

Verifying Pointer and String Analyses with Region Type Systems

Lennart Beringer*

Robert Grabowski[†]

Martin Hofmann[‡]

Abstract

Pointer analysis statically approximates the heap pointer structure during a program execution in order to track heap objects or to establish alias relations between references, and usually contributes to other analyses or code optimizations. In recent years, a number of algorithms have been presented that provide an efficient, scalable, and yet precise pointer analysis. However, it is unclear how the results of these algorithms compare to each other semantically.

In this paper, we present a general region type system for a Java-like language and give a formal soundness proof. The system is subsequently specialized to obtain a platform for embedding the results of various existing context-sensitive pointer analysis algorithms, thereby equipping the computed relations with a common interpretation and verification. We illustrate our system by outlining an extension to a string value analysis that builds on pointer information.

1 Introduction

Pointer (or points-to) analysis is a static program analysis technique that determines an over-approximation of possible points-to relations that occur during the execution of a program. More precisely, it chooses an abstraction of the pointers and references, and computes which pointers may possibly point to which data. A conservative approximation of this structure is also an alias analysis, as the computed points-to relation directly includes the information which pointers may point to the same object. A pointer analysis is often used for compiler optimizations, but may also serve as the basis for other analyses, such as the computation of possible string values in order to prevent string-based security holes.

There exists a large number of pointer analysis algorithms [13, 1, 28] for different languages. Each algorithm faces the trade-off between precision and efficiency of the analysis: it should choose the right abstractions in order to produce as much useful information as possible while at the same time being able to process large code bases in a reasonable time. These analyses have different techniques, implementations, and complexities. Especially BDD-based algorithms [7, 8] have been shown to be very efficient and precise at the same time.

While several of these analyses also consider soundness, it appears that there does not yet exist a uniformly agreed-upon formal framework that encompasses the interpretations of at least a substantial subset of the analyses. We argue that such a unifying treatment is important, for theoretical as well as pragmatic and practical reasons. First, it is a basis for fair comparisons regarding the precision, flexibility or expressivity of different analyses, the theoretical complexity of the associated algorithms, and their experimental evaluation on common benchmarks. Second, once the analyses agree on the formal property they guarantee, we can safely replace one analysis by another in a compiler or verification tool. Third, a uniform guarantee provides the basis for the formal verification of security-relevant properties that rely on pointer analysis results, as is required in proof-carrying code scenarios.

The first purpose of the present paper is to provide such a framework for Java-like languages, given by a hierarchy of region-based type systems for a language in the style of Featherweight Java [17]. Uniformity (i.e. semantic agreement) is guaranteed by equipping the bottom-most layer in the hierarchy

*Princeton University, Princeton, NJ, USA, beringer@tcs.ifi.lmu.de

[†]Ludwig-Maximilians-Universität, München, Germany, robert.grabowski@ifi.lmu.de

[‡]Ludwig-Maximilians-Universität, München, Germany, hofmann@ifi.lmu.de

with a formal interpretation and soundness result, and by deriving the higher levels by specializing this bottom-most layer to move towards calculi representing concrete analyses.

Second, we demonstrate that a number of existing pointer analyses for object-oriented programs are based on abstraction disciplines that arise as specialisations of a generic parametrized refinement of our base-level type system. We focus on disciplines that specify the abstraction of references and execution points [25], and therefore enable different forms of field-sensitive and context-sensitive analyses. For example, objects may be abstracted to their allocation sites and their class, and execution points may be abstracted by receiver-object or call-site stack contexts.

As one moves higher in the hierarchy, different choices for these abstractions regarding expressivity arise, corresponding to the above abstraction disciplines, and decidability issues become more prominent, in particular the question of algorithmic type checking. Our third contribution consists of showing how the parametrized type system can be reformulated algorithmically to yield a type checking algorithm. Thanks to the hierarchical structure of our framework, we thus immediately obtain algorithms for automatically validating the correctness of the results of concrete analyses, as long as the results have been interpreted in the framework by instantiating its parameters. As we only consider the results, our verification is also independent of implementation details of concrete analysis algorithms.

Finally, we apply our framework to the analysis of string values, in order to lay the foundations for eliminating SQL injections and cross-site scripting attacks. We extend the language to a simple string model, and outline how data flow analyses for possible string values that build on pointer analyses [9, 11] can be verified with the correspondingly extended type system.

Related work We include concepts of context-sensitivity to distinguish different analyses of the same method implementation. In particular, k -CFA [27] is a higher-order control flow analysis where contexts are call stack abstractions of finite length k . The k -CFA mainly addresses control flows in functional languages, where functions are first-class values. It requires a combination of value (data flow) analysis and control flow analysis to approximate the possible lambda abstractions that an expression may evaluate to. The similar dynamic dispatch problem for methods in object-oriented languages is easier to solve, as the possible implementation targets of a method invocation can be retrieved from the class information. k -CFA has been extended with polyvariant types for functions [2], such that different types can be used for the function at different application sites. In our type system, we borrow this concept of polyvariance.

Hardekopf and Lin [15] transform variables that are not address-taken into SSA form, such that running a flow-insensitive analysis on these converted variables has the effect of running a flow-sensitive analysis on the original variables. Our framework assumes SSA-transformed code input, presented in the form of a functional program. Identifying opportunities for SSA transformation, which is a central concern of [15], can thus be seen as a preprocessing phase for our framework to apply.

This paper uses “regions” in the sense of Lucassen and Gifford [21], i.e. as representations of disjoint sets of memory locations. Equivalently, regions thus partition or color the memory. In literature, this disjointness is used to show that certain memory manipulations do not influence other parts of a program, in order to e.g. show semantic equivalences [4], to enable a safe garbage collection [29], or to infer properties for single objects by tracking unique objects in a region [10, 14]. In the pointer analysis setting presented here, regions are simply seen as abstract memory locations that summarize one or more concrete locations, thereby helping to discover may-alias relations. We do not consider uniqueness or must-alias information, and do not aim to justify garbage collection.

Paddle [19] and Alias Analysis Library [18] are implementation frameworks that embed concrete pointer analyzers by factoring out different parameters, as is done here. However, the works do not aim at a formal soundness result. Indeed, they embed the *algorithms* into a common *implementation* framework, while our type system embeds the *results* of such algorithms into a common *semantic* framework.

Synopsis The next section introduces the FJEU language and its semantics. We introduce the base-level region type system and the soundness proof in section 3. In section 4, we specialize the type system to a parametrized version, such that abstraction principles found in pointer analyses can be modelled explicitly as instantiations of the parameters, and outline a type-checking algorithm. Finally, section 5 extends the system, such that results of a string analysis based on pointer analysis can also be verified.

2 Featherweight Java with Updates

We examine programs of the language FJEU [16], a simplified formal model of the sequential fragment of Java that is relevant for pointer analysis. FJEU extends Featherweight Java (FJ) [17] with attribute updates, such that programs may have side effects on a heap. Object constructors do not take arguments, but initialize all fields with *null*, as they can be updated later. Also, the language adds `let` constructs and conditionals to FJ.

2.1 Preliminaries

We write $\mathcal{P}^{fin}(X)$ for the set of all finite subsets of X , and $\mathcal{P}(X)$ for the set of all finite and infinite subsets of X . The notations $A \rightarrow B$ and $A \dashrightarrow B$ stand for the set of total and partial functions from A to B , respectively. We write $[x \mapsto v]$ for the partial function that maps x to v and is else undefined. The function $f[x \mapsto v]$ is equal to f , except that it returns v for x . Both function notations may be used in an indexed fashion, e.g. $f[x_i \mapsto v_i]_{\{1, \dots, n\}}$, to define multiple values. In typing rules, we sometimes write $x : v$ and $f, x : v$ corresponding to the above notation. Finally, we write \bar{a} for a sequence of entities a .

2.2 Syntax

The following table summarizes the (infinite) abstract identifier sets in the language, the meta-variables we use to range over them, and the syntax of FJEU expressions:

$$\begin{array}{ll}
 \text{variables: } x, y \in \mathcal{X} & \text{classes: } C, D, E \in \mathcal{C} \\
 \text{fields: } f \in \mathcal{F} & \text{methods: } m \in \mathcal{M} \\
 \mathcal{E} \ni e ::= & \text{null} \mid x \mid \text{new } C \mid \text{let } x = e \text{ in } e \mid x.f \mid x.f := y \mid x.m(\bar{y}) \mid \\
 & \text{if } x \text{ instanceof } E \text{ then } e \text{ else } e \mid \text{if } x = y \text{ then } e \text{ else } e
 \end{array}$$

To keep our calculus minimal and focussed on pointer alias analysis, we omit primitive data types such as integers or booleans. However, such data types and their operations can easily be added to the language. We also omit type casts found in FJ, as they do not provide new insights to pointer analysis. Nevertheless, casts can be included in a straight-forward manner. In order to simplify the proofs, we require programs to be in let normal form. The somewhat unusual conditional constructs for dynamic class tests and value equality are included to have reasonable if-then-else expressions in the language while avoiding the introduction of booleans.

An FJEU program is defined by the following relations and functions:

$$\begin{array}{ll}
 \text{subclass relation: } & \prec \in \mathcal{P}^{fin}(\mathcal{C} \times \mathcal{C}) \\
 \text{field list: } & \text{fields} \in \mathcal{C} \rightarrow \mathcal{P}^{fin}(\mathcal{F}) \\
 \text{method list: } & \text{methods} \in \mathcal{C} \rightarrow \mathcal{P}^{fin}(\mathcal{M}) \\
 \text{method table: } & \text{mtable} \in \mathcal{C} \times \mathcal{M} \rightarrow \mathcal{E} \\
 \text{FJEU program: } & P = (\prec, \text{fields}, \text{methods}, \text{mtable})
 \end{array}$$

FJEU is a language with nominal subtyping: $D \prec C$ means D is an immediate subclass of C . The relation is well-formed if it is a tree successor relation; multiple inheritance is not allowed. We write

\preceq for the reflexive and transitive hull of \prec . The functions *fields* and *methods* describe for each class C which fields and method objects of that class have. The functions are well-formed if for all classes C and D such that $D \preceq C$, $\text{fields}(C) \subseteq \text{fields}(D)$ and $\text{methods}(C) \subseteq \text{methods}(D)$, i.e. classes inherit fields and methods from their superclasses. A method table *mtable* gives for each class and each method identifier its implementation, i.e. the FJEU expression that forms the method's body. To simplify the presentation, we assume that formal argument variables in the body of a method m are named x_1^m, x_2^m , etc., abbreviated to $\overline{x^m}$, besides the implicit and reserved variable *this*. All free variables of an implementation of m must be from the set $\{\text{this}, x_1^m, x_2^m, \dots\}$. A method table is well-formed if $\text{mtable}(C, m)$ is defined whenever $m \in \text{methods}(C)$. In other words, all methods declared by *methods* must be implemented, though the implementation may be overridden in subclasses for the same number of formal parameters. In the following, we assume a fixed FJEU program P whose components are all well-formed.

2.3 Semantics

A state consists of a store (variable environment or stack) and a heap (memory). Stores map variables to values, while heaps map locations to objects. An object consists of a class identifier and a valuation of its fields. The only kind of values in FJEU are locations and *null* references.

$$\begin{array}{ll} \text{locations: } l \in \mathcal{L} & \text{stores: } s \in \mathcal{X} \rightarrow \mathcal{V} \\ \text{values: } v \in \mathcal{V} = \mathcal{L} \cup \{\text{null}\} & \text{heaps: } h, k \in \mathcal{L} \rightarrow \mathcal{O} \\ \text{objects: } \mathcal{O} = \mathcal{C} \times (\mathcal{F} \rightarrow \mathcal{V}) \end{array}$$

The semantics of FJEU is defined as a standard big-step relation $(s, h) \vdash e \Downarrow v, h'$, which means that an FJEU expression e evaluates in store s and heap h to the value v and modifies the heap to h' . Figure 1 shows the defining rules of the operational semantics. A premise involving a partial function, like $s(x) = l$, implies the side condition $x \in \text{dom}(s)$.

2.4 Class tables

A class table $\mathfrak{C}_0 = (A_0, M_0)$ models FJ's standard type system, where types are simply classes. The *field typing* $A_0 : (\mathcal{C} \times \mathcal{F}) \rightarrow \mathcal{C}$ assigns to each class C and each field $f \in \text{fields}(C)$ the class of the field. The field class is required to be invariant with respect to subclasses of C . The *method typing* $M_0 : (\mathcal{C} \times \mathcal{M}) \rightarrow \overline{\mathcal{C}} \times \mathcal{C}$ assigns to each class C and each method $m \in \text{methods}(C)$ a *method type*, which specifies the classes of the formal argument variables and of the result value. It is required to be contravariant in the argument classes and covariant in the result class with respect to subclasses of C .

2.5 Example program

As a very small example, look at the Java program in figure 2 that implements a list via `Nil` and `Cons` classes, with a `copy` method that duplicates the list structure, but not the actual elements. The same program in FJEU and its class table is shown in figure 3, where ε denotes the empty sequence.

3 Region Type System

In this section we define the base region type system, which serves as a main unifying calculus for pointer analysis and is given an interpretation and soundness proof. We assume an infinite set \mathcal{R} of *regions* r , which are abstract memory locations. Each region stands for zero or more concrete locations. Different regions represent disjoint sets of concrete locations, hence they partition or *color* the memory. Two pointers to different regions can therefore never alias.

$$\boxed{(s, h) \vdash e \Downarrow v, h'}$$

$$\frac{}{(s, h) \vdash \text{null} \Downarrow \text{null}, h} \quad \frac{}{(s, h) \vdash x \Downarrow s(x), h} \quad \frac{l \notin \text{dom}(h) \quad F = [f \mapsto \text{null}]_{f \in \text{fields}(C)}}{(s, h) \vdash \text{new } C \Downarrow l, h[l \mapsto (C, F)]}$$

$$\frac{(s, h) \vdash e_1 \Downarrow v_1, h_1 \quad (s[x \mapsto v_1], h_1) \vdash e_2 \Downarrow v_2, h_2}{(s, h) \vdash \text{let } x = e_1 \text{ in } e_2 \Downarrow v_2, h_2}$$

$$\frac{s(x) = l \quad h(l) = (C, -) \quad C \preceq E \quad (s, h) \vdash e_1 \Downarrow v, h'}{(s, h) \vdash \text{if } x \text{ instanceof } E \text{ then } e_1 \text{ else } e_2 \Downarrow v, h'}$$

$$\frac{s(x) = l \quad h(l) = (C, -) \quad C \not\preceq E \quad (s, h) \vdash e_2 \Downarrow v, h'}{(s, h) \vdash \text{if } x \text{ instanceof } E \text{ then } e_1 \text{ else } e_2 \Downarrow v, h'}$$

$$\frac{s(x) = s(y) \quad (s, h) \vdash e_1 \Downarrow v, h'}{(s, h) \vdash \text{if } x = y \text{ then } e_1 \text{ else } e_2 \Downarrow v, h'} \quad \frac{s(x) \neq s(y) \quad (s, h) \vdash e_2 \Downarrow v, h'}{(s, h) \vdash \text{if } x = y \text{ then } e_1 \text{ else } e_2 \Downarrow v, h'}$$

$$\frac{s(x) = l \quad h(l) = (C, F)}{(s, h) \vdash x.f \Downarrow F(f), h} \quad \frac{s(x) = l \quad h(l) = (C, F) \quad h' = h[l \mapsto (C, F[f \mapsto s(y)])]}{(s, h) \vdash x.f := y \Downarrow s(y), h'}$$

$$\frac{s(x) = l \quad h(l) = (C, -) \quad |\bar{x}^m| = |\bar{y}| = n \quad s' = [\text{this} \mapsto s(x)] \cup [x_i^m \mapsto s(y_i)]_{i \in \{1, \dots, m\}} \quad (s', h) \vdash \text{mtable}(C, m) \Downarrow v, h'}{(s, h) \vdash x.m(\bar{y}) \Downarrow v, h'}$$

Figure 1: Operational semantics of FJEU (extract)

```

class Elem {}
class List { List copy() { return new List(); } }
class Nil extends List { Nil copy() { return new Nil(); } }
class Cons extends List {
  Elem elem; List next;
  Cons copy() { Cons res = new Cons(); res.elem = elem;
              res.next = next.copy(); return res; } }

```

Figure 2: List copy example in Java

$P = (\preceq, \text{fields}, \text{methods}, \text{mtable})$	
$\preceq = \{(\text{Nil}, \text{List}), (\text{Cons}, \text{List})\}$	
$\text{fields}(\text{Elem}) = \emptyset$	$\text{methods}(\text{Elem}) = \emptyset$
$\text{fields}(\text{List}) = \emptyset$	$\text{methods}(\text{List}) = \{\text{copy}\}$
$\text{fields}(\text{Nil}) = \emptyset$	$\text{methods}(\text{Nil}) = \{\text{copy}\}$
$\text{fields}(\text{Cons}) = \{\text{elem}, \text{next}\}$	$\text{methods}(\text{Cons}) = \{\text{copy}\}$
$A_0(\text{Cons}, \text{elem}) = \text{Elem}$	$\text{mtable}(\text{List}, \text{copy}) = \text{new List}$
$A_0(\text{Cons}, \text{next}) = \text{List}$	$\text{mtable}(\text{Nil}, \text{copy}) = \text{new Nil}$
$M_0(\text{List}, \text{copy}) = (\varepsilon, \text{List})$	$\text{mtable}(\text{Cons}, \text{copy}) = \text{let } \text{res} = \text{new Cons in}$
$M_0(\text{Nil}, \text{copy}) = (\varepsilon, \text{Nil})$	$\text{let } \text{copyTail} = \text{next.copy}() \text{ in}$
$M_0(\text{Cons}, \text{copy}) = (\varepsilon, \text{Cons})$	$\text{let } _ = \text{res.elem} := \text{elem in}$
	$\text{let } _ = \text{res.next} := \text{copyTail in res}$

Figure 3: List copy example in FJEU

3.1 Refined types and subtyping

The region type system is a refinement of the plain type system: we equip classes with (possibly infinite) subsets from \mathcal{R} . For example, a location l is typed with the *refined type* $C_{\{r,s\}}$ if it points to an object of class C (or a subclass of C), and if l is abstracted to either r or s , but no other region. The *null* value can be given any type, while the type of locations must have a non-empty region set. The following table summarizes the definitions:

Regions:	$r, s, t \in \mathcal{R}$
Region sets:	$R, S, T \in \mathcal{P}(\mathcal{R})$
Refined types:	$\sigma, \tau \in \mathcal{T} = \mathcal{C} \times \mathcal{P}(\mathcal{R})$

In the following, we use the notation C_R instead of (C, R) for types. Though the region identifier s is already used for variable stores, the difference should be clear from the context. Since region sets are an over-approximation of the possible locations where an object resides, we can easily define a subtyping relation $<$: based on set inclusion:

$$C_R <: D_S \iff R \subseteq S \wedge C \preceq D$$

We also extend subtyping to method types:

$$\begin{aligned} \bar{\sigma} <: \bar{\tau} &\iff |\bar{\sigma}| = |\bar{\tau}| \wedge \forall i \in 1, \dots, |\bar{\sigma}|. \sigma_i <: \tau_i \\ (\bar{\sigma}, \tau) <: (\bar{\sigma}', \tau') &\iff \bar{\sigma}' <: \bar{\sigma} \wedge \tau <: \tau' \end{aligned}$$

3.2 Annotated class tables

We extend (plain) class tables \mathcal{C}_0 to *annotated class tables* $\mathcal{C} = (A^{get}, A^{set}, M)$.

- The *annotated field typings* $A^{get}, A^{set} : (\mathcal{C} \times \mathcal{R} \times \mathcal{F}) \rightarrow \mathcal{T}$ assign to each class C , region r and field $f \in \text{fields}(C)$ the refined type of the field for all objects of class C in region r . The field type is split into a covariant get-type A^{get} for the data read from the field, and a contravariant set-type A^{set} that is needed for data to be written to the field. This technique improves precision and is borrowed from Hofmann and Jost [16]. More formally, annotated field typings are *well-formed* if for all classes C , subclasses $D \preceq C$, regions r and fields $f \in \text{fields}(C)$,

$$\begin{array}{c}
\text{T-SUB} \frac{\Gamma \vdash e : \sigma \quad \sigma <: \tau}{\Gamma \vdash e : \tau} \quad \text{T-LET} \frac{\Gamma \vdash e_1 : \sigma \quad \Gamma, x : \sigma \vdash e_2 : \tau}{\Gamma \vdash \text{let } x = e_1 \text{ in } e_2 : \tau} \quad \text{T-VAR} \frac{}{\Gamma, x : \tau \vdash x : \tau} \\
\\
\text{T-NUL} \frac{}{\Gamma \vdash \text{null} : \tau} \quad \text{T-IFINST} \frac{x \in \text{dom}(\Gamma) \quad \Gamma \vdash e_1 : \tau \quad \Gamma \vdash e_2 : \tau}{\Gamma \vdash \text{if } x \text{ instanceof } E \text{ then } e_1 \text{ else } e_2 : \tau} \\
\\
\text{T-IFEQ} \frac{\Gamma, x : C_{R \cap S}, y : D_{R \cap S} \vdash e_1 : \tau \quad \Gamma, x : C_R, y : D_S \vdash e_2 : \tau}{\Gamma, x : C_R, y : D_S \vdash \text{if } x = y \text{ then } e_1 \text{ else } e_2 : \tau} \\
\\
\text{T-NEW} \frac{}{\Gamma \vdash \text{new } C : C_{\{r\}}} \quad \text{T-INVOKE} \frac{\forall r \in R. \exists (\bar{\sigma}', \tau') \in M(C, r, m). (\bar{\sigma}', \tau') <: (\bar{\sigma}, \tau)}{\Gamma, x : C_R, \bar{y} : \bar{\sigma} \vdash x.m(\bar{y}) : \tau} \\
\\
\text{T-GETF} \frac{\forall r \in R. A^{\text{get}}(C, r, f) <: \tau}{\Gamma, x : C_R \vdash x.f : \tau} \quad \text{T-SETF} \frac{\forall r \in R. \tau <: A^{\text{set}}(C, r, f)}{\Gamma, x : C_R, y : \tau \vdash x.f := y : \tau}
\end{array}$$

Figure 4: Region type system

- $A^{\text{set}}(C, r, f) <: A^{\text{get}}(C, r, f)$, and
- $A^{\text{get}}(D, r, f) <: A^{\text{get}}(C, r, f)$ and $A^{\text{set}}(C, r, f) <: A^{\text{set}}(D, r, f)$.

Also, the class component of $A^{\text{get}}(C, r, f)$ and $A^{\text{set}}(C, r, f)$ must be $A_0(C, f)$, i.e. invariant.

- The *annotated method typing* $M : (\mathcal{C} \times \mathcal{R} \times \mathcal{M}) \rightarrow \mathcal{P}(\bar{\mathcal{T}} \times \mathcal{T})$ assigns to each class C , region r and method $m \in \text{methods}(C)$ an unbounded number of refined method types for objects of class C in region r , enabling infinite polymorphic method types. This makes it possible to use a different type at different invocation sites (program points) of the same method. Even more importantly, the same invocation site can be checked in different type derivations with different method types. For every *well-formed* annotated method type, there must be an improved method type in each subclass: for all classes C , subclasses $D \preceq C$, regions r , and methods $m \in \text{methods}(C)$, we require
 - $\forall (\bar{\sigma}, \tau) \in M(C, r, m). \exists (\bar{\sigma}', \tau') \in M(D, r, m). (\bar{\sigma}', \tau') <: (\bar{\sigma}, \tau)$.

Again, the class components of the refined types $M(C, r, m)$ have to match the classes of the underlying unannotated method type $M_0(C, m)$.

In the following, we assume a fixed annotated class table \mathcal{C} with well-formed field and method typings.

3.3 Region type system

The type system (see figure 4) derives judgements $\Gamma \vdash e : \tau$, meaning FJEU expression e has type τ with respect to a *variable context* (store typing) $\Gamma : \mathcal{X} \rightarrow \mathcal{T}$ that maps variables to types.

The rule T-SUB is used to obtain weaker types for the expression. The rules T-LET, T-VAR, and T-IFINST are standard. The rule T-IFEQ exploits the fact that the two variables must point to the same object (or be *null*) in the then branch, therefore the intersection of the region sets can be assumed. In T-NUL, the *null* value may have any type (any class and any region set). In the rule T-NEW, we may choose any region r , which is an abstract location that includes (possibly among others) the concrete location of the object allocated by this expression.

When a field is read (T-GETF), we look up the type of the field in the A^{set} table. As the variable x may point to a number of regions, we need to ensure that τ is an upper bound of the get-types of f over all $r \in R$. In contrast, when a field is written (T-SETF), the written value must have a subtype of the types allowed for that field by the A^{set} table with respect to each possible region $r \in R$. Finally, the rule T-INVOKE requires that for all regions $r \in R$ where the receiver object x may reside, there must exist a method typing that is suitable for the argument and result types.

An FJEU program $P = (\prec, fields, methods, mtable)$ is *well-typed* if for all classes C , regions r , methods m and method types $(\bar{\sigma}, \tau)$ such that $(\bar{\sigma}, \tau) \in M(C, r, m)$, the following judgement is derivable:

$$[this \mapsto C] \cup [x_i^m \mapsto \sigma_i]_{i \in \{1, \dots, |x^m|\}} \vdash mtable(C, m) : \tau$$

The polymorphic method types make the region type system very expressive in terms of possible analyses of a given program. Each method may have many types, each corresponding to a derivation of the respective typing judgment. In the different derivations, different regions may be chosen for new objects, and different types may be chosen for called methods. This flexibility provides the basis for our embedding of external pointer analyses. Moreover, since there may be infinitely many method types for each method, this system is equivalent to one which allows infinite unfoldings of method calls.

3.4 Interpretation

We now give a formal interpretation of the typing judgement in form of a soundness theorem. Afterwards, we continue by restricting the expressivity of the system and reformulating the rules in order to move towards an actual type checking algorithm. The underlying idea is that we only have to prove soundness once for the general system; for later systems, it suffices to show that they are special cases of the general system, such that the soundness theorem applies to these systems as well.

A *heap typing* $\Sigma, \Pi : \mathcal{L} \rightarrow (C \times \mathcal{R})$ assigns to heap locations a static class (an upper bound of the actual class found at that location) and a region. Heap typings, a standard practice in type systems for languages with dynamic memory allocations [24, chapter 13.4], separate the well-typedness definitions of locations in stores and objects from the actual heap, thereby avoiding the need for a co-inductive definition for well-typed heaps in the presence of cyclic structures. Heap typings map locations to very specific types, namely those where the region set is a singleton. A heap typing thus partitions a heap into (disjoint) regions.

We define a typing judgment for values $\Sigma \vdash v : \tau$, which means that according to heap typing Σ , the value v may be typed with τ . In particular, the information in $\Sigma(l)$ specifies the type of l . Also, the typing judgment of locations is lifted to stores and variable contexts.

$$\frac{}{\Sigma \vdash \text{null} : \tau} \quad \frac{\Sigma(l) = (C, r)}{\Sigma \vdash l : C_{\{r\}}} \quad \frac{\Sigma \vdash v : \sigma \quad \sigma <: \tau}{\Sigma \vdash v : \tau}$$

$$\Sigma \vdash s : \Gamma \iff \forall x \in \text{dom}(\Gamma). \Sigma \vdash s(x) : \Gamma(x)$$

A heap h is *well-typed* with respect to a heap typing Σ and implicitly a field typing A^{set} , written $h \models \Sigma$, if the type for all locations given by Σ are actually “valid” with respect to the classes of the objects, and if the field values are well-typed with respect to A^{set} and Σ :

$$h \models \Sigma \iff \forall l \in \text{dom}(\Sigma). l \in \text{dom}(h) \wedge \Sigma \models h(l) : \Sigma(l)$$

where

$$\Sigma \models (C, F) : (D, r) \iff C \preceq D \wedge \text{dom}(F) = \text{fields}(C) \wedge \forall f \in \text{fields}(C). \Sigma \vdash F(f) : A^{set}(C, r, f)$$

As the memory locations are determined at runtime, the heap typings cannot be derived statically. Instead, our interpretation of the typing judgement $\Gamma \vdash e : \tau$ states that whenever a well-typed program is executed on a heap that is well-typed with respect to some typing Σ , then the final heap after the execution is well-typed with respect to some possibly larger heap typing Π . The typing Π may be larger to account for new objects that may have been allocated during execution, but the type of locations that already existed in Σ may not change. More formally, a heap typing Π *extends* a heap typing Σ , written $\Pi \sqsupseteq \Sigma$, if $\text{dom}(\Sigma) \subseteq \text{dom}(\Pi)$ and $\forall l \in \text{dom}(\Sigma). \Sigma(l) = \Pi(l)$.

Theorem 1 (Soundness Theorem). *Fix a well-typed program P . For all $\Sigma, \Gamma, \tau, s, h, e, v, k$ with*

$$\Gamma \vdash e : \tau \quad \text{and} \quad \Sigma \vdash s : \Gamma \quad \text{and} \quad (s, h) \vdash e \Downarrow v, k \quad \text{and} \quad h \models \Sigma$$

there exists some $\Pi \sqsupseteq \Sigma$ such that

$$\Pi \vdash v : \tau \quad \text{and} \quad k \models \Pi.$$

Proof. By induction over the derivations of the operational semantics and the typing judgement. The full proof can be found in appendix A. We have also developed a formalisation in Isabelle/HOL [5]. \square

4 Parametrized Region Type System

Our next goal is to use the type system for the automatic algorithmic verification of results of external analyses. In this section, we first show to *interpret* given analysis results in the type system, and then present an algorithmic type checking algorithm for the automatic verification.

The interpretation of given results requires the reformulation of the above type system to explicitly model abstraction principles that are fundamental for a number of different pointer analysis techniques [25]. We focus on two general classes of abstractions: the abstraction of the call graph using contexts, and the abstraction of objects on the heap. The region type system is equipped with parameters that can be instantiated to specific abstraction principles. We show that the parametrized version of the type system arises as a specialisation of the general region type system.

In the following, the notion of program points is made explicit by annotating expressions with *expression labels* $i \in \mathcal{I}$: we write $[e]^i$ for FJEU expressions, where e is defined as before. An FJEU program is well-formed if each expression label i appears at most once in it. In the following, we only consider well-formed programs, and simply write e instead $[e]^i$ if the expression label i is not important.

4.1 Abstraction principles

Context-insensitive pointer analyses examine each method exactly once. Their result provides for each method a single pointer structure specification which is used for every call of that method. *Context-sensitive* algorithms improve precision: each method may be analyzed multiple times under different *contexts*, so that different specifications can be used for different calls to the same method. A context-sensitive algorithm settles on a specific (finite) set of contexts, and produces for each method one pointer structure specification per context. Pointer structure specifications correspond to our method types. We therefore model the concept of contexts by introducing a finite abstract set of contexts \mathcal{Z} , and by parametrizing the method typing function M to associate one method type per context in \mathcal{Z} .

The choice of contexts and specifications for each method call depends on the analysis in question. For example, call-site sensitive algorithms [26] process each method once for each program point where the method is called. Receiver-object sensitive analyses [22] differentiate pointer structure specifications of a method according to the abstraction of the invoking object. More powerful analyses use *call stacks* as contexts to differentiate method calls. For example, a method m may be analyzed for each possible

call-site stack that may occur at the time when m is called. Similarly, receiver-object stacks can be used. In other words, one considers the *call graph* of the program. A method m is then represented by a node in this graph, and a context corresponds to a possible path that leads to the node. As there may be infinitely many paths in recursive programs, the number of paths needs to be restricted by some mechanism. A common way is to consider only the last k entries on the stack (k -CFA [27]), or to collapse each strongly connected component into a node, thereby eliminating recursive cycles from the set of possible paths [13, 31]. Following this observation, we employ a general *context transfer function* ϕ which represents the edges in the abstract call graph. The function selects a context for the callee based on the caller's context, the class of the receiver object, its region, the method name, and the call site.

Another abstraction principle is the object abstraction, i.e. the abstract location assigned to allocated objects. This corresponds to our concepts of regions. As pointer analysis algorithms differentiate only finitely many abstract objects, we can restrict the set of regions \mathcal{R} to a finite size.

A common abstraction is to distinguish objects according to their allocation site and/or their class. More precise analyses also take into account the context under which allocation takes place. For example, in object-sensitive analysis by Milanova et al. [22], objects are distinguished by their own allocation site and the allocation site of the `this` object. Objects may also be distinguished by the call site of the invoked method, a technique called *heap specialization* [23]. We model these concepts by an *object abstraction function* ψ that assigns the region for the new object, given the allocation site and the current method context. Our system is by design class-sensitive, as the class information is part of the type.

Figure 5 summarizes the four parameters of our system: contexts, context transfer function, regions, and object abstraction function. Also, it shows how to instantiate the parameters to obtain various standard abstraction principles.

The parametrized method typing $\hat{M} : (\mathcal{C} \times \mathcal{R} \times \mathcal{Z} \times \mathcal{M}) \rightarrow \overline{\mathcal{T}} \times \mathcal{T}$ replaces the annotated polymorphic method typing M in the annotated class table. It is well-formed if for all classes C and subclasses $D \preceq C$, regions $r \in \mathcal{R}$, methods $m \in \text{methods}(C)$, and contexts $z \in \mathcal{Z}$, it holds $\hat{M}(D, r, z, m) <: \hat{M}(C, r, z, m)$.

4.2 The parametrized type system

The parametrized type system extends the typing judgment of the general region type system by a context component z . With the exception of the following two rules, all rules remain as presented in section 3 (with the addition of the context z in each judgement):

$$\text{TP-NEW} \frac{r = \psi(z, i)}{\Gamma; z \vdash [\text{new } C]^i : C_{\{r\}}} \quad \text{TP-INVOKE} \frac{\forall r \in \mathcal{R}. \hat{M}(C, r, \phi(z, C, r, m, i), m) <: (\overline{\sigma}, \tau)}{\Gamma, x : C_R, \bar{y} : \overline{\sigma}; z \vdash [x.m(\bar{y})]^i : \tau}$$

While in the previous system any region r could be chosen for new objects, we have restricted this flexibility in the TP-NEW rule to the region specified by ψ . Moreover, instead of allowing arbitrarily many types per method, from which any type could be selected for invocations, we now have one type determined by the context z that is selected by the ϕ function in the TP-INVOKE rule.

For a given parametrized method typing \hat{M} , we define a the corresponding polymorphic method typing $M(C, r, m) := \bigcup_{z \in \mathcal{Z}} \{\hat{M}(C, r, z, m)\}$. The rules of the parametrized system are derivable in the previous system: The only changed rules TP-NEW and TP-INVOKE have more restrictive premises than their counterparts T-NEW and T-INVOKE. Hence if $\Gamma; z \vdash e : \tau$ can be derived from some \hat{M} in the parametrized system, then $\Gamma \vdash e : \tau$ can be derived in the previous system with respect to the corresponding method typing M .

A method table is *well-typed* for \hat{M} if for all classes C , contexts z , regions r , and methods m such that $\hat{M}(C, r, z, m) = (\overline{\sigma}, \tau)$, the judgement $\Gamma; z \vdash \text{mtable}(C, m) : \tau$ can be derived with $\Gamma = [\text{this} \mapsto C] \cup [x_i^m \mapsto \sigma_i]_{i \in \{1, \dots, |\overline{x}^m\}}$. It is easy to see that if a method table is well-typed with respect to \hat{M} in the

Regions (finite):	$r, s, t \in \mathcal{R}$
Contexts (finite):	$z \in \mathcal{Z}$
Context transfer function:	$\phi \in \mathcal{Z} \times \mathcal{C} \times \mathcal{R} \times \mathcal{M} \times \mathcal{I} \rightarrow \mathcal{Z}$
Object abstraction function:	$\psi \in \mathcal{Z} \times \mathcal{I} \rightarrow \mathcal{R}$

<i>set of regions</i>	<i>object abstraction function</i>	<i>principle</i>
$\mathcal{R} = \mathcal{I}$	$\psi(z, i) = i$	allocation site abstraction
$\mathcal{R} = \mathcal{Z} = \mathcal{I} \times \mathcal{I}$	$\psi((i_1, i_2), i_0) = (i_0, i_1)$	object-sensitive allocation site abstraction
$\mathcal{R} = \mathcal{Z} = \mathcal{I}$	$\psi(i_c, i) = i_c$	heap specialization

<i>set of contexts</i>	<i>context transfer function</i>	<i>principle</i>
$\mathcal{Z} = \{z_0\}$	$\phi(z, C, r, m, i) = z_0$	context-insensitivity
$\mathcal{Z} = \mathcal{R}$	$\phi(z, C, r, m, i) = r$	object-sensitive 1-CFA
$\mathcal{Z} = \bigcup_{n \in \{1, \dots, k\}} \mathcal{M}^n$	$\phi(z, C, r, m, i) = (m :: z) _k$	method identifier k -CFA
$\mathcal{Z} = \{\bar{i} \in \sigma(\mathcal{I}') \mid \mathcal{I}' \subseteq \mathcal{I}\}$	$\phi(z_1, C, r, m, i) = z_2$ s.th. $IE_c(z_1, i, z_2, m)$	CFA with eliminated recursive cycles

where

- $\sigma(X)$ is the set of all permutations of X
- $L|_k$ is the truncation of list L after the first k elements
- IE_c is the call graph relation of [31] where recursive cycles have been replaced by single nodes

Figure 5: The four type system parameters, and possible instantiations to common abstraction principles

parametrized system, then it is also well-typed with respect to the corresponding method typing M in the general system. Therefore, the soundness theorem is applicable to the parametrized region type system.

4.3 Algorithmic type checking

We now present a syntax-directed form of the previous type system, from which one can directly read off an algorithm $A(\Gamma, e) = \tau$ that computes “from left to right” the type τ of an expression e based on a store typing Γ , similar to the approach taken by Pierce [24]. For this, we eliminate the subtyping rule, and instead specify the most precise resulting type τ for each expression. The full algorithmic type system is shown in figure 6.

For the TA-NULL rule, the type system needs some guidance for the choice of the type. We solve this problem by assuming that the syntax is annotated with a class: the expression null_C means that the type system shall use the class C for that particular *null* value. The reason why we have not introduced this annotated syntax from the beginning is that the inference of class information is orthogonal to pointer analysis and the verification of region sets.

For the rules TA-IF, TA-GETF, and TA-INVOKE, we compute the least upper bound of two or more types. We define the join $C_R \vee D_S = E_{R \cup S}$ where E is the most specific common superclass of C and D . The notation $\bigvee X$ denotes join of a set of types X . The functions *argtypes* and *restype* used in TA-INVOKE return the argument and result types of a method type, respectively.

The soundness proof shows that the internalisation of the subtyping rule is correct, i.e. that the judgements derived with the algorithmic type system can also be derived with the parametrized type system. We only outline the proof: The rule TA-IFEQ is derivable by a combination of the rules TP-IFEQ and

$$\begin{array}{c}
\text{TA-VAR} \frac{}{\Gamma, x : \tau; z \vdash x : \tau} \qquad \text{TA-LET} \frac{\Gamma; z \vdash e_1 : \sigma \quad \Gamma, x : \sigma; z \vdash e_2 : \tau}{\Gamma; z \vdash \text{let } x = e_1 \text{ in } e_2 : \tau} \\
\\
\text{TA-IFEQ} \frac{\Gamma, x : C_{R \cap S}, y : D_{R \cap S}; z \vdash e_1 : \tau_1 \quad \Gamma, x : C_R, y : D_S; z \vdash e_2 : \tau_2}{\Gamma, x : C_R, y : D_S; z \vdash \text{if } x \text{ instanceof } E \text{ then } e_1 \text{ else } e_2 : \tau_1 \vee \tau_2} \\
\\
\text{TA-IFINST} \frac{\Gamma; z \vdash e_1 : \tau_1 \quad \Gamma; z \vdash e_2 : \tau_2}{\Gamma; z \vdash \text{if } x \text{ instanceof } E \text{ then } e_1 \text{ else } e_2 : \tau_1 \vee \tau_2} \\
\\
\text{TA-NULL} \frac{}{\Gamma; z \vdash \text{null}_C : C_\emptyset} \qquad \text{TA-NEW} \frac{r = \psi(z, i)}{\Gamma; z \vdash \text{new } C : C_{\{r\}}} \\
\\
\text{TA-GETF} \frac{\bigvee_{r \in R} A^{\text{get}}(C, r, f) = \tau}{\Gamma, x : C_R; z \vdash x.f : \tau} \qquad \text{TA-PUTF} \frac{\forall r \in R. \tau <: A^{\text{set}}(C, r, f)}{\Gamma, x : C_R, y : \tau; z \vdash x.f := y : \tau} \\
\\
\text{TA-INVOKE} \frac{\forall r \in R. \text{argtypes}(\hat{M}(C, r, \phi(z, C, r, m, i), m)) <: \bar{\sigma} \quad \bigvee_{r \in R} \text{restype}(\hat{M}(C, r, \phi(z, C, r, m, i), m)) = \tau}{\Gamma, x : C_R, \bar{y} : \bar{\sigma}; z \vdash x.m(\bar{y}) : \tau}
\end{array}$$

Figure 6: Algorithmic type system

TP-SUB; similarly for TA-IFINST. The rule TA-NULL is a specialisation of TP-NULL. In TA-GETF, the computed type τ satisfies the premise of the TP-GETF. The same holds for the type τ in the TA-INVOKE rule, which is a candidate for τ in TP-INVOKE. All other rules as well as the definition of annotated class tables remain unchanged. We can therefore apply the soundness theorem to the algorithmic type system.

We have also developed a context-sensitive pointer analysis algorithm for FJEU to demonstrate the verification capabilities of the algorithmic type system. Building on the work by Whaley and Lam [31], we specified the algorithm declaratively as a set of recursive Datalog rules, which can be translated into an efficient BDD-based algorithm. As presented in their work, we chose the IE_c call graph abstraction relation for the contexts, and the program points (expression labels) as the object abstraction. The algorithm computes several relations that express which variables, fields, and method arguments may point to which regions, depending on the context. Finally, we have interpreted the resulting relations for the region type system as annotated class tables, and modelled the abstractions using appropriate functions ψ and ϕ . Details can be found in appendix B.

5 String Analysis

String analysis is a dataflow analysis technique to determine possible string values (character sequences) that may occur during the execution of a program. Since strings appear as objects in Java, it is natural to implement a string analysis by building on pointer analysis: string objects are identified and tracked by the pointer analysis, while their possible values are determined by the string analysis.

We now equip the FJEU language with special string objects and operations to give a simplified formalisation of Java's `String` class, and extend the region type system to enable the verification of

pointer analyses of string objects. Afterwards, we show how to use the region information to interpret the results of a specific string analysis.

5.1 FJEU with strings

The language FJEUS is an extension of FJEU with operations to create and concatenate strings. While the `String` class in Java is just another class in the Java class hierarchy, it is regarded as a separate type in FJEUS. This allows us to treat string objects differently: An object of class `String` in FJEUS is simply a string value (character sequence) on the heap. The meta-variable w ranges over character sequences \mathcal{W} , and W ranges over sets of string values. We rely on a given sequence concatenation function $+$.

$$\begin{array}{ll}
 \text{character sequences:} & w \in \mathcal{W} \\
 \text{sets of character sequences:} & W \in \mathcal{P}(\mathcal{W}) \\
 \text{extended expressions:} & \mathcal{E} \ni e ::= \dots \mid \text{new String}(w) \mid x.\text{concat}(y) \\
 \\
 \text{character sequence concatenation:} & + \in \mathcal{W} \times \mathcal{W} \rightarrow \mathcal{W} \\
 \text{heaps:} & h, k \in \mathcal{L} \rightarrow \mathcal{O} \cup \mathcal{W}
 \end{array}$$

The expression `new String(w)` allocates a new string object with the character sequence w on the heap. The string operation `$x.\text{concat}(y)$` has its own special semantics and is implemented with the $+$ operator. As we only model strings with non-mutable values in the language, a string concatenation always creates a new string object on the heap. The operational semantics is extended as follows:

$$\frac{l \notin \text{dom}(h)}{(s, h) \vdash \text{new String}(w) \Downarrow l, h[l \mapsto w]} \qquad \frac{\begin{array}{ll} s(x) = l_1 & s(y) = l_2 & h(l_1) = w_1 & h(l_2) = w_2 \\ & l \notin \text{dom}(h) & w = w_1 + w_2 \end{array}}{(s, h) \vdash x.\text{concat}(y) \Downarrow l, h[l \mapsto w]}$$

Note that this rather modest extension is intended to keep the formalisation simple. Other extensions could include more string operations, or mutable string objects that model Java’s `StringBuffer` class.

5.2 Pointer analysis for string objects

We extend the region type system from section 3 to accommodate the new string objects. We distinguish references to “proper” objects and to string objects: a type is either a class with a region set (C_R), or the special `String` class with a region set (String_R). The `String` class is independent from other classes in the class hierarchy.

$$\begin{array}{l}
 \text{types: } \sigma, \tau \in \mathcal{T} = (\mathcal{C} \cup \{\text{String}\}) \times \mathcal{P}(\mathcal{R}) \\
 \\
 \frac{C \preceq D \quad R \subseteq S}{C_R <: D_S} \qquad \frac{R \subseteq S}{\text{String}_R <: \text{String}_S}
 \end{array}$$

Fields and method typings may now include the String_R type. Also, all existing typing rules from section 3 remain unchanged. In particular, the `null` value may be assigned a `String` type. A field `$x.f$` may only be accessed if x is a (non-`String`) class C ; similarly for method calls `$x.m$` . These are the two additional typing rules for `new String(str)` and `$x.\text{concat}(y)$` :

$$\frac{}{\Gamma \vdash \text{new String}(w) : \text{String}_{\{r\}}} \qquad \frac{}{\Gamma, x : \text{String}_R, y : \text{String}_S \vdash x.\text{concat}(y) : \text{String}_{\{t\}}}$$

The heap typings $\Sigma, \Pi : \mathcal{L} \rightarrow (\mathcal{C} \cup \{\text{String}\}) \times \mathcal{R}$ may now also map locations to the `String` class. We extend the well-typed value relation $\Sigma \vdash v : \tau$ accordingly:

$$\frac{}{\Sigma \vdash \text{null} : \tau} \quad \frac{\Sigma(l) = (C, r)}{\Sigma \vdash l : C_{\{r\}}} \quad \frac{\Sigma(l) = (\text{String}, r)}{\Sigma \vdash l : \text{String}_{\{r\}}} \quad \frac{\Sigma \vdash v : \sigma \quad \sigma <: \tau}{\Sigma \vdash v : \tau}$$

The definition of well-typed heaps additionally requires well-typed character sequences (last line):

$$\begin{aligned} h \models \Sigma &\iff \forall l \in \text{dom}(\Sigma). l \in \text{dom}(h) \wedge \Sigma \models h(l) : \Sigma(l) \\ \Sigma \models (C, F) : (D, r) &\iff C \preceq D \wedge \dots \text{ (as before)} \\ \Sigma \models w : (\text{String}, r) &\iff \text{PROP}(w, r) \end{aligned}$$

In other words, the property that a heap h is well-typed with respect to Σ now includes the condition that for all locations l such that $\Sigma(l) = (\text{String}, r)$, $h(l)$ contains a string value w that satisfies a certain property `PROP` with respect to r . For the moment, assume `PROP` is simply `True`. The proof of the soundness theorem is extended in a straight-forward way for the extensions described above. Moreover, the type system can be parametrized in the same fashion as described in section 4: as both string operations create new objects, we use the ψ function to determine the region of these objects.

In the following subsection, we present an analysis that can help to prevent cross-site scripting attacks, and give a semantic formalization by instantiating the string property `PROP`.

5.3 String analysis with operation contexts

In a typical cross-site scripting scenario, a user input is embedded into a string that is executed or interpreted. To prevent the injection of malicious code, one wants to track how a string is constructed, and ensure that its executable parts originate from trusted sources, like string literals in the program code.

We therefore add a string operation context Ω , which expresses the possible string operations that objects in specific regions may be the result of. The typing rules are extended with constraints on Ω .

$$\begin{aligned} \text{string operations: } \mathcal{O} \ni \omega &::= \text{newstring } w r \mid \text{concat } r s t \mid \text{unknown } W r \\ \text{string operation context: } \Omega &\in \mathcal{P}(\mathcal{O}) \end{aligned}$$

$$\frac{\text{newstring } w r \in \Omega}{\Gamma \vdash \text{new String}(w) : \text{String}_{\{r\}}} \quad \frac{\forall r \in R, s \in S. \text{concat } r s t \in \Omega}{\Gamma, x : \text{String}_R, y : \text{String}_S \vdash x.\text{concat}(y) : \text{String}_{\{t\}}}$$

Informally, an operation `newstring` $w r \in \Omega$ means that region r may include a string w . `concat` $r s t$ means that region t may include a string object that is obtained by concatenating some strings from regions r and s . The operation `unknown` $W r$ means that strings in region r may be from the set W . This is useful for external methods whose types are given to but not verified by the type system. Whenever more primitive string operations are added to the language, the set \mathcal{O} may be extended accordingly.

Formally, we define a semantic interpretation $\Omega[[r]]$ that gives the possible values for string objects in region r . It is defined as the smallest set satisfying the following conditions:

$$\begin{aligned} \text{newstring } w r \in \Omega &\Rightarrow w \in \Omega[[r]] \\ \text{concat } r s t \in \Omega &\Rightarrow \forall w_1 \in \Omega[[r], w_2 \in \Omega[[s]]. w_1 + w_2 \in \Omega[[t]] \\ \text{unknown } W r \in \Omega &\Rightarrow W \subseteq \Omega[[r]] \end{aligned}$$

After instantiating the string property as $\text{PROP}(w, r) \equiv w \in \Omega[[r]]$, the relation $h \models \Sigma$ ensures that string values on the heap are indeed in the interpretation of the string operation context.

Apart from this extensional interpretation, the string operation context and the typing of external methods also contain intensional information about the origin and the possible constructions of strings in a specific region, which enables the verification of more complex string policies.

For example, consider the following string-manipulating program that relies on external functions *getUserInput()* to retrieve data from the user, *escapeHTML(s)* that quotes all HTML tags in string *s* and returns the result as a new string, and *output(s)* that outputs the string *s*.

```
let firstPart = new String("<sometag>")
  in let contents = getUserInput()
    in let escContents = escapeHTML(contents)
      in output(firstPart.concat(sanContents))
```

The security policy is that the output may only be the result of a concatenation of a string literal with a string that does not contain HTML tags. The policy is expressed using the following types for external methods and the string operation context:

$$\begin{array}{ll}
 \textit{getUserInput} & : \text{unit} \longrightarrow \text{String}_{\{q\}} & \Omega = \{ \text{newstring } "<\text{sometag}>" \ l, \\
 \textit{escapeHTML} & : \text{String}_{\{q\}} \longrightarrow \text{String}_{\{s\}} & \text{unknown } \mathcal{W} \ q, \text{ unknown } \hat{\mathcal{W}} \ s, \\
 \textit{output} & : \text{String}_{\{t\}} \longrightarrow \text{unit} & \text{concat } s \ l \ t, \text{ concat } l \ s \ t \ }
 \end{array}$$

($\hat{\mathcal{W}}$ is the set of all strings that do not contain HTML tags, and *unit* is a unit type, which could be modelled in FJEUS as a null type like $\text{String}_{\emptyset}$.) The external function *getUserInput* returns strings with arbitrary values of the set \mathcal{W} in region *q* (“questionable”), which are converted by *escapeHTML* into strings of region *s* (“sanitized”), which are assumed to not contain any HTML tags (set $\hat{\mathcal{W}}$). For the literal *firstPart*, the type checker can assign the region *l* (“literals”), as the literal value in Ω matches. The *output* function only accepts strings from region *t* (“trusted”), which must be, according to Ω , a concatenation of strings from region *l* and (HTML tag-free) region *s*. Therefore, the typability of the program proves that the security policy is indeed fulfilled. The example demonstrates that handling trusted sanitizing functions is actually a strength of type-based presentations: simply assign an appropriate type to a function if you believe the associated semantic property of the function.

The approach is related to the work by Christensen et al. [9] on the analysis of string values using context-free grammars with operation productions. The symbolic string operations in the context correspond to nodes in their annotated flow graph, and their semantics of the flow graph resembles our interpretation of $\Omega[[r]]$. Similarly, Crégut and Alvarado [11] have presented an algorithm that tracks string objects with pointer analysis, and collects intensional information about the string operations applied to them. We thus expect that aspects of the results of these algorithms are verifiable in our system. Our approach is also related to taint analysis [30], as the region identifiers can convey information about the trustworthiness of strings, which is preserved throughout assignments and method invocations.

6 Discussion

We presented a framework for classifying alias analyses for Java-like languages, given by a hierarchy of region-based type systems. We demonstrated how existing disciplines arise as instantiations of our framework and may be given a uniform interpretation by embedding their results in a single type system. We also gave an algorithmic variant of the type system, thus enabling syntax-directed type-checking and hence validation of analyses results. Finally, we showed how our framework may be extended to string analyses. In the following, we briefly discuss specific design decisions, and outline future work.

To our knowledge, most existing pointer analyses express their results at a coarse-grain level of syntactic structure such as methods. In accordance with this, we employed a phrase-based formulation of type systems and interpreted judgements with respect to a big-step evaluation semantics. An extension of the interpretation to include non-terminating executions appears possible, using techniques as in [6].

Our framework does not aim to be flow- or path-sensitive. We see these concepts as orthogonal to the central idea of our paper, namely interpreting context-sensitivity using polyvariant types. Nevertheless, we acknowledge the increasing relevance of flow and path sensitivity in recent work on pointer analysis. The T-IFEQ rule illustrates a possible extension of the type system with path-sensitive capabilities: the information from the branching expression is used to refine the analysis for the “then” branch of the conditional. Moreover, subderivations of a judgement contain implicitly more fine-grained (non-)alias relationships applicable at intermediate program points, and include aspects of flow-sensitivity as any variable may be associated with different types in a derivation. Arguably, local alias assertions could be made more explicit by moving to a small-step operational regime and/or formulations of the type systems that are inspired by abstract interpretation and yield a global specification table with entries for all program points [3]. However, the use of evaluation-style judgements greatly simplifies soundness proofs at least at the level of methods, as recursive calls follow the type derivation structure.

The Doop framework [8] enables the definition of highly precise and efficient pointer analyses declaratively using Datalog rules. While the authors do not seem to aim at a fully formal correctness proof that interprets the Datalog rules and relations with respect to the semantics of the Java language, they take great care to separate the essential aspects of analysis techniques from implementation details of the algorithm. We intend to look at the ideas in their work in order to find ways to adapt our type system to more language features and pointer analyses.

In addition to type systems for pointer alias analysis, the literature contains the terminology “type-based alias analysis” [12]. The latter term appears to mean that datatype or static class information is used to *improve the precision* of an alias analysis, while region type systems directly *include* the points-to relations in the types. However, as our system extends the ordinary type system of Java, it arguably also encompasses type-based alias analyses.

Having been developed in order to embed static analysis results, it is not surprising that the type systems over-approximate their semantic guarantee. Thus, failure to validate the *legitimate* result of a specific analysis may be rooted in either an incorrect interpretation of the analysis into our framework or the fact that the analysis is more precise than the host type system. A particular direction in which our type system (and some analyses) might be generalized is shape analysis [20]. Another interesting recent development that our work does not accommodate is the analysis of programs that use reflection [30]; this would require a fundamental understanding of the semantic analysis of reflection.

Although we have only examined the *results* of pointer analysis algorithms, the algorithms can be seen as an external means of type inference. It seems promising to further investigate the *implementations* of these algorithms, and to recreate their logic in the type system in order to obtain a parametric type system with a fully automatic (internal) type inference. Alternatively, the identification of abstraction principles could also propose a way to parametrize existing pointer analysis implementations.

Regarding the string analysis, we have concentrated on the previously-noted observation that the precision of the analysis benefits from the availability of (non-)aliasing information [11]. In principle, the benefits may be mutual. For example, the method call $x.concat(y)$ on a `String` in Java actually returns the reference x if y has a length of zero. If the length of y can be obtained from a string analysis, this information helps to improve the region set for the result type in the rule for concatenation. Mutual dependencies between aliasing and string analyses may thus be an interesting topic for future work.

Also, we have only outlined the use of a verified string analysis for security policies. In collaboration with SAP’s Sophia-Antipolis-based research lab on security we plan to extend the type system with string operation contexts to verify the absence of cross-site scripting attacks in concrete scenarios.

Acknowledgements We would like to thank Pierre Crégut for making us aware of the interplay between pointer and string analyses and providing us with several pointers to the literature, as well as Keqin Li (SAP research) for giving us insights to concrete cross-site scripting scenarios and the corresponding string policies. This work was supported by the Information Society Technologies programme of the European Commission, Future and Emerging Technologies under the IST-2005-015905 MOBIUS project.

References

- [1] L. O. Andersen. *Program Analysis and Specialization for the C Programming Language*. PhD thesis, DIKU, University of Copenhagen, 1994.
- [2] A. Banerjee and T. Jensen. Modular control-flow analysis with rank 2 intersection types. *Mathematical Structures in Computer Science*, 13(1):87–124, 2003.
- [3] G. Barthe, D. Pichardie, and T. Rezk. A certified lightweight non-interference Java bytecode verifier. In *European Symposium on Programming (ESOP 2007)*, volume 4421 of *Lecture Notes in Computer Science*, pages 125 – 140. Springer-Verlag, 2007.
- [4] N. Benton, A. Kennedy, L. Beringer, and M. Hofmann. Relational semantics for effect-based program transformations with dynamic allocation. In *Principles and Practice of Decl. Prog. (PPDP '07)*. ACM, 2007.
- [5] L. Beringer, R. Grabowski, and M. Hofmann. Verifying Pointer and String Analyses with Region Type Systems: Soundness proofs and other material. <http://www.tcs.uni-lmu.de/~grabow/regions>, 2010.
- [6] L. Beringer, M. Hofmann, and M. Pavlova. Certification Using the Mobius Base Logic. In *Formal Methods for Components and Objects: 6th International Symposium (FMCO 2007)*, pages 25–51, Berlin, Heidelberg, 2008. Springer-Verlag.
- [7] M. Berndt, O. Lhoták, F. Qian, L. Hendren, and N. Umanee. Points-to analysis using BDDs. In *2003 Conference on Programming Language Design and Implementation (PLDI '03)*, New York, NY, USA, 2003. ACM.
- [8] M. Bravenboer and Y. Smaragdakis. Strictly declarative specification of sophisticated points-to analyses. In *24th Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA 2009)*, pages 243–262. ACM, 2009.
- [9] A. S. Christensen, A. Møller, and M. I. Schwartzbach. Precise analysis of string expressions. In *Static Analysis: 10th International Symposium (SAS 2003)*, volume 2694 of *Lecture Notes in Computer Science*. Springer-Verlag, 2003.
- [10] K. Crary, D. Walker, and G. Morrisett. Typed memory management in a calculus of capabilities. In *26th Symposium on Principles of programming languages (POPL '99)*, pages 262–275. ACM, 1999.
- [11] P. Crégut and C. Alvarado. Improving the Security of Downloadable Java Applications With Static Analysis. *Electronic Notes in Theoretical Computer Science*, 141(1):129–144, 2005.
- [12] A. Diwan, K. S. McKinley, and J. E. B. Moss. Type-based alias analysis. In *1998 Conference on Programming Language Design and Implementation (PLDI '98)*, pages 106–117, New York, NY, USA, 1998. ACM.
- [13] M. Emami, R. Ghiya, and L. J. Hendren. Context-sensitive interprocedural points-to analysis in the presence of function pointers. In *1994 Conference on Programming language design and implementation (PLDI'94)*, pages 242–256, New York, NY, USA, 1994. ACM.
- [14] J. S. Foster, T. Terauchi, and A. Aiken. Flow-sensitive type qualifiers. In *2002 Conference on Programming language design and implementation (PLDI '02)*, pages 1–12, New York, NY, USA, 2002. ACM.
- [15] B. Hardekopf and C. Lin. Semi-sparse flow-sensitive pointer analysis. *SIGPLAN Not.*, 44(1):226–238, 2009.
- [16] M. Hofmann and S. Jost. Type-based amortised heap-space analysis. In *15th European Symposium on Programming (ESOP 2006)*, volume 3924 of *Lecture Notes in Computer Science*, pages 22–37, Vienna, Austria, 2006. Springer.
- [17] A. Igarashi, B. Pierce, and P. Wadler. Featherweight Java: A minimal core calculus for Java and GJ. In *1999 Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA 1999)*. ACM, 1999.

- [18] T. Lenherr. Taxonomy and Applications of Alias Analysis. Master's thesis, ETH Zürich, 2008.
- [19] O. Lhoták. *Program Analysis using Binary Decision Diagrams*. PhD thesis, McGill University, Jan. 2006.
- [20] S. Loncaric. A survey of shape analysis techniques. *Pattern Recognition*, 31:983–1001, 1998.
- [21] J. M. Lucassen and D. K. Gifford. Polymorphic effect systems. In *15th Symposium on Principles of programming languages (POPL '88)*, pages 47–57, New York, NY, USA, 1988. ACM.
- [22] A. Milanova, A. Rountev, and B. G. Ryder. Parameterized object sensitivity for points-to analysis for Java. *ACM Trans. Softw. Eng. Methodol.*, 14(1):1–41, 2005.
- [23] E. M. Nystrom, H.-S. Kim, and W. W. Hwu. Importance of heap specialization in pointer analysis. In *5th Workshop on Program Analysis for Software Tools and Engineering (PASTE '04)*, New York, USA, 2004. ACM.
- [24] B. C. Pierce. *Types and Programming Languages*. MIT Press, 2002.
- [25] B. G. Ryder. Dimensions of precision in reference analysis of object-oriented programming languages. In *12th International Conference on Compiler Construction*, volume 2262, pages 126–137. Springer-Verlag, 2003.
- [26] M. Sharir and A. Pnueli. Two approaches to interprocedural data flow analysis. In Muchnick and Jones, editors, *Program Flow Analysis: Theory and Applications*. Prentice Hall International, 1981.
- [27] O. Shivers. *Control-Flow Analysis of Higher-Order Languages, or Taming Lambda*. PhD thesis, Carnegie Mellon University, Pittsburgh, PA, USA, 1991. Technical Report CMU-CS-91-145.
- [28] B. Steensgaard. Points-to analysis in almost linear time. In *23rd Symposium on Principles of programming languages (POPL '96)*, pages 32–41, New York, NY, USA, 1996. ACM.
- [29] M. Tofte and J.-P. Talpin. Region-based memory management. *Information and Computation*, 132(2):109–176, 1997.
- [30] O. Tripp, M. Pistoia, S. J. Fink, M. Sridharan, and O. Weisman. Taj: effective taint analysis of web applications. In *2009 Conference on Programming Language Design and Implementation (PLDI '09)*, pages 87–97, New York, NY, USA, 2009. ACM.
- [31] J. Whaley and M. S. Lam. Cloning-based context-sensitive pointer alias analysis using binary decision diagrams. *SIGPLAN Not.*, 39(6):131–144, 2004.

A Soundness Proof

Theorem 2. Fix a well-typed FJEU program P . If for all $\Sigma, \Gamma, \tau, s, h, e, v, k$

$$\begin{aligned} \Gamma &\vdash e : \tau \\ \Sigma &\vdash s : \Gamma \\ (s, h) &\vdash e \Downarrow v, k \\ h &\models \Sigma \end{aligned}$$

then, for some $\Pi \sqsupseteq \Sigma$,

$$\begin{aligned} \Pi &\vdash v : \tau \\ k &\models \Pi \end{aligned}$$

Proof. By induction over derivation of the operational semantics and the typing judgement.

We first consider the case where $\Gamma \vdash e : \tau$ has been derived by the subtyping rule. By rule inversion, we get $\Gamma \vdash e : \sigma$ and $\sigma <: \tau$. With this (smaller) derivation of the typing judgement, we can apply the theorem inductively, and get $\Pi \vdash v : \sigma$ and $k \models \Pi$ for some $\Pi \sqsupseteq \Sigma$. As $\sigma <: \tau$, we can deduce $\Pi \vdash v : \tau$.

In the following, we assume that the typing judgement has not been derived by a subtyping rule, and continue with a case distinction over the possible forms of the big-step semantics relation.

- $(s, h) \vdash x \Downarrow s(x), h$.

Then we have $\Gamma, x : \tau \vdash x : \tau$, and $k = h, v = s(x), \Gamma = [x \mapsto \tau]$. With $\Pi = \Sigma$, we get $k \models \Pi$. As $\Pi \vdash s : (\Gamma, x : \tau)$, we can deduce $\Pi \vdash v : \tau$.

- $(s, h) \vdash \text{null} \Downarrow \text{null}, h$.

Then we have $\Gamma \vdash \text{null} : \tau, v = \text{null}, k = h$. With $\Pi = \Sigma$, we get $\Pi \vdash v : \tau$ by definition of well-typed values, and $k \models \Pi$ follows directly from the assumption.

- $(s, h) \vdash \text{new } C \Downarrow l, h[l \mapsto (C, F)]$.

Then $\Gamma \vdash \text{new } C : C_{\{r\}}$ and $v = l$ and $k = h[l \mapsto (C, F)]$ where $l \notin \text{dom}(h)$ and $F = [f \mapsto \text{null}]_{f \in \text{fields}(C)}$. Since $\text{dom}(\Sigma) \subseteq \text{dom}(h)$, we have $l \notin \text{dom}(\Sigma)$. We choose $\Pi = \Sigma[l \mapsto (C, r)]$, and thus have $\Pi \vdash l : C_{\{r\}}$. To show $k \models \Pi$, it suffices to show $\Pi \models k(l) : C_{\{r\}}$. This holds trivially, as $k(l) = (C, F)$ and $C \preceq C$ and $F(f) = \text{null}$ for all $f \in \text{dom}(F)$.

- $(s, h) \vdash \text{let } x = e_1 \text{ in } e_2 \Downarrow v_2, h_2$.

Then the typing judgement is $\Gamma \vdash \text{let } x = e_1 \text{ in } e_2 : \tau$. Then by rule inversion of the typing and semantics rule, we get $\Gamma \vdash e_1 : \sigma$ and $\Gamma, x : \sigma \vdash e_2 : \tau$, as well as $(s, h) \vdash e_1 \Downarrow v_1, h_1$ and $(s[x \mapsto v_1], h_1) \vdash e_2 \Downarrow v_2, h_2$. By induction, we get $\Sigma_1 \vdash v_1 : \sigma$ and $h_1 \models \Sigma_1$ for some $\Sigma_1 \sqsupseteq \Sigma$. As s can include at most locations in $\text{dom}(\Sigma)$ and since Σ_1 is a disjoint extension of Σ , we also get $\Sigma_1 \vdash s : \Gamma$, and thereby $\Sigma_1 \vdash s[x \mapsto v_1] : \Gamma[x \mapsto \sigma]$. Again by induction, we get $\Sigma_2 \vdash v_2 : \tau$ and $h_2 \models \Sigma_2$ for some $\Sigma_2 \sqsupseteq \Sigma_1$. We can transitively deduce $\Sigma_2 \sqsupseteq \Sigma$, and have proven the case for $\Pi = \Sigma_2$.

- $(s, h) \vdash \text{if } x \text{ instanceof } E \text{ then } e_1 \text{ else } e_2 \Downarrow v, h'$.

The typing judgement is $\Gamma \vdash \text{if } x \text{ instanceof } E \text{ then } e_1 \text{ else } e_2 : \tau$. As there are two rules for the semantics of if-then-else constructs, we either get $(s, h) \vdash e_1 \Downarrow v, h'$ or $(s, h) \vdash e_2 \Downarrow v, h'$ by rule inversion. Without loss of generality, we only treat the first case here. By inversion of the typing rule, we get $\Gamma \vdash e_1 : \tau$. By induction, there exists a $\Pi \sqsupseteq \Sigma$ such that $\Pi \vdash v : \tau$ and $h' \models \Pi$.

- $(s, h) \vdash \text{if } x = y \text{ then } e_1 \text{ else } e_2 \Downarrow v, h'$.

The typing judgement is $\Gamma, x : C_R, y : D_S \vdash \text{if } x = y \text{ then } e_1 \text{ else } e_2 : \tau$. Let $s(x) = s(y)$. Then we get from the operational semantics $(s, h) \vdash e_1 \Downarrow v, h'$. We have to show that the variable context for the “then” branch is well-formed, i.e. that $\Sigma \vdash s(x) : C_{\{R \cap S\}}$ and $\Sigma \vdash s(y) : D_{\{R \cap S\}}$. If $s(x) = s(y) = \text{null}$, this holds trivially. If $s(x) = s(y) = l$, then we know $\Sigma(s(x)) = \Sigma(s(y)) = \Sigma(l) = (E, r)$, and $r \in R$ and $r \in S$, hence $r \in R \cap S$. By induction, there exists a $\Pi \sqsupseteq \Sigma$ such that $\Pi \vdash v : \tau$ and $h' \models \Pi$. If $s(x) \neq s(y)$, then variable context remains the same, therefore the induction can be applied immediately.

- $(s, h) \vdash x.f \Downarrow F(f), h$.

Then $e = x.f$ and $k = h$ and there is some $l \in \text{dom}(h)$ and some value v such that $s(x) = l$ and $h(l) = (D, F)$ and $F(f) = v$. The typing judgement is $\Gamma, x : C_R \vdash x.f : \tau$. With $\Pi = \Sigma$, we get $\Pi \vdash s : (\Gamma, x : C_R)$, which implies $\Pi \vdash s(x) : (C_R)$. We have $k \models \Pi$ and need to show $\Pi \vdash F(f) : \tau$. With $k \models \Pi$, we have $\Pi \models (D, F) : (C_R)$ and therefore by definition $D \preceq C$ and there is some $r \in R$ such that $\Pi \vdash F(f) : A^{\text{get}}(D, r, f)$. By well-typedness of the class table and premise of the typing rule, we get $A^{\text{get}}(D, r, f) < A^{\text{get}}(C, r, f) < \tau$, hence $\Pi \vdash F(f) : \tau$.

- $(s, h) \vdash x.f := y \Downarrow s(y), h'$.

Then $e = x.f := y$ and $v = s(y)$. The typing judgement is $\Gamma, x : C_R, y : \tau \vdash x.f := y : \tau$. We choose $\Pi = \Sigma$. With $\Sigma \vdash s : (\Gamma, x : C_R, y : \tau)$, we have $\Pi \vdash v : \tau$. By $h \models \Sigma$, we know $D \preceq C$. With the rule of the operational semantics, there is some $l \in \text{dom}(h)$ such that $s(x) = l$ and $h(l) = (D, F)$ and $k = h[l \mapsto (D, F[f \mapsto s(y)])]$. We still need to show $k \models \Pi$. Because k and h are identical with the exception of $k(l)(f)$, we need to show that the semantic heap relation is still preserved for the field value, i.e. that there is a region $r \in R$ such that $\Pi \vdash s(y) : A^{\text{get}}(C, r, f)$. With well-formedness of the class table and the premise of the rule, we have $\tau < A^{\text{set}}(C, r, f) < A^{\text{get}}(C, r, f)$, hence $\Pi \vdash s(y) : \tau$.

- $(s, h) \vdash x.m(\bar{y}) \Downarrow v, h'$.

By inversion of the rule for the operational semantics, we know there is a location $l \in \text{dom}(h)$ such that $s(x) = l$ and $h(l) = (D, -)$, and $(s', h) \vdash \text{mtable}(D, m) \Downarrow v, k$ where $s' = [\text{this} \mapsto l] \cup [x_i^m \mapsto s(y_i)]_{i \in \{1, \dots, |\bar{x}^m\}}$.

The typing judgement is $\Gamma, x : C_R, \bar{y} : \bar{\sigma} \vdash x.m(\bar{y}) : \tau$. As $\Sigma \vdash s : (\Gamma, x : C_R, \bar{y} : \bar{\sigma})$, we have $\Sigma \vdash l : C_R$. By definition of well-typed values, there must exist some class E and some class r with $\Sigma(l) = (E, r)$ such that $r \in R$ and $E \preceq C$. We can derive $\Sigma \vdash l : C_{\{r\}}$. With $h \models \Sigma$, we can also infer $D \preceq C$.

By inversion of the typing rule, we know there is some method typing $(\bar{\sigma}', \tau') \in M(C, r, m)$ such that $\bar{\sigma} < \bar{\sigma}'$ and $\tau' < \tau$. As $D \preceq C$ and the class table is well-formed, there is a method typing $(\bar{\sigma}'', \tau'') \in M(D, r, m)$ such that $\bar{\sigma}' < \bar{\sigma}''$ and $\tau'' < \tau'$. From $\Sigma \vdash s : (\Gamma, x : C_R, \bar{y} : \bar{\sigma})$, it follows $\Sigma \vdash s(y_i) : \sigma_i''$ for $i \in \{1, \dots, |\bar{x}^m\}$.

As the FJEU program is well-typed, we get $\Gamma' \vdash \text{mtable}(D, m) : \tau''$ where $\Gamma' = [\text{this} \mapsto C_{\{r\}}] \cup [x_i^m \mapsto \sigma_i'']_{i \in \{1, \dots, |\bar{x}^m\}}$. With the facts from above, we get $\Sigma \vdash s' : \Gamma'$, so we can finally apply the theorem inductively on the derivation of the semantics and get $\Pi \vdash v : \tau''$ and $k \models \Pi$ for some $\Pi \sqsupseteq \Sigma$. From $\tau'' < \tau$ follows $\Pi \vdash v : \tau$.

□

B FJEU Pointer Analysis with Datalog

We specify a pointer analysis algorithm for the FJEU language in the style of [31]. To simplify matters, we omit class types, and assume the program in question has been verified by some other (class) type checking facility before. Methods are supposed to have only one argument variable x^m . There is no class type analysis, i.e. all implementations of a method are regarded as possible targets of a method invocation.

The Datalog rules look as follows:

$$\begin{aligned}
eP(z, i, r) &: - \text{var}(i, x), vP(z, x, r) \\
eP(z, i, r_2), vP(z, x, r_1) &: - \text{let}(i, x, i_1, i_2), eP(z, i_1, r_1), eP(z, i_2, r_2) \\
eP(z, i, r_1), eP(z, i, r_2) &: - \text{if}(i, i_1, i_2), eP(z, i_1, r_1), eP(z, i_2, r_2) \\
eP(z, i, r) &: - \text{new}(i), \underline{r = i} \\
eP(z, i, r_2) &: - \text{get } f(i, x, f), vP(z, x, r_1), hP(z, r_1, f, r_2) \\
eP(z, i, r_1), hP(z, r_1, f, r_2) &: - \text{put } f(i, x, f, z), vP(z, x, r_1), vP(z, y, r_2) \\
eP(z, i, r), mT(r_1, m, z', r_2, r) &: - \text{call}(i, m, x, y), vP(z, x, r_1), vP(z, y, r_2), \\
&\quad \text{body}(m, i_{\text{body}}), eP(z', i_{\text{body}}, r), \underline{IE_c(z, i, z', m)} \\
vP(z, x^m, r) &: - mT(r, m, z, r_v, r_r)
\end{aligned}$$

where

- i, r, z range over finite sets of expression labels, regions, and contexts
- $\text{var}(i, x), \text{let}(i, x, i_1, i_2), \text{new}(i), \dots$: the expression at i is a variable/let expression/new expression/...
- $eP(z, i, r)$: in context z , the evaluation of the expression at i may point to r
- $vP(z, x, r)$: in context z , variable x may point to region r
- $hP(z, r_1, f, r_2)$: in context z , field f of objects in region r_1 may point to region r_2 ; this corresponds to A^{get} and A^{set}
- $mT(r, m, z, r_v, r_r)$: the type for method m for context z of objects in region r is (r_v, r_r) ; this corresponds to M
- $\text{body}(m, i)$: the expression with label i is an implementation of method m

The underlined sections determine the abstraction principles of the analysis, and show how the abstraction functions of the type system must be instantiated to interpret the result of the analysis correctly: $\psi(z, i) = i$ and $\phi(C, z, r, m, i) = z' \iff IE_c(z, i, z', m)$, where IE_c is the abstract call graph relation of [31].