrbracket$ to the powerset.

The concrete semantics properties, which we are interested in, are the types and the event-based annotations. We need to define a suitable domain for both. One possibility is to define the abstract domain as the set of Hindley's monotypes (terms) with variables [15, 10, 5, 20, 13]. However, this is not possible, since types are annotated by effects. For example, a function type will have the form $T_1 \{ F_1 \} \rightarrow T_2 \{ F_2 \}$, where F_1 are the events which have to be occurred before the function application, whereas F_2 are the events which we can consider occurred afterwards. Hence, we need to define a domain of annotated types. The main problem is that the algebra of annotated terms is not free. In fact, two types can be identified even if their syntax is different. For example, the types $\text{xml} \{ q(10), p(1) \} \rightarrow \text{xml} \{ \}$ and $\text{xml} \{ p(1), q(10) \} \rightarrow \text{xml} \{ \}$ have a different representation, but they are equal because the effects $\{ q(10), p(1) \}$ and $\{ p(1), q(10) \}$ denote the same set. Therefore, we cannot use a syntactic unification algorithm [17] to solve equations between terms.

$$\begin{aligned}
\llbracket \text{get}(V) \rrbracket \rho \phi &= \text{let}^* v' = \mathcal{V} \llbracket V \rrbracket \rho \phi \text{ in} \\
&\quad \text{case } v' \text{ of} \\
&\quad \quad H \text{ref}(f) \rightarrow f \phi \\
&\quad \quad - \rightarrow (\llbracket \text{WrongValue}() \rrbracket, \iota) \\
\llbracket \text{post}(VL, V) \rrbracket \rho \phi &= \text{let}^* v' = \mathcal{V} \llbracket V \rrbracket \rho \phi \text{ in} \\
&\quad \text{let}^* v_2 = \text{checkStringList}(\text{map}(\lambda x. \mathcal{V} \llbracket x \rrbracket \rho \phi) VL) \text{ in} \\
&\quad \text{case } v' \text{ of} \\
&\quad \quad \text{Form}(f) \rightarrow \text{case } v_2 \text{ of} \\
&\quad \quad \quad V(vl) \rightarrow f vl \phi \\
&\quad \quad \quad - \rightarrow (\llbracket \text{WrongValue}() \rrbracket, \iota) \\
&\quad \quad - \rightarrow (\llbracket \text{WrongValue}() \rrbracket, \iota) \\
\llbracket \text{event } q(V) \rrbracket \rho \phi &= \text{let}^* d = \text{evalToDval}(\mathcal{V} \llbracket V \rrbracket \rho \phi) \text{ in} \\
&\quad \text{if } d = \text{dint}(n) \text{ then} \\
&\quad \quad (\llbracket \text{Unit}() \rrbracket, \phi [(d, E)/q]) \\
&\quad \text{else} \\
&\quad \quad (\llbracket \text{WrongValue}() \rrbracket, \iota) \\
\llbracket \text{assert } q(V) \rrbracket \rho \phi &= \text{let}^* ev = \text{evalToDval}(\mathcal{V} \llbracket V \rrbracket \rho \phi) \text{ in} \\
&\quad \text{let } (ev', m) = \phi q \\
&\quad \text{if } ev = ev' \text{ then} \\
&\quad \quad (\llbracket \text{Unit}() \rrbracket, \phi [(ev', EA)/q]) \\
&\quad \text{else} \\
&\quad \quad (\llbracket \text{WrongValue}() \rrbracket, \iota)
\end{aligned}$$

Figure 4: Examples of semantic equations for expressions.

One solution would be to use an algorithm for unifying terms in non-free algebras (semantic unification). Such algorithms do exist [1], but they are not usable in practice.

Our reconstruction does not rely on semantic unification but on another approach described in [22]. This approach exploits special annotated types (simple types), where annotations are replaced by variables (annotation variables), whose values have to satisfy some constraint. For example, the annotated type $\text{xml} \{ q(10), p(1) \} \rightarrow \text{xml} \{ \}$ becomes $\text{xml}(\alpha) \rightarrow \text{xml}(\beta)$, where α and β are the minimal annotations A and B which satisfy the constraints $A \supseteq \{ q(10), p(1) \}$ and $B \supseteq \{ \}$, respectively. The algebra of simple types is free. Hence, the introduction of a new kind of variable in terms requires a simple variation of the unification algorithm: an annotation variable unifies with another annotation variable only.

However, this solution is not completely adequate to define an abstract domain for the properties which we are concerned with, because the effects depend on the values. Hence we need to include them in the abstract domain. Since events in the precondition and post-condition of a function type may depend on the value bound to a formal parameter we need to remember it. We then introduce in the set of terms another kind of variables, called identifier variables. Identifier variables are handled by simple

modification of the unification algorithm: an identifier variable unifies with another identifier variable only.

The domain of abstract values will contain also substitutions as in [13]. The role of substitutions can be explained as follows. At some point in the evaluation of the abstract semantics (for example, in the semantics of function abstraction), we will introduce new type variables, with the meaning "any possible type". During the evaluation (for example, of the function body), this information will be subject to instantiations, computed by unifications and represented as an idempotent substitution. Since the abstract semantic evaluation functions are defined by structural recursion, the easiest way to provide the instantiation information to the caller is to include it in the returned value.

Although we have now all necessary information for defining an adequate abstract domain, there is a problem concerning the representation of effects in the constraints. Intuitively we can simply represent them by using a set of pairs, where the first component is the predicate and the second one is the denotable value. The problem is in partial order, since we should consider both set inclusion and the relative precision of denotable values. We can achieve this by using power domains [14, 24]. We use a different approach: we define an effect as a function from predicates to denotable values (we will name it correspondence function). We can then represent constraints by splitting them in two parts: the first part is a set of pairs (annotation variable, predicate) and the second one is a correspondence function.

Let V_t be a countable set of type variables, V_a be a countable set of annotation variables, Ide be a countable set of identifier variable ($V_t \cap V_a \cap Ide = \emptyset$) and $\Sigma = \{unit : 0, int : 0, string : 0, xml : 1, link : 1, form : 1, list : 1, fun : 5\} \cup \{tuple_n : n \mid n \geq 2\}$ be a numerable set of function symbol, T_s is the set of terms with variables $V_t \cup V_a \cup Ide$ modulo renaming, ordered by the inverse instance relation. It is worth noting that we have introduced two new types *form* and *link* in order to solve the problem relating forms and links which we described in Section 3. Furthermore we will use annotation variables in *xml*, *link*, *form* and *fun* only; in *fun* there are two annotation variables representing the precondition and the post-condition respectively. We further assume that the first argument of *fun* is an identifier variable. We obtain *TypeS* by lifting T_s with idempotent substitutions [13] and by adding a new bottom element *Notype*.

As we described above, the first part of a constraint is a pair (annotation variable, predicate): (δ, q) means that the predicate q is in the effect represented by the variable δ . We use inverse inclusion as partial order: if C_1 is included in C_2 , then C_1 has less information than C_2 , hence, its value is less precise. Let V_a be the set of annotation variables and $Pred$ be the set of predicates. We define $Constr = \wp(V_a \times Pred)$. The second part of a constraint is a correspondence function whose domain is $TPred = Pred \rightarrow Dval$ ordered by using the dual of usual partial order. We assume that $cb : \wp(TPred) \rightarrow TPred$ is the glb operator and ζ is the bottom element.

The domain of abstract values is $TypeA = TypeS \times Dval \times Constr \times TPred$. In the following, we will denote by *Error* the bottom element of this domain.

The domain of abstract environment (type environment) is $AEnv = Ide \rightarrow TypeA$. We are now in the position to define our abstract domains $AV = AEnv \rightarrow EEnv \rightarrow TypeA$ for values and $AE = AEnv \rightarrow EEnv \rightarrow (TypeA \times EEnv)$ for expressions.

To relate the abstract domain to the concrete one we need to define a Galois connection. In [11] we formally built this connection in in various steps, by using properly defined representation functions [22] and propositions.

Some examples of abstract semantic equations are shown in Figures 5, 6 and 7. In these definitions, we assume to have a function *mgu*, which, given a set of term equations, computes a solution by using the unification algorithm. If there exists a solution, it returns the unifier $S(\theta)$; otherwise, it returns F to denote failure. The set of equations is denoted by $\{t_1 = t'_1, \dots, t_n = t'_n\}$. Since idempotent substitutions

$$\begin{aligned}
\mathcal{V} \llbracket \text{href}(\mathbf{E}) \rrbracket^a \rho \phi = & \\
& \gamma \in V_a \text{ fresh} \\
& \text{let } ((ts, -, C, f), \phi') = \llbracket \mathbf{E} \rrbracket^a \rho \phi \text{ in} \\
& \text{let } A = \text{assert}(\phi', \phi) \text{ in} \\
& \text{let } E = \text{event}(\phi', \phi) \text{ in} \\
& \text{if } E = \emptyset \wedge ts \neq \text{NoType} \text{ then} \\
& \quad \text{case mgu}(\{ts.t = \text{xml}(\gamma)\} \cup ts.\theta) \text{ of} \\
& \quad \quad S(\theta) \rightarrow \text{let } C' = C \cup \{(\gamma, q) \mid q \in A\} \\
& \quad \quad \quad \text{let } f' = \text{cb}\{f, \text{eenvToTPred}(\text{diff}(\phi', \phi))\} \text{ in} \\
& \quad \quad \quad ((\theta(\text{link}(\gamma)), \theta), \text{nodval}, \theta(C'), \theta(f')) \\
& \quad \quad \quad _ \rightarrow \text{Error} \\
& \text{else} \\
& \quad \text{Error} \\
\mathcal{V} \llbracket \lambda x. \mathbf{E} \rrbracket^a \rho \phi = & \\
& \alpha \in V_t \quad \gamma_1, \gamma_2 \in V_a \text{ fresh} \quad \varepsilon \text{ identity substitution} \\
& \text{let } ((ts, -, C_1, f_1), \phi') = \llbracket \mathbf{E} \rrbracket^a \rho [((\alpha, \varepsilon), \text{var}(x), \emptyset, \zeta)/x] \phi \text{ in} \\
& \text{if } ts \neq \text{NoType} \text{ then} \\
& \quad \text{let } \phi_d = \theta(\phi) \\
& \quad \text{let } C' = \{(\gamma_1, q) \mid q \in \text{assert}(\phi', \phi_d)\} \text{ in} \\
& \quad \text{let } C'' = \{(\gamma_2, q) \mid q \in \text{event}(\phi', \phi_d)\} \text{ in} \\
& \quad \text{let } f_2 = \text{eenvToTPred}(\text{diff}(\phi', \phi_d)) \text{ in} \\
& \quad \quad ((\theta(\text{fun}(x, \alpha, \gamma_1, ts.t, \gamma_2)), \theta), \\
& \quad \quad \quad \text{nodval}, \theta(C_1 \cup C' \cup C''), \theta(\text{cb}\{f_1, f_2\})) \\
& \text{else} \\
& \quad \text{Error}
\end{aligned}$$

Figure 5: The abstract semantics of links and functional abstractions.

are isomorphic to solved form equations, we will use $\{t_1 = t'_1, \dots, t_n = t'_n\} \cup \theta$ to refer the union of equations in $\{t_1 = t'_1, \dots, t_n = t'_n\}$ and equations defined by θ . For the sake of simplicity, the components of the elements of the domain $TypeS$, will be identified by a notation similar to the one used to access the fields of a structure in an imperative language. Given $ts = (t', \theta') \in TypeS$, then $ts.t = t'$ and $ts.\theta = \theta'$.

Given an element C of $Constr$ and a substitution θ , we will denote by $\theta(C) = \{(\theta(\delta), l) \mid (\delta, l) \in C\}$ the pair obtained by applying θ to all the annotation variables in C .

Given a correspondence function $f \in TPred$ and a substitution θ , we define $\theta(f) = \lambda q. \theta(f q)$, where if $d \neq \text{var}(x)$ for some x then $\theta(d) = d$.

Furthermore we assume that for $f \in TPred$ and $C \in Constr$ $f \downarrow C$ and $f \leftarrow C$ are the correspondence functions achieved by removing from f the predicates occurring and not occurring in C respectively; that

$$\begin{aligned}
\llbracket \text{get}(V) \rrbracket^a \rho \phi = & \\
& \gamma \in V_a \text{ fresh} \\
& \text{let } (ts, d, C, f) = \mathcal{V} \llbracket V \rrbracket^a \rho \phi \text{ in} \\
& \text{if } ts \neq \text{NoType} \text{ then} \\
& \quad \text{case mgu}(\{ts.t = \text{link}(\gamma)\} \cup ts.\theta) \text{ of} \\
& \quad \quad S(\theta) \rightarrow \text{let } C' = \{(\theta(\gamma), q) \in \theta(C)\} \text{ in} \\
& \quad \quad \quad \text{if } \text{check}(\theta(f \leftarrow C'), \phi) \text{ then} \\
& \quad \quad \quad \quad ((\theta(\text{xml}(\gamma)), \theta), \\
& \quad \quad \quad \quad \quad \text{nodval}, \theta(C) \setminus C', \theta(f \downarrow C)), \phi) \\
& \quad \quad \quad \text{else} \\
& \quad \quad \quad \quad (\text{Error}, \iota) \\
& \quad \quad \quad \quad \rightarrow (\text{Error}, \iota) \\
& \quad \text{else} \\
& \quad \quad (\text{Error}, \iota) \\
\llbracket E_1 E_2 \rrbracket^a \rho \phi = & \\
& x \in \text{Ide } \alpha_1 \in V_t \ \gamma_1, \gamma_2 \in V_a \text{ fresh} \\
& \text{let } ((ts_1, -, C_1, f_1), \phi_1) = \llbracket E_1 \rrbracket^a \rho \phi \\
& \text{let } ((ts_2, d_2, C_2, f_2), \phi_2) = \llbracket E_2 \rrbracket^a \rho \phi_1 \\
& \text{if } ts_1 \neq \text{NoType} \wedge ts_2 \neq \text{NoType} \text{ then} \\
& \quad \text{case mgu}(\{ts_1.t = \text{fun}(x, \alpha, \gamma_1, ts_2.t, \gamma_2)\} \cup \\
& \quad \quad \cup ts_1.\theta \cup ts_2.\theta) \text{ of} \\
& \quad \quad S(\theta) \rightarrow \text{let } C' = \{(\theta(\gamma_1), q) \in \theta(C_1)\} \text{ in} \\
& \quad \quad \quad \text{let } C'' = \{(\theta(\gamma_2), q) \in \theta(C_2)\} \text{ in} \\
& \quad \quad \quad \quad \text{let } f'_1 = \theta(f_1)[\theta(x), d_2] \text{ in} \\
& \quad \quad \quad \quad \text{if } \text{check}(\theta(f'_1 \leftarrow C'), \phi_2) \text{ then} \\
& \quad \quad \quad \quad \quad ((\theta(ts_2.t), \theta), \top, \theta(C_1 \cup C_2) \setminus (C' \cup C''), \\
& \quad \quad \quad \quad \quad \quad \text{cb}\{\theta(f_1) \downarrow (C' \cup C''), \theta(f_2)\}, \text{incl}(\phi_2, (\theta(f'_1) \leftarrow C''))) \\
& \quad \quad \quad \quad \text{else} \\
& \quad \quad \quad \quad \quad (\text{Error}, \iota) \\
& \quad \quad \quad \quad \rightarrow (\text{Error}, \iota) \\
& \quad \text{else} \\
& \quad \quad (\text{Error}, \iota)
\end{aligned}$$

Figure 6: The abstract semantics of get expression and function application.

$f[x, d]$ for $x \in \text{Ide}$, $d \in D\text{Val}$ and $f \in TPred$ is the correspondence function achieved by binding d to all predicates which are bound to $\text{var}(x)$ in f ; that given a $f \in TPred$ and $\phi \in EEnv$ the function $\text{check}(f, \phi)$

returns *true* if the events represented by f have been occurred in ϕ , *false* otherwise; that for $\phi \in EEnv$ $eenvToTPred(\phi)$ is the correspondence function achieved from ϕ ; that given ϕ_1, ϕ_2 $assert(\phi_2, \phi_1)$ is the set of predicates of events asserted in ϕ_2 but not in ϕ_1 , that $event(\phi_2, \phi_1)$ is the set of predicates of events generated in ϕ_2 but not in ϕ_1 and that $diff(\phi_2, \phi_1)$ is the events environment which contains the events of ϕ_2 which are not in ϕ_1 and the events of ϕ_1 which changed their value or state in ϕ_2 .

$$\begin{aligned}
& \llbracket \text{assert } q(V) \rrbracket^a \rho \phi = \\
& \quad \text{let } (ts, d, C, f) = \mathcal{V} \llbracket V \rrbracket^a \rho \phi \text{ in} \\
& \quad \text{if } ts \neq NoType \wedge (d = nint(n) \vee d = var(x)) \text{ then} \\
& \quad \quad \text{case mgu}(\{ ts.t = int \} \cup ts.\theta) \text{ of} \\
& \quad \quad \quad S(\theta) \rightarrow \text{if } q \notin dom(\phi) \vee \pi_1(\phi(q)) = d \text{ then} \\
& \quad \quad \quad \quad ((unit, \theta), nodval, \theta(C), \theta(f)), \phi[(d, A)/q] \\
& \quad \quad \quad \text{else} \\
& \quad \quad \quad \quad (Error, \iota) \\
& \quad \quad \quad \rightarrow (Error, \iota) \\
& \quad \text{else} \\
& \quad \quad (Error, \iota) \\
& \llbracket \text{event } q(V) \rrbracket^a \rho \phi = \\
& \quad \text{let } (ts, d, C, f) = \mathcal{V} \llbracket V \rrbracket^a \rho \phi \text{ in} \\
& \quad \text{if } ts \neq NoType \wedge (d = nint(n) \vee d = var(x)) \text{ then} \\
& \quad \quad \text{case mgu}(\{ ts.t = int \} \cup ts.\theta) \text{ of} \\
& \quad \quad \quad S(\theta) \rightarrow \text{if } q \notin dom(\phi) \vee \phi(q) = (d, T) \text{ then} \\
& \quad \quad \quad \quad ((unit, \theta), nodval, \theta(C), \theta(f)), \phi[(d, E)/q] \\
& \quad \quad \quad \text{else} \\
& \quad \quad \quad \quad (Error, \iota) \\
& \quad \quad \quad \rightarrow (Error, \iota) \\
& \quad \text{else} \\
& \quad \quad (Error, \iota)
\end{aligned}$$

for some $n \in \mathcal{L}$, $x \in Ide$ and where $T \in \{E, EA\}$

Figure 7: The abstract semantics of event and assert annotations.

The semantics of links consists in the evaluation of the expression E . If in this evaluation no errors ($ts \neq NoType$) and no new events (this is required by the rule described in Section 2) occur, then we check that the computed value has type *xml*. Since in our reconstruction *xml*, *link* and *form* are different types without any relation, this check rejects all the expressions which return a value of type *link* or *form*. Although this behavior may seem too restrictive, because it rejects some legal expressions like

`href(href(Text("Hello")))`, it guarantees us safety and simplicity in the management of these different and unrelated types. If this check has success, we return an abstract value where the simple type is *link* and the constraint is risen by properly extending the result of the evaluation of `E`.

The semantics of `forms` is similar. We evaluate `E` in a type environment where the labels `ll` are bound to the abstract value with simple type *string* and constraint empty and we return an abstract value where the simple type is *form*.

The semantics of functional abstraction consists in the evaluation of the body `E` in a type environment, where the formal parameter `x` is bound to a generic type. If in this evaluation no errors occur, we compute the events which are included in the precondition (represented by C' and f_2) and in the post-condition (represented by C'' and f_2). We return an abstract value where the simple type is obtained by applying the substitution $ts.\theta$ to the functional type $(fun(x, \alpha, \gamma_1, ts.t, \gamma_2))$ and the constraint is obtained by combining C with C' and C'' and f with f_2 .

The semantics of `get` requires the evaluation of `V` to be successful and yields a value of type *link*. If the preconditions are satisfied, that is if they are in ϕ and have occurred before, we construct an abstract value where the simple type is *xml* and the constraint is obtained from the one returned by the evaluation of `V` by removing the information about preconditions. The pair which is returned has in the first component this abstract value and in the second one the events environment ϕ . This is correct because the semantics of `href` guarantees that no new events have occurred during the evaluation of the suspended expression.

The semantics of `post` is similar except that we ask that the elements of list `VL` are strings and that the value yielded by the evaluation of `V` has type *form*.

In the semantics of function application we evaluate the sub-expressions `E1` and `E2`: if both evaluations do not produce errors, we check that the simple type of `E1` is a function type where the argument has the simple type of `E2` and that the precondition of function is satisfied in the events environment ϕ_2 , obtained from evaluating both the sub-expressions. In order to perform this last check, we substitute the denotable value bound to `x` in f_1 by the one returned by the evaluation of `E2`. Then, by using the function *check*, we ask that the events required by the function body are in ϕ_2 . If we succeed, we construct an abstract value where the simple type is $\theta(\alpha)$ and the constraint is obtained by composing those returned by the evaluation of the sub-expressions, where the events of preconditions and post-conditions are removed. We return a pair composed by this abstract value and by the events environment ϕ_2 extended with the events of the post-condition of the function.

The semantics of `assert` consists in the evaluation of `V`. If it yields an abstract value whose simple type is *int* and whose denotable value is a specific integer or a specific identifier, we check that there is in ϕ at most the same event which we are generating. In this way we are sure that it is impossible to change the value bound to a predicate. If this check has success, we build an abstract value where the simple type is *unit* and the constraint is the one returned by the evaluation of `V`. This abstract value is the first component of returned pair; the second component consists of the events environment ϕ extended with the new event.

The semantics of `event` is similar except that we ask that, if the event is in ϕ , then its state has to be either *E* or *EA*.

5 Implementation and Examples

Both the concrete and the abstract semantics have been implemented as OCaml [16] programs. The language provides a feature, the mechanism of functors, which allows us to have a unique semantic

function (realised by the functor `Semantics`), parametrized with respect to the primitive operations and the semantic domain. We can thus construct the concrete semantics interpreter, which executes programs, and the abstract interpreter, which analyzes programs in terms of types and effects, by instantiating the same functor `Semantics`.

Programs are represented in abstract syntax, although, for the sake of simplicity, we will use in the following LINKS-like syntax. For example, the expression

```
fun buy(value, dbpass) {
  var _ = assert PriceIs(value);
  Text("Hello")
}
```

defines a function which requires that the event `PriceIs(value)` has occurred and which returns an XML value. The result of its evaluation by the abstract semantics interpreter is

```
(type - :
  Function(_#value#var0_, Integer(), _annvar0_,
    Function(_#dbpass#var1_, _typevar1_, _annvar2_,
      Xml(_annvar4_), _annvar3_),
    _annvar1_)
  No_dval [(_annvar2_,PriceIs)] {PriceIs -> _#value#var0_}, {})
```

meaning that the computed type is a function type whose first argument has a type integer and the second one has type variable⁶ where the precondition (represented by the annotation variable `_annvar2_`) includes the event composed by the predicate `PriceIs` and the value bound to the first formal parameter. If we give a value (for example 5) to the first parameter, the abstract semantics is

```
(type - :
  Function(_#dbpass#var3_, _typevar3_, _annvar7_,
    Xml(_annvar9_), _annvar8_)
  Unknown [(_annvar7_,PriceIs)] {PriceIs -> 5}, {})
```

that is the computed type is a specialization of that one computed for `buy` where the predicate `PriceIs` is bound to the value 5 in the precondition. The abstract semantics of the application of the function `buy` to 5 and "a" is an error

```
Exception: No_type "apply_fun: no preconditions"
```

because we are applying a function whose precondition is not satisfied.

6 Conclusions

We have described how to reconstruct a types-and-effects system, proposed to handle some security issues in LINKS, as an abstract interpretation of a denotational semantics which explicitly models the types and the effects. By our reconstruction we have precisely defined the relation between the semantics and the analysis, we have systematically constructed a correct analyser and we have shown that the proposed types-and-effects system was not sound. We have stressed that the unsoundness derived from the fact of considering forms and links as simple XML values forgetting their own differentiating features. In our reconstruction we have solved this problem by using two new specific types and we have managed them in ad-hoc manner. We plan to extend our reconstruction to consider a type system with sub-types so as

⁶since the `dbpass` parameter is not used in the body, the analyzer cannot compute a more precise type

to be able to manage links and forms in a more uniform and elegant way and to use additional values in the effects.

One advantage of abstract interpretation approach on the type system approach is that the analysis is directly derived from the semantics and is sound by construction. This forces one to tackle from the very beginning subtle problems such as the ones described in Section 3 that might only be revealed while trying to prove the soundness theorem following the type system approach. On the other hand we have shown that abstract interpretation can easily handle extensions of types, such as types and effects. There is only one example in the literature of an abstract interpretation reconstruction of a type and effect static analysis [25].

References

- [1] F. Baader & J. H. Siekmann (1994): *Unification theory*. In Dov M. Gabbay, Christopher J. Hogger, J. A. Robinson & Jörg H. Siekmann, editors: *Handbook of Logic in Artificial Intelligence and Logic Programming (2)*. Oxford University Press, pp. 41–126.
- [2] I. G. Baltopoulos & A. D. Gordon (2009): *Secure Compilation of a Multi-Tier Web Language*. In: *TLDI '09: Proceedings of the 4th international workshop on Types in language design and implementation*. ACM, New York, NY, USA, pp. 27–38.
- [3] J. Bengtson, K. Bhargavan, C. Fournet, A. D. Gordon & S. Maffei (2008): *Refinement Types for Secure Implementations*. In: *Proceedings of the 2008 21st IEEE Computer Security Foundations Symposium*. IEEE Computer Society, Washington, DC, USA, pp. 17–32.
- [4] E. Cooper, S. Lindley, P. Wadler & J. Yallop (2006): *Links: Web programming without tiers*. In: *In 5th International Symposium on Formal Methods for Components and Objects (FMCO)*. Springer-Verlag.
- [5] P. Cousot (1997): *Types as Abstract Interpretations, invited paper*. In: *Conference Record of the Twenty-fourth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM Press, New York, NY, Paris, France, pp. 316–331.
- [6] P. Cousot & R. Cousot (1977): *Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints*. In: *Conference Record of the Fourth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM Press, New York, NY, Los Angeles, California, pp. 238–252.
- [7] P. Cousot & R. Cousot (1979): *Systematic design of program analysis frameworks*. In: *Conference Record of the Sixth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM Press, New York, NY, San Antonio, Texas, pp. 269–282.
- [8] P. Cousot & R. Cousot (1992): *Abstract Interpretation and Application to Logic Programs*. *Journal of Logic Programming* 13(2–3), pp. 103–179.
- [9] P. Cousot & R. Cousot (1992): *Abstract Interpretation Frameworks*. *Journal of Logic and Computation* 2(4), pp. 511–547.
- [10] L. Damas & R. Milner (1982): *Principal type-schemes for functional programs*. In: *POPL '82: Proceedings of the 9th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*. ACM, New York, NY, USA, pp. 207–212.
- [11] L. Galletta (2010): *Una semantica astratta per l'inferenza dei tipi ed effetti in un linguaggio multi-tier*. Master's thesis, Università di Pisa.
- [12] D. K. Gifford & J. M. Lucassen (1986): *Integrating functional and imperative programming*. In: *Proceedings of the 1986 ACM conference on LISP and functional programming*. LFP '86, ACM, New York, NY, USA, pp. 28–38.

- [13] R. Gori & G. Levi (2002): *An Experiment in Type Inference and Verification by Abstract Interpretation*. In: *VMCAI '02: Revised Papers from the Third International Workshop on Verification, Model Checking, and Abstract Interpretation*. Springer-Verlag, London, UK, pp. 225–239.
- [14] C. A. Gunter & D. S. Scott (1990): *Semantic Domains*, chapter 12, pp. 634–674. *Handbook of Theoretical Computer Science*, Elsevier Science.
- [15] R. Hindley (1969): *The Principal Type-Scheme of an Object in Combinatory Logic*. *Transactions of the American Mathematical Society* 146, pp. 29–60.
- [16] INRIA: *The Caml Language*. Available at <http://caml.inria.fr>. WWW publication.
- [17] J.-L. Lassez, M. J. Maher & K. Marriott (1988): *Unification revisited*. In: *Foundations of deductive databases and logic programming*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, pp. 587–625.
- [18] X. Leroy & F. Pessaux (2000): *Type-based analysis of uncaught exceptions*. *ACM Trans. Program. Lang. Syst.* 22, pp. 340–377.
- [19] E. Meijer, W. Schulte & G. Bierman (2003): *Programming with circles, triangles and rectangles*. In: *In XML Conference and Exposition*.
- [20] B. Monsuez (1992): *Polymorphic Typing by Abstract Interpretation*. In: *Proceedings of the 12th Conference on Foundations of Software Technology and Theoretical Computer Science*. Springer-Verlag, London, UK, pp. 217–228.
- [21] F. Nielson & H. Nielson (1994): *Constraints for polymorphic behaviours of concurrent ML*. In Jean-Pierre Jouannaud, editor: *Constraints in Computational Logics*. *Lecture Notes in Computer Science* 845, Springer Berlin / Heidelberg, pp. 73–88. 10.1007/BFb0016845.
- [22] F. Nielson, H. Riis Nielson & C. Hankin (2005): *Principles of Program Analysis*, 1st ed. 1999. corr. 2nd printing, 1999 edition. Springer.
- [23] C. Queinnec (2000): *The influence of browsers on evaluators or, continuations to program web servers*. *SIGPLAN Not.* 35, pp. 23–33.
- [24] D. Schmidt (1986): *Denotational Semantics: A Methodology for Language Development*. William C Brown Pub.
- [25] J. Vouillon & P. Jouvelot (1995): *Type and Effect Systems via Abstract Interpretation*. Available at <http://www.cri.ensmp.fr/classement/doc/A-273.pdf>.
- [26] G. Winskel (1993): *The Formal Semantics of Programming Languages*. MIT Press.