

Types and Trace Effects for Object Orientation

Christian Skalka
The University of Vermont

Abstract. *Trace effects* are statically generated program abstractions, that can be model checked for verification of assertions in a temporal program logic. In this paper we develop a type and effect analysis for obtaining trace effects of Object Oriented programs in Featherweight Java. We observe that the analysis is significantly complicated by the interaction of trace behavior with inheritance and other Object Oriented features, particularly overridden methods, dynamic dispatch, and downcasting. We propose an expressive type and effect inference algorithm combining polymorphism and subtyping/subeffecting constraints to obtain a flexible trace effect analysis in this setting, and show how these techniques are applicable to Object Oriented features. We also extend the basic language model with exceptions and stack-based event contexts, and show how trace effects scale to these extensions by structural transformations.

Keywords: Object Oriented Languages, type and effect theory, temporal program logics.

1. Introduction

Program type analysis and model checking have a shared goal: to statically enforce properties of programs. A variety of analyses have been proposed to enforce specific properties, while certain frameworks provide flexibility to enforce a general class of properties. A number of authors (Skalka et al., 2005; K. Marriott and Sulzmann, 2003; Igarashi and Kobayashi, 2002) have observed that these two approaches can play complementary roles in the verification of general *trace based* program properties, which are properties of program *event traces* expressible in temporal logic. Type systems can be used to compute program abstractions, which can in turn be used as inputs to model checking (Steffen and Burkart, 1992). In other words, type analysis can serve as a technique for model extraction (Holzmann and Smith, 2001), for the subsequent verification of a general class of program properties. This paper establishes a foundational theory of type and trace effects for Object Oriented programs in a language model adapted from Featherweight Java (Igarashi et al., 2001), defines a type inference system for automatically reconstructing sound type and trace effects of programs, and shows how trace effect representations can be manipulated to reflect control flow operations such as exceptions.

© 2006 Kluwer Academic Publishers. Printed in the Netherlands.

Trace based program properties are properties of event traces, where events are records of program actions, explicitly inserted into program code either manually (by the programmer) or automatically (by the compiler). Events are intended to be sufficiently abstract to represent a variety of program actions—e.g. opening a file, access control privilege activation, or entry to or exit from critical regions. Event traces maintain the ordered sequences of events that occur during program execution. Assertions enforce properties of event traces—e.g. certain privileges should be activated before a file can be opened. Results in (Skalka and Smith, 2004; K. Marriott and Sulzmann, 2003; Bartoletti et al., 2005c) have demonstrated that static approximations of program event traces can be generated by type and effect analyses (Talpin and Jouvelot, 1992; Amtoft et al., 1999), in a form amenable to existing model-checking techniques for verification. We call these approximations *trace effects*.

Trace based analyses have been shown capable of statically enforcing flow-sensitive security properties such as safe locking behavior (Foster et al., 2002) and resource usage policies such as file usage protocols and memory management (K. Marriott and Sulzmann, 2003; Igarashi and Kobayashi, 2002). In (Bartoletti et al., 2005a), a trace effect analysis is used to enforce secure service composition. The history-based access control model of (Abadi and Fournet, 2003) can be implemented with event traces and checks (Skalka and Smith, 2004), as can be the policies realizable in that model, e.g. sophisticated Chinese Wall policies (Abadi and Fournet, 2003). Stack-based security policies are also amenable to this form of analysis, as shown in (Skalka and Smith, 2004; Skalka et al., 2005) and this paper. In short, the combination of a primitive notion of program events with a temporal program logic for asserting properties of event traces yields a powerful and general tool for enforcing program properties.

The analyses cited above have been developed in functional language settings, but practical use of these tools require adaptation to realistic languages. In this paper we address technical considerations for application of trace effects to Object Oriented languages, particularly Java. As discussed more thoroughly in Sect. 2, inheritance, dynamic dispatch, and downcasts present significant challenges to trace effect analysis. The effect is that scaling the analysis to Object Oriented programs is a foundational problem, not simply an engineering one. To study these issues in isolation, we extend Featherweight Java (FJ) (Igarashi et al., 2001) with events, traces, and checks, and a polymorphic type and effect inference analysis for static enforcement of checks, yielding the language FJ_{trace} . Technical results in this paper extend and enhance material presented in (Skalka, 2005) and (Skalka et al., 2005).

1.1. A FLEXIBLE TYPE ANALYSIS

Type theory provides a variety of useful tools in this setting. As will be discussed in detail in later sections, a combination of parametric polymorphism and subtyping can be used to provide the right abstractions and flexibility for addressing issues associated with inheritance and dynamic dispatch, and type constraint representation allows a precise treatment of object downcasting. We extend an Object Oriented type language with *trace effects* that approximate trace behavior of programs. Trace effects are a form of label transition system (LTS), which are amenable to model checking. Therefore, the type analysis serves as a technique for extracting a verifiable abstraction of program trace behavior, with type inference automating the process. We also show that trace effect representations are amenable to transformations that reflect the impact of control flow modifications on trace behavior, particularly exceptions.

The metatheory of types also provides an appealing language for characterizing the analysis and proving its correctness. A type safety result guarantees that programs satisfying the analysis will not have run-time errors, in particular all specified properties of program traces are guaranteed to hold. This result is established via subject reduction and progress arguments. The type and effect system is shown to be a conservative extension of the underlying Featherweight Java type system, ensuring backwards compatibility with existing programs. We show that type inference is sound, and so-called *trace approximation* is demonstrated, formalizing the idea that trace effects conservatively approximate program trace behavior. We also develop an extended language model with exceptions, and show that a post-processing transformation of inferred trace effects is sufficient to capture the effect of exceptions on control flow, via another type safety result.

1.2. OUTLINE OF THE PAPER

The remainder of the paper is organized as follows. In Sect. 2, the central issues of our type and effect analysis in relation to Object Oriented programming are described and discussed, clarifying the contribution of this paper. In Sect. 3, the FJ_{trace} language is formally defined, which is FJ extended with primitives for a security logic of program traces. In Sect. 4, we formalize the language and meaning of trace effects. In Sect. 5 a logical type system for FJ_{trace} is presented, with features and examples discussed in Sect. 5.2.1 and Sect. 5.3.1, and properties including type safety proved in Sect. 5.4. A type inference algorithm is defined in Sect. 6, that is shown to be sound with respect to the logical type system in Sect. 6.4, implying type safety in the implementation.

In Sect. 7, we study variations on the basic language model, including exceptions in Sect. 7.1 and stack based trace contexts in Sect. 7.2. We conclude with more discussion of related work and a final summary in Sect. 8.

2. Trace Effects and Object Orientation

Subtyping is a common discipline for relating behavior of objects in an inheritance hierarchy. However, as we illustrate below, imposing a subsumption relation on the trace behavior of methods in an inheritance hierarchy is overly restrictive for applications such as access control. It is possible and useful to extend the definition of subtyping to trace effects, as we do in Sect. 5, but a realistic analysis requires that we develop some mechanism for allowing independence of inheritance and effects, and accommodate this independence in the presence of dynamic dispatch. We propose the use of *parametric polymorphism* for this purpose. We also propose a *type constraint* representation; along with known benefits of this approach in application to Object Oriented programming (Eifrig et al., 1995; Bruce et al., 1995), we show how type constraints can be used for a novel soft-typing of downcasts. In this section we discuss and illustrate these issues, before providing formal details in Sect. 5.

2.1. EFFECTS AND INHERITANCE

The manner in which inheritance and dynamic dispatch complicates trace effect analysis is best illustrated by example. Consider the application of event traces to enforce a history-based access control mechanism, as in (Skalka and Smith, 2004; Abadi and Fournet, 2003), where code is statically signed by its owner (an authorization event), and a local access control list \mathcal{A} associates owners with their authorizations. A **demand** predicate ensures that the intersection of all authorizations encountered up to the point of the check contains a specified privilege. Let \mathbf{r} denote the specified privilege, and let \mathbf{P} range over owners. Thus, given an authorization trace $\mathbf{P}_1; \dots; \mathbf{P}_n$, we require that $\mathbf{r} \in \mathcal{A}(\mathbf{P}_1) \cap \dots \cap \mathcal{A}(\mathbf{P}_n)$ in order for the trace to satisfy **demand**(\mathbf{r}). Note that standard practice allows class owners to extend classes owned by others, i.e. ownership need not be consistent throughout an inheritance hierarchy. In our encoding, distinct ownership implies method override, at least with regard to authorization events; an accurate analysis therefore requires independence of trace effects of different method versions in the inheritance hierarchy.

The issue is complicated further by dynamic dispatch. Imagine a class `Writer` that implements a `safewrite` method signed with a `System` authorization event, where `safewrite` takes a `Formatter` and a `File` as arguments, and requires that the `FileWrite` privilege be active before writing the formatter output to the file, via an access control check `demand(FileWrite)`. Note especially that the specification of the check requires that `FileWrite` must be among the authorizations of the `x.format` method, since these will affect the flow of control and therefore appear in the `safewrite` event trace:

```
class Writer extends Object {
    void safewrite(Formatter x, File f){
        System;
        String s = x.format()
        demand(FileWrite);
        write(s, f);
    }
}
```

Thus, we can statically approximate the trace generated by the method `safewrite` as:

```
System; H; demand(FileWrite)
```

where `H` represents the trace effect approximation of `x.format()`. The central issue is, what is `H`? Note that in a language with dynamic dispatch such as Java, the trace generated by `x.format()` could be generated by any version of `format` among the subclasses of `Formatter`, so it is unsound to imagine `H` as just the approximation of `Formatters` version.

In the FJ subtyping system, the type `Formatter` subsumes its subclass types via the subtyping relation. So as a first approximation, we can imagine extending subtyping to trace effects, meaning that `H` should subsume the effect of the `format` method in the `Formatter` class, *as well as the effects of every format method in Formatter subclasses*. This approach is taken in (Higuchi and Ohori, 2006) as part of a related analysis for Java stack inspection, and a standard definition of effect subtyping (Skalka and Smith, 2004; Amtoft et al., 1999; K. Marriott and Sulzmann, 2003) approximates the effect of `x.format()` as the non-deterministic choice of trace effects of the `format` method in every `Formatter` subclass. For example, suppose that there exist only two such classes, one which is owned by `System`, and the other which is owned by an `Applet`. This would be implemented by prepending the former's `format` method with a `System` event, and the latter's with an `Applet` event. For simplicity, we assume that these methods are

otherwise event-free. In this case, we would have $H = \text{System|Applet}$, where $|$ is a choice constructor. But since it is natural to assume that **Applets** are not **FileWrite** authorized, verification of:

$$\text{System}; (\text{System|Applet}); \text{demand}(\text{FileWrite})$$

will fail. This means that *any* invocation of **safewrite** would be statically rejected, even if invoked with a **System** formatter. Further, the scheme requires the entire **Formatter** class hierarchy be known in advance for static analysis, since any addition would require re-computation of its effects. This would disallow modularity.

We address this problem by using polymorphism, rather than subtyping, to approximate the effects of method parameters; to wit, the effect H in question is represented by a universally quantified type variable. This is accomplished via the object type form $[TC]$, where T contains the inferred type and effects of a given object's methods, and C is the declared object class—so that the nominal type language of FJ_{trace} is “superimposed” over the type language of FJ , as a conservative extension of the latter (as is discussed more extensively in Sect. 5). Let **StringT**, and **FileT** be the types of **String**, and **File** objects respectively, the details of which are unimportant to the example. Then, when typing the **safewrite** method in the **Writer** class, the FJ_{trace} type system will assign an abstract effect h to its x parameter, as in the following type we abbreviate as **AbsFormatterT**:

$$\text{AbsFormatterT} \triangleq [\text{format} : () \xrightarrow{h} \text{StringT} \text{ Formatter}]$$

and **safewrite** may be assigned the type we abbreviate as T :

$$T \triangleq (\text{AbsFormatterT}, \text{FileT}) \xrightarrow{\text{System}; h; \text{demand}(\text{FileWrite})} \text{void}$$

and the typing $\text{Writer} : \forall h. [\text{safewrite} : T \text{ Writer}]$ may be assigned, where the abstract effect h of $x.\text{format}()$ is quantified. At specific application points, h can then be instantiated with the accurate trace effect of the substituant of x . This example is extended and discussed in Sect. 5.3.1 following formal development of the type system.

2.2. CONSTRAINT SUBTYPING AND CASTING

To maintain decidability in the type system, we propose only first-order parametric polymorphism. This means that if x is a formal parameter of some method m , any method $x.m'$ cannot be invoked within m in a polymorphic fashion. To obtain the flexibility necessary to statically allow application of abstracted methods to objects of multiple types, we propose a subtyping relation, similar to that discussed above, that

can be used where parametric polymorphism cannot due to first-orderly restrictions.

A number of considerations motivate subtyping in our type and effect system, beyond the fact that it integrates neatly with FJ subtyping analysis. Firstly, while a top-level effect weakening rule, as in (Skalka and Smith, 2004), is sufficient for a flexible type and effect analysis, a subtyping rule that incorporates weakening of latent effects on function types is more precise and complete, as observed in (Amtoft et al., 1999). Also, we implement subtyping via a recursive constraint representation, which has been shown to allow precise typing of common object-oriented idioms, such as binary methods (Eifrig et al., 1995; Bruce et al., 1995).

A constraint type representation also supports a soft typing analysis of downcasts, which combines static and dynamic checks to ensure soundness (Cartwright and Fagan, 1991). For example, suppose some expression e has a type T , where T is constrained to be a supertype of both `Triangle` and `Polygon` objects, where `Triangle` is a subclass of `Polygon`, and where R and S are the field and method types of the `Triangle` and `Polygon` objects, respectively:

$$[R \text{ Triangle}] <: T \qquad [S \text{ Polygon}] <: T$$

Then, given the cast $(\text{Triangle})e$, we first observe that an FJ dynamic cast check will ensure that any run-time scenario in which e evaluates to an object strictly in a superclass of `Triangle` will be stuck (Igarashi et al., 2001). Guided by the intuition that constraints represent data flow paths, we further observe that any constraint representing flow of an object strictly in a superclass of `Triangle` to the program point represented by T can be ignored when analyzing the cast $(\text{Triangle})e$, without compromising type safety, since this unsafe flow will be caught at run time by a dynamic cast check. In our type analysis, we implement this idea by positing a type T' such that:

$$(\text{Triangle})e : [T' \text{ Triangle}]$$

with the condition that T be a *soft subtype* of T' , written:

$$T < [T' \text{ Triangle}]$$

This means that if some object type $[UC]$ is a subtype of T , then $[UC] <: [T' \text{ Triangle}]$ only if $C <: \text{Triangle}$; see Definition 5.2 for a formalization of the idea. This implies:

$$[R \text{ Triangle}] <: [T' \text{ Triangle}]$$

and *not* $[S \text{ Polygon}] <: [T' \text{ Triangle}]$. Note that requiring the latter would yield an inconsistent constraint set in any case, so a benefit of this

$L ::= \text{class } C \text{ extends } C \{ \bar{C} \bar{f}; K \bar{M} \}$	<i>class definitions</i>
$K ::= C(\bar{C} \bar{f}) \{ \text{super}(\bar{f}); \text{this}.\bar{f} = \bar{f}; \}$	<i>constructors</i>
$M ::= C \text{ m}(\bar{C} \bar{x}) \{ \text{return } e; \}$	<i>methods</i>
$e ::= x \mid e.f \mid e.m(\bar{e}) \mid \text{new } C(\bar{e}) \mid (C)e \mid \text{ev}[i] \mid \text{chk}[i]$	<i>expressions</i>
$\eta ::= \epsilon \mid \eta; \eta \mid \text{ev}[i] \mid \text{chk}[i]$	<i>traces</i>

Figure 1. FJ_{trace} language syntax

$C <: C$	$\frac{B <: C \quad C <: D}{B <: D}$	$\frac{CT(C) = \text{class } C \text{ extends } D \{ \dots \}}{C <: D}$
----------	--------------------------------------	---

Figure 2. Nominal subtyping for FJ

approach is completeness in the presence of dynamically checked downcasts. We believe that a constraint representation yields a distinctly precise type analysis; it is hard to see how the precision obtainable by selective pruning of the constraint graph used in our implementation of soft subtyping (Sect. 6) can be recreated with e.g. a unification-based approach. Subtyping and soft subtyping is further discussed in Sect. 5.2.1 following formal development of the interpretation of subtyping constraints.

3. The Language FJ_{trace}

In this section we define the syntax and semantics of FJ_{trace} , by extending the syntax of FJ with primitives for specifying events and local checks, and the semantics of FJ with event traces as configuration components.

3.1. SYNTAX

The syntax of FJ_{trace} is defined in Fig. 1. We let A, B, C, D range over class names, f range over field names, m range over method names, x, y range over both field and method names, and d, e range over expressions. *Values*, denoted v or u , are objects, i.e. expressions of the form $\text{new } C(v_1, \dots, v_n)$. A *class table* CT is a mapping from class names C to definitions L , and a *program* is a pair (CT, e) ; for brevity in the

following, we assume a fixed class table CT . As in (Igarashi et al., 2001) we assume **Object** values that have no fields or methods; we let $()$ denote the value $\text{new Object}()$. Since constructor definitions K are hard-coded by the language grammar, for brevity we do not explicitly include them in example class definitions in this paper.

The essential distinguishing features of FJ_{trace} are *events* $\text{ev}[i]$ and *local checks* $\text{chk}[i]$. Both events and checks are distinguished by labels i . Events and checks encountered during execution are accrued in linear order in *traces* η , with execution blocking if unsuccessful checks are encountered. Events and checks are therefore side-effecting instructions; the value $()$ is the direct evaluation result of checks and events, as specified in the next section.

In this presentation we leave the logic of checks abstract, specifying only that checks are predicates on traces, and write $\eta \vdash \text{chk}[i]$ to denote that η satisfies $\text{chk}[i]$. We could for example instantiate the language of checks with the linear mu-calculus, as in (Skalka and Smith, 2004), but in this presentation we are mainly concerned with typing.

3.1.1. Vector Notations

For brevity in numerous instances, we adopt the vector notations of (Igarashi et al., 2001). We write $\bar{\mathbf{f}}$ to denote the sequence $\mathbf{f}_1, \dots, \mathbf{f}_n$, similarly for $\bar{\mathbf{C}}, \bar{\mathbf{m}}, \bar{\mathbf{x}}, \bar{\mathbf{e}}$, etc., and we write $\bar{\mathbf{M}}$ as shorthand for $\mathbf{M}_1 \cdots \mathbf{M}_n$. We write the empty sequence as \emptyset , and we write $|\bar{\mathbf{x}}|$ to denote the length of $\bar{\mathbf{x}}$. If and only if \mathbf{m} is one of the names in $\bar{\mathbf{m}}$, we write $\mathbf{m} \in \bar{\mathbf{m}}$, similarly for $\mathbf{f} \in \bar{\mathbf{f}}$. Given some $\bar{\mathbf{f}}$, we write \mathbf{f}_i to denote the i th element of $\bar{\mathbf{f}}$. Vector notation is also used to abbreviate sequences of declarations; we let $\bar{\mathbf{C}} \bar{\mathbf{f}}$ and $\bar{\mathbf{C}} \bar{\mathbf{f}}$; denote $\mathbf{C}_1 \mathbf{f}_1, \dots, \mathbf{C}_n \mathbf{f}_n$ and $\mathbf{C}_1 \mathbf{f}_1; \dots; \mathbf{C}_n \mathbf{f}_n$; respectively, and we write $\mathbf{C} \mathbf{f} \in \bar{\mathbf{C}} \bar{\mathbf{f}}$ iff $\mathbf{C} \mathbf{f}$ is one of the declarations in $\bar{\mathbf{C}} \bar{\mathbf{f}}$. The notation $\text{this}.\bar{\mathbf{f}} = \bar{\mathbf{f}}$; abbreviates the sequence of initializations $\text{this}.\mathbf{f}_1 = \mathbf{f}_1; \dots; \text{this}.\mathbf{f}_n = \mathbf{f}_n$; Sequences of names and declarations are assumed to contain no duplicate names.

3.2. OPERATIONAL SEMANTICS

The operational semantics of FJ_{trace} are defined in Fig. 3. The small-step reduction relation \rightarrow is defined on closed *configurations*, which are pairs of traces and expressions η, \mathbf{e} . As in (Igarashi et al., 2001), we divide the operational rules into *computation* and *congruence* rules; the former (resp. latter) are those whose names are prefixed by R- (resp. RC-). The semantics are defined in terms of a number of auxiliary functions and a nominal subtyping relation $<$: taken from (Igarashi et al., 2001) and recalled in Fig. 2 and Fig. 5. We let \rightarrow^* denote the reflexive, transitive closure of \rightarrow .

$$\begin{array}{c}
\text{R-FIELD} \\
\frac{\text{fields}(\mathbf{C}) = \bar{\mathbf{c}} \bar{\mathbf{f}}}{\eta, (\text{new } \mathbf{C}(\bar{\mathbf{v}})).\mathbf{f}_i \rightarrow \eta, \mathbf{v}_i} \\
\\
\text{R-INVK} \\
\frac{\text{mbody}(\mathbf{m}, \mathbf{C}) = \bar{\mathbf{x}}.\mathbf{e}}{\eta, (\text{new } \mathbf{C}(\bar{\mathbf{v}})).\mathbf{m}(\bar{\mathbf{u}}) \rightarrow \eta, [\bar{\mathbf{u}}/\bar{\mathbf{x}}, \text{new } \mathbf{C}(\bar{\mathbf{v}})/\text{this}]\mathbf{e}} \\
\\
\text{R-CAST} \qquad \mathbf{C} <: \mathbf{D} \qquad \qquad \qquad \text{R-EVENT} \\
\frac{}{\eta, (\mathbf{D})(\text{new } \mathbf{C}(\bar{\mathbf{v}})) \rightarrow \eta, \text{new } \mathbf{C}(\bar{\mathbf{v}})} \qquad \eta, \mathbf{ev}[\mathbf{i}] \rightarrow \eta; \mathbf{ev}[\mathbf{i}], () \\
\\
\text{R-CHECK} \\
\frac{\eta \vdash \mathbf{chk}[\mathbf{i}]}{\eta, \mathbf{chk}[\mathbf{i}] \rightarrow \eta; \mathbf{chk}[\mathbf{i}], ()} \\
\\
\text{RC-FIELD} \qquad \qquad \qquad \text{RC-INVK-RECV} \\
\frac{\eta, \mathbf{e} \rightarrow \eta', \mathbf{e}'}{\eta, \mathbf{e}.\mathbf{f} \rightarrow \eta', \mathbf{e}'.\mathbf{f}} \qquad \frac{\eta, \mathbf{e} \rightarrow \eta', \mathbf{e}'}{\eta, \mathbf{e}.\mathbf{m}(\bar{\mathbf{e}}) \rightarrow \eta', \mathbf{e}'.\mathbf{m}(\bar{\mathbf{e}})} \\
\\
\text{RC-INVK-ARG} \\
\frac{\eta, \mathbf{e}_i \rightarrow \eta', \mathbf{e}'_i}{\eta, \mathbf{v}.\mathbf{m}(\bar{\mathbf{v}}, \mathbf{e}_i, \bar{\mathbf{e}}) \rightarrow \eta', \mathbf{v}.\mathbf{m}(\bar{\mathbf{v}}, \mathbf{e}'_i, \bar{\mathbf{e}})} \\
\\
\text{RC-NEW-ARG} \qquad \qquad \qquad \text{RC-CAST} \\
\frac{\eta, \mathbf{e}_i \rightarrow \eta', \mathbf{e}'_i}{\eta, \text{new } \mathbf{C}(\bar{\mathbf{v}}, \mathbf{e}_i, \bar{\mathbf{e}}) \rightarrow \eta', \text{new } \mathbf{C}(\bar{\mathbf{v}}, \mathbf{e}'_i, \bar{\mathbf{e}})} \qquad \frac{\eta, \mathbf{e} \rightarrow \eta', \mathbf{e}'}{\eta, (\mathbf{C})\mathbf{e} \rightarrow \eta', (\mathbf{C})\mathbf{e}'}
\end{array}$$

Figure 3. FJ_{trace} operational semantics

The operational semantics are largely the same as FJ, with the addition of run-time traces and the treatment of events and checks. The rules that directly affect the trace include R-EVENT, which appends an event $\mathbf{ev}[\mathbf{i}]$ encountered during execution to the end of the trace. The R-CHECK rule is defined similarly, except the check $\mathbf{chk}[\mathbf{i}]$ is required to be satisfied by the current trace, otherwise computation becomes stuck. Each of the congruence rules propagates changes to the trace effected by the reduction of subterms.

$$\begin{array}{c}
\text{ev}[\mathbf{i}] \xrightarrow{\text{ev}[\mathbf{i}]} \epsilon \qquad \text{chk}[\mathbf{i}] \xrightarrow{\text{chk}[\mathbf{i}]} \epsilon \qquad \epsilon; \mathbf{H} \xrightarrow{\epsilon} \mathbf{H} \\
\mathbf{H}_1; \mathbf{H}_2 \xrightarrow{a} \mathbf{H}'_1; \mathbf{H}_2 \quad \text{if } \mathbf{H}_1 \xrightarrow{a} \mathbf{H}'_1 \qquad \mathbf{H}_1 | \mathbf{H}_2 \xrightarrow{\epsilon} \mathbf{H}_1 \qquad \mathbf{H}_1 | \mathbf{H}_2 \xrightarrow{\epsilon} \mathbf{H}_2 \\
\mu \mathbf{h}. \mathbf{H} \xrightarrow{\epsilon} \mathbf{H}[\mu \mathbf{h}. \mathbf{H} / \mathbf{h}] \\
\llbracket \mathbf{H} \rrbracket = \begin{array}{c} \{a_1 \cdots a_n \mid \mathbf{H} \xrightarrow{a_1} \cdots \xrightarrow{a_n} \mathbf{H}'\} \\ \cup \\ \{a_1 \cdots a_n \downarrow \mid \mathbf{H} \xrightarrow{a_1} \cdots \xrightarrow{a_n} \epsilon\} \end{array}
\end{array}$$

Figure 4. Interpretation of trace effects

4. Semantics of Trace Effects

The aim of our analysis is to statically guarantee the satisfaction of run-time checks in programs. To this end, our analysis infers an approximation of the trace that will be generated during program execution, by reconstructing the *trace effect* of programs. In essence, trace effects \mathbf{H} conservatively approximate traces η that may develop during execution, by representing a set of traces containing at least η . The grammar of trace effects is given in Fig. 6. A trace effect may be an event $\text{ev}[\mathbf{i}]$ or check $\text{chk}[\mathbf{i}]$, or a sequencing of trace effects $\mathbf{H}_1; \mathbf{H}_2$, a nondeterministic choice of trace effects $\mathbf{H}_1 | \mathbf{H}_2$, or a μ -bound trace effect $\mu \mathbf{h}. \mathbf{H}$ which finitely represents the set of traces that may be generated by a recursive function. Noting that the syntax of traces η is the same as linear, variable-free trace effects, we abuse syntax and let η also range over linear, variable-free trace effects.

We define a Labeled Transition System (LTS) interpretation of trace effects as sets of *abstract traces*, which are strings of events and checks that may be delimited by a \downarrow symbol to denote termination. The interpretation is defined via strings denoted θ , characterized by the following grammar, where ϵ denotes the empty string in the usual manner and the alphabet consists of events and checks:

$$\begin{array}{l}
a ::= \text{ev}[\mathbf{i}] \mid \text{chk}[\mathbf{i}] \mid \epsilon \\
s ::= a \mid s s \\
\theta ::= s \mid s \downarrow
\end{array}$$

The interpretation of an effect \mathbf{H} , denoted $\llbracket \mathbf{H} \rrbracket$ and specified in Fig. 4, is then taken to be the prefix-closed, finite approximation of the trace

sets that may be generated by H , when viewed as a program in a transition semantics defined by relations \xrightarrow{a} on closed effects. Since programs may not terminate, their traces may be infinite, but this possibility is captured via finite approximation— all finite prefixes of an infinite trace are contained in the interpretation. Trace effect equivalence is defined via the interpretation, i.e. $H_1 = H_2$ iff $\llbracket H_1 \rrbracket = \llbracket H_2 \rrbracket$. This relation is in fact undecidable: trace effects are equivalent to BPA's (basic process algebras) (Skalka and Smith, 2004), and their trace equivalence is known to be undecidable (Burkart et al., 2001). Note that prefix closure does not cause any loss of information in the interpretation, since the post-pending of \downarrow to terminating traces allows them to be distinguished from their prefixes. In particular, it is not necessarily true that $\llbracket H \rrbracket \subseteq \llbracket H; H' \rrbracket$ for arbitrary H and H' .

Trace effects predict the history of both events and checks. In order to ensure that program checks will succeed, it is not sufficient for trace effects to approximate run time traces, but also success of these checks must be guaranteed. We base *validity* of trace effects on validity of checks that occur in traces in its interpretation. In particular, for any given check in a trace, that check must hold for its prefix:

DEFINITION 4.1. *H is valid iff for all $(a_1 \cdots a_n \text{chk}[i]) \in \llbracket H \rrbracket$ it is the case that $a_1; \dots; a_n \vdash \text{chk}[i]$ holds.*

This definition is logically sufficient to obtain our desired type safety result. It is not algorithmic, but trace effects are equivalent to basic process algebras (BPAs) as observed in (Skalka et al., 2006), for which known model-checking techniques exist (Burkart et al., 2001), allowing automated verification of effect validity (Skalka and Smith, 2004). In this paper, we focus on type and effect inference, rather than the effect verification component of our analysis. For details on the latter, the reader is referred to (Skalka and Smith, 2004; Skalka et al., 2006).

4.1. PROPERTIES

Various properties of trace effect equivalence are enumerated as follows. The equivalences will be exploited for brevity and clarity in examples throughout the text, as well as for later proofs:

LEMMA 4.1. *We note the following properties of trace effect equivalence for closed H , H_1 , H_2 , and H_3 :*

1. $H|H = H$
2. $\epsilon; H = H = H; \epsilon$
3. $\mu h.H = H$

4. $H_1|H_2 = H_2|H_1$
5. $H_1;(H_2;H_3) = (H_1;H_2);H_3$
6. $H_1|(H_2|H_3) = (H_1|H_2)|H_3$
7. $H_1;(H_2|H_3) = (H_1;H_2)|(H_1;H_3)$
8. $(H_1|H_2);H_3 = (H_1;H_3)|(H_2;H_3)$
9. *Trace effect equivalence is homomorphic for all constructors.*

We also note some properties related to trace effect interpretation containment; these properties are important, since our type analyses will allow *weakening* of trace effects for flexibility. That is, if a trace effect H approximates the traces generated by a program, and $\llbracket H \rrbracket \subseteq \llbracket H' \rrbracket$, then H' is also a sound approximation. Like equality, containment is known to be undecidable.

LEMMA 4.2. *Writing $H <: H'$ iff $\llbracket H \rrbracket \subseteq \llbracket H' \rrbracket$, the following properties hold for arbitrary closed H, H_1 , and H_2 :*

1. $H <: H_1$ is undecidable.
2. $H <: H|H_1$
3. For closed $\mu h.H_0$, we have that $H_0[\mu h.H_0/h] <: \mu h.H$.
4. If $H <: H_1$ then $H|H_2 <: H_1|H_2$ and $H_2;H <: H_2;H_1$ and $H;H_2 <: H_1;H_2$.
5. If $H_1 <: H_2$ then validity of H_2 implies validity of H_1 .

5. Types for FJ_{trace}

Featherweight Java is equipped with a declarative, nominal type system; the type language is based on class names, which annotate function return and argument types, casts, and object creation points. Method and field types are not explicit in the FJ type of objects, which are just class names, but rather can be looked up given the object class name and its definition in the class table. The lookup functions *mtype* and *fields* are defined in Fig. 5. The FJ type system is algorithmically checkable, and enjoys a type safety result (Igarashi et al., 2001). Our intent is to not to redo the type system of FJ, but to “superimpose” a type and effect analysis on it, thereby subsuming type safety for the FJ subset of FJ_{trace} . This superimposition should be conservative and

$$\begin{array}{c}
fields(\text{Object}) = \emptyset \\
\hline
CT(\mathbf{C}) = \text{class } \mathbf{C} \text{ extends } \mathbf{D} \{ \bar{\mathbf{C}} \bar{\mathbf{f}}; \mathbf{K} \bar{\mathbf{M}} \} \quad fields(\mathbf{D}) = \bar{\mathbf{D}} \bar{\mathbf{g}} \\
\hline
fields(\mathbf{C}) = \bar{\mathbf{D}} \bar{\mathbf{g}}, \bar{\mathbf{C}} \bar{\mathbf{f}} \\
\\
CT(\mathbf{C}) = \text{class } \mathbf{C} \text{ extends } \mathbf{D} \{ \bar{\mathbf{C}} \bar{\mathbf{f}}; \mathbf{K} \bar{\mathbf{M}} \} \\
\mathbf{B} \ \mathbf{m}(\bar{\mathbf{B}} \bar{\mathbf{x}}) \{ \text{return } \mathbf{e}; \} \in \bar{\mathbf{M}} \\
\hline
mtype(\mathbf{m}, \mathbf{C}) = \bar{\mathbf{B}} \rightarrow \mathbf{B} \\
\\
CT(\mathbf{C}) = \text{class } \mathbf{C} \text{ extends } \mathbf{D} \{ \bar{\mathbf{C}} \bar{\mathbf{f}}; \mathbf{K} \bar{\mathbf{M}} \} \quad \mathbf{m} \notin \bar{\mathbf{M}} \\
\hline
mtype(\mathbf{m}, \mathbf{C}) = mtype(\mathbf{m}, \mathbf{D}) \\
\\
CT(\mathbf{C}) = \text{class } \mathbf{C} \text{ extends } \mathbf{D} \{ \bar{\mathbf{C}} \bar{\mathbf{f}}; \mathbf{K} \bar{\mathbf{M}} \} \\
\mathbf{B} \ \mathbf{m}(\bar{\mathbf{B}} \bar{\mathbf{x}}) \{ \text{return } \mathbf{e}; \} \in \bar{\mathbf{M}} \\
\hline
mbody(\mathbf{m}, \mathbf{C}) = \bar{\mathbf{x}}.\mathbf{e} \\
\\
CT(\mathbf{C}) = \text{class } \mathbf{C} \text{ extends } \mathbf{D} \{ \bar{\mathbf{C}} \bar{\mathbf{f}}; \mathbf{K} \bar{\mathbf{M}} \} \quad \mathbf{m} \notin \bar{\mathbf{M}} \\
\hline
mbody(\mathbf{m}, \mathbf{C}) = mbody(\mathbf{m}, \mathbf{D})
\end{array}$$

Figure 5. Auxiliary functions

transparent to the programmer, both for ease of use, and for backwards compatibility with Java. Thus, we reuse the declared, nominal type system of FJ, but add machinery to infer trace effects, for static verification of checks— although our type language will explicitly represent the field and method types of objects, rather than implicitly via lookup functions.

We define our type and effect analysis via subtyping constraints interpreted in a regular tree model. This representation promotes type reconstruction for common Object Oriented idioms such as object self-reference and binary methods (Eifrig et al., 1995; Bruce et al., 1995), since it possesses precisely the expressiveness of recursive types (Palsberg and Smith, 1996). A constraint type representation also yields an elegant definition of the *soft subtyping* relation for static analysis of casts, discussed in Sect. 2 and formalized below, which supports a simple form of soft typing (Cartwright and Fagan, 1991). Furthermore, as has been observed frequently in previous related type and effect analyses (Skalka and Smith, 2004; Igarashi and Kobayashi, 2002; K. Marriott and Sulzmann, 2003), some flavor of subeffecting is necessary to conservatively extend underlying type structure, with a subtyping approach

$H ::= \epsilon \mid \text{ev}[i] \mid \text{chk}[i] \mid \mathbf{h} \mid H;H \mid H H \mid \mu\mathbf{h}.H$	<i>trace effects</i>
$T ::= H \mid X \mid [TC] \mid (\bar{x} : \bar{T}) \mid \bar{T} \xrightarrow{H} T$	<i>types</i>
$X ::= \mathbf{h} \mid \mathbf{t}$	<i>type variables</i>
$C ::= T <: T \mid T \triangleleft T \mid C \wedge C \mid \mathbf{true}$	<i>constraints</i>

Figure 6. FJ_{trace} type and constraint syntax

being particularly flexible (Amtoft et al., 1999). A constraint representation is an effective implementation of subtyping, providing the expressiveness of intersection and union types (Eifrig et al., 1995). While a recursive constraint representation is not the most human-readable type abstraction, our goal here is a transparent program analysis, and automatic extraction of trace effects for verification.

As discussed in Sect. 2, we also incorporate *effect polymorphism* (Skalka and Smith, 2004), in a manner that allows flexibility and modularity of trace effect analysis in the presence of method override and dynamic dispatch.

5.1. TYPE AND CONSTRAINT LANGUAGE

The type and constraint grammar of FJ_{trace} is given in Fig. 6. The type language includes method types $\bar{T} \xrightarrow{H} T$ where H is the latent effect of the method. Object types are denoted $[TC]$, where C is the class name of the object and T is either a type variable \mathbf{t} or a vector of bindings $\bar{x} : \bar{T}$ for the field and method types of the object. The language of constraints is mostly standard, though in addition to subtyping constraints $T <: T$, we include weaker *soft subtyping* constraints $T \triangleleft T$, to address downcasting in the type analysis.

The type language specified in the grammar is more liberal than what is actually allowed in type judgements. Formally, we will define well-formedness of types via interpretation in a regular tree model, defined and discussed in Sect. 5.2. The model is endowed with a primitive subtyping relation, which also supplies meaning for constraints via the interpretation. We will require that constraints and types C, T assigned to expressions have a sensible (*well-kinded* and *realizable*, in our terminology) interpretation in the model. Informally, in any object type $[TC]$, if T is a vector of type bindings $\bar{x} : \bar{T}$, the set of bound names \bar{x} are restricted to be exactly the field and method names of C , providing an “inlined” width constraint on the form of the type.

$$\begin{array}{c}
\top^\emptyset, \perp^\emptyset : k \qquad \frac{\text{fv}(\mathbf{H}) = \emptyset}{\mathbf{H}^\emptyset : \text{Eff}} \\
\\
\frac{\begin{array}{l} \text{mtype}(\mathbf{m}, \mathbf{C}) = \bar{\mathbf{D}} \rightarrow \mathbf{D} \\ \varsigma = \text{Body}_{\bar{\mathbf{D}}}, \text{Eff}, \text{Body}_{\mathbf{D}} \quad \text{CT}(\mathbf{C}) = \text{class } \mathbf{C} \text{ extends } \mathbf{B} \{ \dots \} \\ \text{if } \text{mtype}(\mathbf{m}, \mathbf{B}) = \bar{\mathbf{E}} \rightarrow \mathbf{E} \text{ then } \bar{\mathbf{E}} = \bar{\mathbf{D}} \text{ and } \mathbf{E} = \mathbf{D} \end{array}}{[\cdot \bar{\mathbf{D}}] \dot{\mapsto} [\cdot \mathbf{D}]^\varsigma : \text{Meth}_{\mathbf{m}, \mathbf{C}}} \\
\\
\frac{\bar{\mathbf{x}} \text{ distinct} \quad |\varsigma| = |\bar{\mathbf{x}}| \quad k \in \varsigma \Rightarrow k \in \{ \text{Type}, \text{Meth}_{\mathbf{m}, \mathbf{C}} \}}{(\bar{\mathbf{x}} : \cdot)^\varsigma : \text{Type}} \\
\\
\frac{\text{fields}(\mathbf{C}) = \bar{\mathbf{D}} \bar{\mathbf{f}} \quad \varsigma = \text{Body}_{\bar{\mathbf{D}}}, \text{Meth}_{\bar{\mathbf{m}}, \mathbf{C}}}{(\bar{\mathbf{f}} : [\cdot \bar{\mathbf{D}}] \bar{\mathbf{m}} : \cdot)^\varsigma : \text{Body}_{\mathbf{C}}} \quad [\cdot \mathbf{C}]^{\text{Body}_{\mathbf{C}}} : \text{Type}
\end{array}$$

Figure 7. Regular tree ranked alphabet kinding rules

5.1.1. Vector Notations

For brevity we extend vector notation to the language of types and constraints. The type $\bar{\mathbf{T}}$ is a vector $\mathbf{T}_1, \dots, \mathbf{T}_n$, with \emptyset denoting the empty vector. We also write $[\bar{\mathbf{T}} \mathbf{C}]$ to denote a vector of class types $[\mathbf{T}_1 \mathbf{C}_1], \dots, [\mathbf{T}_n \mathbf{C}_n]$, while $\bar{\mathbf{x}} : \bar{\mathbf{T}}$ denotes a vector of bindings $\mathbf{x}_1 : \mathbf{T}_1 \cdots \mathbf{x}_n : \mathbf{T}_n$. We abbreviate constraints using vector notation, writing $\bar{\mathbf{S}} <: \bar{\mathbf{T}}$ for $\mathbf{S}_1 <: \mathbf{T}_1 \wedge \cdots \wedge \mathbf{S}_n <: \mathbf{T}_n$.

Looking ahead to the next section, vector notations are also used to abbreviate the kinding rules in Fig. 7 and Fig. 8. We write $\text{Body}_{\bar{\mathbf{C}}}$ for $\text{Body}_{\mathbf{C}_1}, \dots, \text{Body}_{\mathbf{C}_n}$, and $\text{Meth}_{\bar{\mathbf{m}}, \mathbf{C}}$ for $\text{Meth}_{\mathbf{m}_1, \mathbf{C}}, \dots, \text{Meth}_{\mathbf{m}_n, \mathbf{C}}$. We write $\bar{\varphi} \preceq_{\text{fin}} \bar{\varphi}'$ to denote $\varphi_1 \preceq_{\text{fin}} \varphi'_1, \dots, \varphi_n \preceq_{\text{fin}} \varphi'_n$.

5.2. REGULAR TREE MODEL AND SUBTYPING

To accommodate recursive constraints, we define subtyping in FJ_{trace} via primitive subtyping in a regular tree model, using techniques adapted from (Trifonov and Smith, 1996). In our model, function type nodes in regular trees are labeled with trace effects, and rather than being constructed from ranked alphabets as in (Trifonov and Smith, 1996) and elsewhere, our regular trees are constructed from kinded alphabets, imposing well-formedness of trees. Trace effect labelings of regular trees require an extension of the primitive subtyping relation to trace effects, which is based on set containment of effect interpretations. Otherwise, the interpretation is essentially standard.

$$\boxed{
\begin{array}{c}
\frac{\varphi \text{ is finite}}{\perp \preceq_{\text{fin}} \varphi} \qquad \frac{\varphi \text{ is finite}}{\varphi \preceq_{\text{fin}} \top} \qquad \frac{[[\mathbf{H}]] \subseteq [[\mathbf{H}']]}{\mathbf{H} \preceq_{\text{fin}} \mathbf{H}'} \\
\\
\frac{\bar{\varphi}_1 \preceq_{\text{fin}} \bar{\varphi}_1 \quad \varphi_2 \preceq_{\text{fin}} \varphi'_2 \quad \mathbf{H} \preceq \mathbf{H}' \quad \varphi_1, \varphi'_1, \varphi_2, \varphi'_2 \text{ finite}}{\bar{\varphi}_1 \xrightarrow{\mathbf{H}} \varphi_2 \preceq_{\text{fin}} \bar{\varphi}_1 \xrightarrow{\mathbf{H}} \varphi'_2} \\
\\
\frac{\bar{\varphi}_1 \preceq_{\text{fin}} \bar{\varphi}'_1 \quad \bar{\varphi}_1, \bar{\varphi}'_1 \text{ finite} \quad \varphi \preceq_{\text{fin}} \varphi' \quad \varphi, \varphi' \text{ finite} \quad \mathbf{C} <: \mathbf{D}}{\bar{\mathbf{x}} : \bar{\varphi}_1, \bar{\mathbf{y}} : \bar{\varphi}_2 \preceq_{\text{fin}} \bar{\mathbf{x}} : \bar{\varphi}'_1 \quad [\varphi \mathbf{C}] \preceq_{\text{fin}} [\varphi' \mathbf{D}]} \\
\\
\frac{\varphi|_n \preceq_{\text{fin}} \varphi'|_n \text{ for all } n \in \mathbb{N}}{\varphi \preceq \varphi'}
\end{array}
}$$

Figure 8. Primitive subtyping for regular trees

DEFINITION 5.1 (Regular Tree Model). *Let the tree constructor kinds be defined as:*

$$k ::= \text{Type} \mid \text{Eff} \mid \text{Meth}_{\mathbf{m}, \mathbf{c}} \mid \text{Body}_{\mathbf{c}}$$

and let signatures ς range over ordered sequences of kinds, where \emptyset denotes the empty sequence and $\varsigma(n)$ denotes the 0-indexed n th kind in ς . The alphabet L of tree constructors is built from the following grammar:

$$c ::= \top \mid \perp \mid \mathbf{H} \mid (\bar{\mathbf{x}} : \cdot) \mid [\cdot \mathbf{C}] \mid [\cdot \bar{\mathbf{C}}] \dot{\rightarrow} [\cdot \mathbf{C}]$$

where each element of the alphabet is indexed by a signature, written c^ς , and must be well-kinded according to the rules given in Fig. 7.

A tree φ is a partial function from finite sequences (paths) π of natural numbers \mathbb{N}^* to L such that $\text{dom}(\varphi)$ is prefix-closed. Furthermore, for all $\pi n \in \text{dom}(\varphi)$, with $c^\varsigma = \varphi(\pi)$, it is the case that $\varphi(\pi n) : \varsigma(n)$. Trees of the form $(\bar{\mathbf{x}} : \bar{\varphi})$ are equated up to reordering of labels. The subtree at $\pi \in \text{dom}(\varphi)$ is the function $\lambda \pi'. \varphi(\pi \pi')$, while $|\pi|$ is the level of that subtree. A tree is regular iff the set of its subtrees is finite, and we define \mathbb{T} as the set of regular trees over L .

A partial order over \mathbb{T} is then defined via an approximate relation over finite $\varphi \in \mathbb{T}$. First, define a level- n cut $\varphi|_n$ for $\varphi \in \mathbb{T}$ as the finite tree obtained by replacing all subtrees at level n of φ with \top . Then, \preceq_{fin} is the partial order over finite $\varphi \in \mathbb{T}$ axiomatized in Fig. 8, and \preceq is the partial order over \mathbb{T} approximated by \preceq_{fin} axiomatized in Fig. 8.

The meaning of subtyping constraints is then defined via interpretation in the regular tree model. The principal novelties here are the

$$\begin{array}{ll}
\rho(\mathbf{h}, hs) = \rho(\mathbf{h}) & \mathbf{h} \notin hs \\
\rho(\mathbf{h}, hs) = \mathbf{h} & \mathbf{h} \in hs \\
\rho(\mathbf{ev}[i], hs) = \mathbf{ev}[i] & \\
\rho(\mathbf{chk}[i], hs) = \mathbf{chk}[i] & \\
\rho(\epsilon, hs) = \epsilon & \\
\rho(\mathbf{H}_1; \mathbf{H}_2, hs) = \rho(\mathbf{H}_1, hs); \rho(\mathbf{H}_2, hs) & \\
\rho(\mathbf{H}_1 | \mathbf{H}_2, hs) = \rho(\mathbf{H}_1, hs) | \rho(\mathbf{H}_2, hs) & \\
\rho(\mu \mathbf{h}. \mathbf{H}, hs) = \mu \mathbf{h}. \rho(\mathbf{H}, hs \cup \{h\}) & \\
\\
\rho([\mathbf{T} \mathbf{C}]) = [\rho(\mathbf{T}) \mathbf{C}] & \\
\rho(\mathbf{x} : \mathbf{T}) = \mathbf{x} : \rho(\mathbf{T}) & \\
\rho(\bar{\mathbf{T}} \xrightarrow{\mathbf{H}} \mathbf{T}) = \rho(\bar{\mathbf{T}}) \xrightarrow{\rho(\mathbf{H})} \rho(\mathbf{T}) & \\
\rho(\bar{\mathbf{T}}, \mathbf{T}) = \rho(\bar{\mathbf{T}}), \rho(\mathbf{T}) & \\
\rho(\emptyset) = \emptyset & \\
\rho(\mathbf{H}) = \rho(\mathbf{H}, \emptyset) &
\end{array}$$

Figure 9. Interpretations extended to types and effects

extension of subtyping to trace effects, and the interpretation of the soft subtyping relation.

DEFINITION 5.2 (Interpretation of Constraints). Interpretations ρ are total mappings from type variables \mathbf{X} to \mathbb{T} . Interpretations are extended to types in Fig. 9, and also to effects by an abuse of notation allowing parameterization by sets hs of effect variables, to prevent substitution of μ -bound variables. The relation $\rho \vdash C$, pronounced ρ satisfies or solves C , is axiomatized as follows:

$$\begin{array}{c}
\rho \vdash \mathbf{true} \quad \frac{\rho(\mathbf{S}) \preceq \rho(\mathbf{T})}{\rho \vdash \mathbf{S} <: \mathbf{T}} \quad \frac{\rho \vdash C \quad \rho \vdash D}{\rho \vdash C \wedge D} \\
\\
\frac{\forall \mathbf{BR}. (\mathbf{B} <: \mathbf{D} \wedge \rho \vdash [\mathbf{R}\mathbf{B}] <: \mathbf{S}) \Rightarrow \rho \vdash [\mathbf{R}\mathbf{B}] <: [\mathbf{T}\mathbf{D}]}{\rho \vdash \mathbf{S} <: [\mathbf{T}\mathbf{D}]}
\end{array}$$

The relation $C \Vdash D$ holds iff $\rho \vdash C$ implies $\rho \vdash D$ for all interpretations ρ . Constraints C and D are equivalent, written $C = D$, iff $C \Vdash D$ and $D \Vdash C$.

We immediately observe that transitivity of primitive subtyping is reflected in the type and constraint representation.

LEMMA 5.1 (Transitivity of Entailment and Subtyping). $C \Vdash D$ and $D \Vdash E$ implies $C \Vdash E$, and $C \Vdash R <: S \wedge S <: T$ implies $C \Vdash R <: T$.

We also note the relevance of soft subtyping for the type analysis, that a soft subtyping relation entails an analogous subtyping relation just in case the latter could be sound.

LEMMA 5.2. If $C \Vdash [TC] < [SD]$, then $C \Vdash [TC] < [SD]$ iff $C < D$.

5.2.1. Discussion

The kinding rules in the interpretation of types impose certain well-formedness qualities on types, in particular the class name component C of a type $[TC]$ fixes any fields and methods appearing in the type to be exactly those that are listed in the class definition. This is because the underlying FJ type system allows any object to be treated as an object in its statically known class— thus, types for all known fields and methods for that class must be available, and no others. For example, imagining a class `Foo` with two fields `a` and `b` of `Object` type:

```
class Foo extends Object { Object a; Object b; }
```

The type of `Foo` objects are labeled with the class name, and types for each of the fields:

```
[(a : [Object]; b : [Object]) Foo]
```

While the type system will allow any `Foo` object to be viewed as an `Object` by subsumption, any object type term labeled `Foo` with anything other than both `a` and `b` fields has no interpretation due to the kinding rules. For example, the following type term has no interpretation:

```
[(a : [Object]) Foo]
```

The kinding rule for method typings also deserves attention for its contribution to the independence of inheritance and effects in the analysis. Recalling that the type system of (Igarashi et al., 2001) requires that method overrides in a subclass have the same type signature as the overridden methods in their superclass, we observe that the kinding rule for constructors of kind $Meth_{m,C}$ imposes this restriction, but only on the declared class name component of the type. The inferred effect component H , on the other hand, is not restricted to relate to superclass method effects in any way. This issue is revisited with examples in Sect. 5.3.1.

The interpretation of type constraints defines a system of object width and depth subtyping, with method subtyping predicated on subsumption of trace effects. This extension, defined in Fig. 8, is based on

containment of trace effect interpretations, reflecting a type soundness requirement that if \mathbf{T} subsumes \mathbf{S} , it must also subsume \mathbf{S} 's trace effects. Since constraints are defined on the basis of interpretations, constraints on types with abstract components are meaningful; for example, we could assert:

$$\mathbf{S} <: \mathbf{T} \wedge \bar{\mathbf{T}} <: \bar{\mathbf{S}} \Vdash \bar{\mathbf{S}} \xrightarrow{\mathbf{h}} \mathbf{S} <: \bar{\mathbf{T}} \xrightarrow{\mathbf{h}(\mathbf{ev}[1];\mathbf{ev}[2])} \mathbf{T}$$

since any interpretation of \mathbf{h} must be contained in the same interpretation $\mathbf{h}(\mathbf{ev}[1];\mathbf{ev}[2])$.

The soft subtyping relation is useful in application to downcasts, allowing constraints to be relaxed for downcasting. For example, assuming $\mathbf{C} <: \mathbf{D}$, we could meaningfully assert:

$$[\mathbf{RC}] <: \mathbf{T} \wedge [\mathbf{SD}] <: \mathbf{T} \wedge \mathbf{T} \triangleleft [\mathbf{UC}] \Vdash [\mathbf{RC}] <: [\mathbf{UC}]$$

but not:

$$[\mathbf{RC}] <: \mathbf{T} \wedge [\mathbf{SD}] <: \mathbf{T} \wedge \mathbf{T} \triangleleft [\mathbf{UC}] \Vdash [\mathbf{SD}] <: [\mathbf{UC}]$$

which formalizes the idea discussed in Sect. 2. This relaxation is allowable at casting points, since dynamic checks will “pick up the slack”, preserving type safety as discussed in the next section.

This definition of subtyping may raise questions about the decidability of typing, since subtyping is predicated on trace effect containment, but equivalence of trace effects was shown to be undecidable in Sect. 4. However, while trace effect containment is undecidable in the general case, we show in Sect. 6 that constraints generated by type inference are in a normal form that is amenable to algorithmic solution.

5.3. LOGICAL TYPE JUDGEMENTS AND PROPERTIES

To define the logical type system, we introduce constrained type schemes $\forall \bar{\mathbf{X}}[C].\mathbf{T}$ and type environments Γ binding class names and variables to type schemes. If $\bar{\mathbf{X}} \cap \text{fv}(\mathbf{T}) = \emptyset$ we abbreviate $\forall \bar{\mathbf{X}}[\mathbf{true}].\mathbf{T}$ as \mathbf{T} . We also introduce three forms of type judgements: $\Gamma, C, \mathbf{H} \vdash \mathbf{e} : \mathbf{T}$ for expressions, $\Gamma, C \vdash \mathbf{m}, \mathbf{C} : \bar{\mathbf{T}} \xrightarrow{\mathbf{H}} \mathbf{T}$ for method \mathbf{m} in class \mathbf{C} , and $\Gamma \vdash \mathbf{C} : \forall \bar{\mathbf{X}}[C].\mathbf{T}$ for classes. Sensibility of types and constraints in judgements is imposed by realizability conditions, as follows.

DEFINITION 5.3 (Realizability of Types). *We say that a constraint, type pair C, \mathbf{T} or type scheme $\forall \bar{\mathbf{X}}[C].\mathbf{T}$ is realizable iff there exists a solution ρ of C such that $\rho(\mathbf{T})$ is well-kinded, and we impose the following sanity conditions on judgements: for any $\Gamma, C, \mathbf{H} \vdash \mathbf{e} : \mathbf{T}$ or $\Gamma, C \vdash \mathbf{m}, \mathbf{C} : \bar{\mathbf{T}} \xrightarrow{\mathbf{H}} \mathbf{T}$ we require that C, \mathbf{T} is realizable, and every type scheme in Γ must be realizable, in which case we say that Γ is realizable.*

T-VAR $\Gamma, C, \epsilon \vdash \mathbf{x} : \Gamma(\mathbf{x})$	
T-FIELD $\frac{\Gamma, C, \mathbf{H} \vdash \mathbf{e} : [\mathbf{S}\mathbf{C}]}{\Gamma, C, \mathbf{H} \vdash \mathbf{e}.f : [\mathbf{T}\mathbf{D}]}$	$C \Vdash \mathbf{S} <: (f : [\mathbf{T}\mathbf{D}]) \quad \mathbf{D} f \in \text{fields}(C)$
T-SEQNIL $\frac{}{\Gamma, C, \epsilon \vdash \emptyset : \emptyset}$	T-SEQCONS $\frac{\Gamma, C, \mathbf{H}_1 \vdash \bar{\mathbf{e}} : \bar{\mathbf{T}} \quad \Gamma, C, \mathbf{H}_2 \vdash \mathbf{e} : \mathbf{T}}{\Gamma, C, \mathbf{H}_1; \mathbf{H}_2 \vdash \bar{\mathbf{e}}, \mathbf{e} : \bar{\mathbf{T}}, \mathbf{T}}$
T-INVK $\frac{\Gamma, C, \mathbf{H}_1 \vdash \mathbf{e} : [\mathbf{T}\mathbf{C}] \quad C \Vdash \mathbf{T} <: (\mathbf{m} : \bar{\mathbf{S}} \xrightarrow{\mathbf{H}_3} [\mathbf{S}\mathbf{D}]) \quad \text{mtype}(\mathbf{m}, C) = \bar{\mathbf{D}} \rightarrow \mathbf{D}}{\Gamma, C, \mathbf{H}_2 \vdash \bar{\mathbf{e}} : \bar{\mathbf{S}} \quad \Gamma, C, \mathbf{H}_1; \mathbf{H}_2; \mathbf{H}_3 \vdash \mathbf{e}.m(\bar{\mathbf{e}}) : [\mathbf{S}\mathbf{D}]}$	
T-NEW $\frac{\Gamma(C) = \forall \bar{\mathbf{X}}[D].[\mathbf{T}\mathbf{C}] \quad C \Vdash D[\bar{\mathbf{S}}/\bar{\mathbf{X}}] \quad \text{fieldsig}[\mathbf{T}\mathbf{C}] = \bar{\mathbf{T}} \quad \Gamma, C, \mathbf{H} \vdash \bar{\mathbf{e}} : \bar{\mathbf{R}} \quad C \Vdash \bar{\mathbf{R}} <: \bar{\mathbf{T}}[\bar{\mathbf{S}}/\bar{\mathbf{X}}]}{\Gamma, C, \mathbf{H} \vdash \text{new } C(\bar{\mathbf{e}}) : [\mathbf{T}[\bar{\mathbf{S}}/\bar{\mathbf{X}}] C]}$	
T-EVENT $\Gamma, C, \text{ev}[i] \vdash \text{ev}[i] : \text{Unit}$	T-CHECK $\Gamma, C, \text{chk}[i] \vdash \text{chk}[i] : \text{Unit}$
T-CAST $\frac{\Gamma, C, \mathbf{H} \vdash \mathbf{e} : \mathbf{T} \quad C \Vdash \mathbf{T} <: [\mathbf{S}\mathbf{D}]}{\Gamma, C, \mathbf{H} \vdash (\mathbf{D})\mathbf{e} : [\mathbf{S}\mathbf{D}]}$	T-WEAKEN $\frac{\Gamma, C, \mathbf{H}' \vdash \mathbf{e} : \mathbf{T} \quad C \Vdash \mathbf{H}' <: \mathbf{H}}{\Gamma, C, \mathbf{H} \vdash \mathbf{e} : \mathbf{T}}$
T-METH $\frac{C \Vdash \mathbf{S} <: \mathbf{T} \wedge \bar{\mathbf{T}} <: \bar{\mathbf{S}} \quad \Gamma; \bar{\mathbf{x}} : \bar{\mathbf{S}}, C, \mathbf{H} \vdash \mathbf{e} : \mathbf{S} \quad \Gamma(\text{this}).m = \bar{\mathbf{T}} \xrightarrow{\mathbf{H}} \mathbf{T} \quad \text{mbody}(\mathbf{m}, C) = \bar{\mathbf{x}}.e}{\Gamma, C \vdash \mathbf{m}, C : \bar{\mathbf{T}} \xrightarrow{\mathbf{H}} \mathbf{T}}$	
T-CLASS $\frac{\Gamma; \mathbf{C} : [\mathbf{T}\mathbf{C}]; \text{this} : [\mathbf{T}\mathbf{C}], C \vdash m_i, C : \mathbf{T}_i \text{ for all } m_i \in \text{meths}(C)}{\Gamma \vdash C : \forall \bar{\mathbf{X}}[C].[\mathbf{T}\mathbf{C}]}$	

Figure 10. FJ_{trace} logical typing rules

$$\begin{array}{c}
\frac{\{\mathbf{m}_1, \dots, \mathbf{m}_n\} = \{\mathbf{m} \mid \mathit{mbody}(\mathbf{m}, \mathbf{C}) \text{ is defined}\}}{\mathit{meths}(\mathbf{C}) = \mathbf{m}_1, \dots, \mathbf{m}_n} \\
\\
\frac{\mathit{fields}(\mathbf{C}) = \mathbf{D}_1 \mathbf{f}_1, \dots, \mathbf{D}_n \mathbf{f}_n \quad \mathbf{f}_1 : \mathbf{T}_1 \in \mathbf{T}, \dots, \mathbf{f}_n : \mathbf{T}_n \in \mathbf{T}}{\mathit{fieldsig}[\mathbf{T} \mathbf{C}] = \mathbf{T}_1, \dots, \mathbf{T}_n}
\end{array}$$

Figure 11. Auxiliary functions

Type derivation rules are given in Fig. 10, with some auxiliary functions defined in Fig. 11. For these rules and later examples, we posit a **Unit** type, which is the type of objects in the **Object** class, that possess no fields or methods; i.e.:

$$\mathbf{Unit} \triangleq [\mathbf{Object}]$$

In these and later rules we also write $[\bar{\mathbf{x}} : \bar{\mathbf{T}} \mathbf{C}].\mathbf{x}_i$ to denote \mathbf{T}_i . We note that \forall -intro and -elim are located at class definition and object construction points, respectively. Subeffecting of top level effects is allowable at any point in a derivation via **T-WEAKEN**. The soft subtyping relation is used for casts $(\mathbf{C})\mathbf{e}$; due to the interpretation of soft subtyping defined in the previous section, this rule will only track the flow of objects to the casting point that are in a subclass of \mathbf{C} . Any other objects will cause a dynamic cast check exception, and can therefore be ignored statically without compromising type safety, as in (Igarashi et al., 2001). Ignoring “junk” in this manner gives a more precise analysis, and allows typing of downcasts. We omit the distinction of “stupid casts” entailing “stupid warnings” as in (Igarashi et al., 2001), noting that we essentially follow their approach, and can easily adopt this distinction. We make the following definitions:

DEFINITION 5.4. *The judgement $\Gamma, \mathbf{C}, \mathbf{H} \vdash \mathbf{e} : \mathbf{T}$ is valid iff it is derivable and there exists a solution ρ of \mathbf{C} such that $\rho(\mathbf{H})$ is valid. An environment Γ is well-formed iff $\Gamma \vdash \mathbf{C} : \forall \mathbf{X}[C].\mathbf{T}$ is derivable for all $(\mathbf{C} : \forall \mathbf{X}[C].\mathbf{T}) \in \Gamma$.*

It is demonstrable that for closed, event- and check-free expressions \mathbf{e} , there exists a derivable typing for \mathbf{e} in $\mathbf{FJ}_{\text{trace}}$ iff \mathbf{e} is well-typed in the type system of (Igarashi et al., 2001), as we show in Sect. 5.4, Lemma 5.8. This means that our system absorbs properties of **FJ**, including type safety for the **FJ** subset of $\mathbf{FJ}_{\text{trace}}$, implying that the only case we really need to consider, to extend type safety to $\mathbf{FJ}_{\text{trace}}$, is progress for trace checks. That is, we must show that checks can be statically enforced by our type analysis. We establish this by a subject

reduction argument, showing that trace effect approximations are preserved by computation. Note that the statement of subject reduction must apply to configurations; the proof, given in Sect. 5.4, follows by induction on derivations:

LEMMA 5.3 (Subject Reduction). *If $\Gamma, C, H \vdash e : [TC]$ is derivable for closed e and well formed Γ , and $\eta, e \rightarrow \eta', e'$, then $\Gamma, C, H' \vdash e' : [T' C']$ is derivable with $C \Vdash [T' C'] <: [TC]$ and $C \Vdash \eta'; H' <: \eta; H$.*

A corollary of this result formalizes the intuition that trace effects approximate event traces, insofar as any event trace generated during evaluation of an expression must be contained in the trace effect assigned to that expression by the type analysis:

COROLLARY 5.1. *If $\Gamma, C, H \vdash e : [TC]$ is derivable for closed e and well-formed Γ , and $\epsilon, e \rightarrow \eta, e'$, then $C \Vdash \eta <: H$.*

To state our type safety result, we must define what we mean by “run-time checks”. In short, these are checks in the hole of an evaluation context encountered during execution:

DEFINITION 5.5. Evaluation contexts are defined as follows:

$$E ::= [] \mid E.f \mid E.m(\bar{e}) \mid v.m(\bar{v}, E, \bar{e}) \mid \mathbf{new} C(\bar{v}, E, \bar{e}) \mid (C)E$$

We can then formally assert that run-time checks in well-typed programs are guaranteed to succeed:

LEMMA 5.4 (Static Enforcement of Checks). *Given closed e and well-formed Γ , if the judgement $\Gamma, C, H \vdash e : T$ is valid, and $\epsilon, e \rightarrow^* \eta, E[\mathbf{chk}[x]]$, then $\eta \vdash \mathbf{chk}[x]$.*

Proofs are demonstrated in Sect. 5.4.

5.3.1. Discussion

In this section we discuss some examples that illustrate properties of the type system. We assume the definitions and types given in Sect. 2. We also assume the trivial extension of the language with a sequencing construct $e; e$ and a lexically scoped name-to-value binding construct $C x = \mathbf{new} C()$. To provide intuitions more easily, we present types and effects in a unified form. We posit the definition of a class `PasswdTmp1`, which is owned by `System`; in class `PasswdTmp1` the `format` method is overridden, defined to format a password file template string:

```
class PasswdTmp1 extends Formatter {
  String format() { System; ... }
}
```

Assuming that the body of this version of `format` is effect free other than the initial `System` event, the type of `PasswdTpl` can be given as follows, where `StringT` is the type of a `String` object, the details of which are unimportant to the example:

$$\text{PTmplT} \triangleq [\text{format} : () \xrightarrow{\text{System}} \text{StringT} \text{ PasswdTpl}]$$

We also posit the definition of a class `Backdoor`, devised as an untrusted `Applet` extending `PasswdTpl`, that formats a password file template containing a `uname/passwd` combination known to the attacker:

```
class Backdoor extends PasswdTpl {
  String format() { Applet;...}
}
```

Assuming that the body of this version of `format` is also effect free other than the initial `Applet` event, the type of `Backdoor` can be given as follows, which we abbreviate as `BdoorT`:

$$\text{BdoorT} \triangleq [\text{format} : () \xrightarrow{\text{Applet}} \text{StringT} \text{ Backdoor}]$$

Notice that the effects assigned to the `Backdoor` and `PasswdTpl` versions of `format` are incomparable, despite their inheritance relation.

For the purposes of the example, we further imagine that there exists an open `PasswdFile` object in the current namespace. Now, since `Writer.safewrite` can be assigned a type which is polymorphic in the effects of its arguments, the code:

```
Writer w = new Writer();
PasswdTpl p = new PasswdTpl();
w.safewrite(p, PasswdFile);
```

could be assigned the effect:

```
System; System; demand(FileWrite);
```

which we assume is verifiable, whereas the following application of `safewrite` could be treated independently:

```
Writer w = new Writer();
Backdoor b = new Backdoor();
w.safewrite(b, PasswdFile);
```

and assigned the following effect, which we assume is not verifiable:

```
System; Applet; demand(FileWrite)
```

Since polymorphism is restricted to be first order, method parameters themselves cannot be polymorphic. This means that in code such as the following, the domain effects of functionally abstracted objects will be merged by subtyping:

```
class C extends Object {
  void m(Writer w, File f){
    Backdoor b = new Backdoor();
    PasswdTpl p = new PasswdTpl();
    w.safewrite(b, f);
    w.safewrite(p, f);
  }
}
```

This implies that the effects at each calling site of `w.safewrite` generated by the following code cannot be distinguished:

```
Writer w = new Writer();
C c = new C();
c.m(w, PasswdFile);
```

so that the following effect would be assigned:

```
System; )(System|Applet); demand(FileWrite);
System; (System|Applet); demand(FileWrite)
```

Soft subtyping can be used at downcasts to ignore unsound flow that will be caught by dynamic cast checks. Suppose that type analysis predicts that some expression `e` may evaluate to either a `PasswdTpl` object or a `Backdoor` object; i.e. $\Gamma, C, H \vdash e : T$ with:

$$C \Vdash \text{PTmplT} <: T \wedge \text{BdoorT} <: T$$

where `PTmplT` and `BdoorT` are defined as above. By the `T-CAST` typing rule and properties of soft subtyping discussed above, we therefore may assert (as a somewhat contrived example; a `Backdoor` cast is not likely to occur in practice):

$$\Gamma, C, \epsilon \vdash (\text{Backdoor})e : \text{BdoorT}$$

meaning that the following:

```
Writer w = new Writer();
w.safewrite((Backdoor)e, f);
```

may be assigned the following effect:

```
H; System; Applet; demand(FileWrite)
```

5.4. PROPERTIES

In this section we demonstrate our main type safety result, based on a subject reduction argument. We also show that FJ_{trace} typings conserve FJ typings, i.e. any program that is typable in the FJ type system is also typable in the FJ_{trace} type system. In the course of proving these results, we demonstrate a standard suite of results (weakening, substitution, canonical forms, etc.) for the FJ_{trace} type system.

We begin by demonstrating conservation of FJ typing in FJ_{trace} . FJ typing is nominal, and object types are represented “implicitly” via the *mtype* and *fields* functions. In FJ_{trace} , types are represented “explicitly” via type terms and constraints. Thus, we define a translation, or “expansion”, from types of the former to the latter. The translation expands every class name \mathbf{C} into a representation $[\mathbf{T}\mathbf{C}]$, where \mathbf{T} is a listing of field and method types. Each method is assigned an ϵ effect, since FJ is effect free. Expansion of any class name \mathbf{C} is only one level deep— that is, class names \mathbf{D} in \mathbf{C} ’s field and method types are not recursively expanded, but assigned a special type variable that is constrained to be the lower bounds of the expansion of \mathbf{D} . This ensures termination of expansion in the presence of recursive types. Since constraints can be recursive, a corollary of this definition is the representability of recursive class types in the expansion— i.e., the constraint is solvable.

DEFINITION 5.6. *Assume given a fixed type variable $\mathfrak{t}_{\mathbf{C}}$ for each $\mathbf{C} \in \text{dom}(CT)$. Then the expansion of method types, class names, and vectors of class names is defined as follows:*

$$\frac{\text{mtype}(\mathfrak{m}, \mathbf{C}) = \bar{\mathbf{D}} \rightarrow \mathbf{D} \quad \forall \mathbf{S}_i \in \bar{\mathbf{S}}. \mathbf{S}_i = \mathfrak{t}_{\mathbf{D}_i}}{\text{expand}(\mathfrak{m}, \mathbf{C}) = [\bar{\mathbf{S}}\bar{\mathbf{D}}] \xrightarrow{\epsilon} [\mathfrak{t}_{\mathbf{D}}\mathbf{D}]}$$

$$\frac{\begin{array}{l} \text{fields}(\mathbf{C}) = \bar{\mathbf{D}} \bar{\mathbf{f}} \quad \forall \mathbf{S}_i \in \bar{\mathbf{S}}. \mathbf{S}_i = \mathfrak{t}_{\mathbf{D}_i} \\ \text{meths}(\mathbf{C}) = \bar{\mathbf{m}} \quad \forall \mathfrak{m}_i \in \bar{\mathbf{m}}. \text{expand}(\mathfrak{m}_i, \mathbf{C}) = \mathbf{T}_i \end{array}}{\text{expand}(\mathbf{C}) = [\bar{\mathbf{f}} : [\bar{\mathbf{S}}\bar{\mathbf{D}}] \bar{\mathbf{m}} : \bar{\mathbf{T}} \mathbf{C}]}$$

$$\text{expand}(\bar{\mathbf{C}}) = \text{expand}(\mathbf{C}_1), \dots, \text{expand}(\mathbf{C}_n)$$

We also construct a possibly recursive system of lower bounds on type variables $\mathfrak{t}_{\mathbf{C}}$, as follows. Given CT such that $\text{dom}(CT) = \mathbf{C}_1, \dots, \mathbf{C}_n$. Then:

$$C_{CT} \triangleq \text{expand}(\mathbf{C}_1) <: [\mathfrak{t}_{\mathbf{C}_1} \mathbf{C}_1] \wedge \dots \wedge \text{expand}(\mathbf{C}_n) <: [\mathfrak{t}_{\mathbf{C}_n} \mathbf{C}_n]$$

LEMMA 5.5. $C_{CT}, \text{expand}(\mathbf{C})$ is realizable for all $\mathbf{C} \in \text{dom}(CT)$.

We observe that expansion preserves subtyping, as desired:

LEMMA 5.6. *If $\mathbf{C} <: \mathbf{D}$ then $C_{CT} \Vdash \text{expand}(\mathbf{C}) <: \text{expand}(\mathbf{D})$.*

Now, we show that the expansion allows a simulation of FJ field and method typing.

LEMMA 5.7. *Given $\text{expand}(\mathbf{D}) = [\mathbf{T}\mathbf{D}]$. Then the following properties hold:*

1. $C_{CT} \Vdash \mathbf{T} <: (\mathbf{m} : \text{expand}(\bar{\mathbf{B}}) \xrightarrow{\epsilon} \text{expand}(\mathbf{B}))$ if $\text{mtype}(\mathbf{m}, \mathbf{D}) = \bar{\mathbf{B}} \rightarrow \mathbf{B}$.
2. $C_{CT} \Vdash \mathbf{T} <: (\mathbf{f} : \text{expand}(\mathbf{B}))$ if $\mathbf{B} \mathbf{f} \in \text{fields}(\mathbf{D})$.

Now we can prove the conservativity result by induction on FJ type derivations. To properly frame the result, we need to extend expansion to environments, both to correctly state the induction, and to gather class type representations into the FJ_{trace} type environment in the appropriate manner.

LEMMA 5.8 (Conservation of FJ Typing). *Define:*

$$\frac{\text{expand}(\Gamma) = \Gamma' \quad \text{expand}(\mathbf{C}) = \mathbf{T}}{\text{expand}(\Gamma; x : \mathbf{C}) = \Gamma'; x : \mathbf{T}}$$

$$\frac{\text{dom}(CT) = \mathbf{C}_1, \dots, \mathbf{C}_n}{\text{expand}(\emptyset) = \forall \emptyset [C_{CT}]. \text{expand}(\mathbf{C}_1); \dots; \forall \emptyset [C_{CT}]. \text{expand}(\mathbf{C}_n)}$$

If $\Gamma \vdash \mathbf{e} : \mathbf{C}$ is derivable, then so is $\text{expand}(\Gamma), C_{CT}, \epsilon \vdash \mathbf{e} : \text{expand}(\mathbf{C})$.

Proof. The result follows by induction on $\Gamma \vdash \mathbf{e} : \mathbf{C}$ and case analysis on the last step in the derivation. Here, we give only the T-INVK case, wherein $\mathbf{e} = \mathbf{e}_0.\mathbf{m}(\bar{\mathbf{e}})$, and by inversion of the rule we have:

$$\Gamma \vdash \mathbf{e}_0 : \mathbf{C}_0 \quad \text{mtype}(\mathbf{m}, \mathbf{C}_0) = \bar{\mathbf{D}} \rightarrow \mathbf{C} \quad \Gamma \vdash \bar{\mathbf{e}} : \bar{\mathbf{C}} \quad \bar{\mathbf{C}} <: \bar{\mathbf{D}}$$

Hence by the induction hypothesis the judgements:

$$\text{expand}(\Gamma), C_{CT}, \epsilon \vdash \mathbf{e}_0 : \text{expand}(\mathbf{C}_0)$$

$$\text{expand}(\Gamma), C_{CT}, \epsilon \vdash \bar{\mathbf{e}} : \text{expand}(\bar{\mathbf{C}})$$

are derivable. Let $\text{expand}(\mathbf{C}_0) = [\mathbf{T}\mathbf{C}_0]$. Then by Lemma 5.6 and Lemma 5.7, the following are valid:

$$C_{CT} \Vdash \mathbf{T} <: (\mathbf{m} : \text{expand}(\bar{\mathbf{D}}) \xrightarrow{\epsilon} \text{expand}(\mathbf{C}))$$

$$C_{CT} \Vdash \text{expand}(\bar{\mathbf{C}}) <: \text{expand}(\bar{\mathbf{D}})$$

But then by Definition 5.2 and Lemma 5.1:

$$C_{CT} \Vdash \mathbf{T} <: (\mathbf{m} : \mathit{expand}(\bar{\mathbf{C}}) \xrightarrow{\epsilon} \mathit{expand}(\mathbf{C}))$$

so the judgement:

$$\mathit{expand}(\Gamma), C_{CT}, \epsilon; \epsilon \vdash \mathbf{e} : \mathit{expand}(\mathbf{C})$$

is derivable by T-INVK in FJ_{trace} , hence this case follows by T-WEAKEN. \square

Next, we turn to subject reduction and type soundness. We start by showing that valid typing judgements can always be specialized, either by weakening the top-level constraint in the judgement, instantiating type variables, or by adding extraneous bindings to the type environment. Each of these results follows by a straightforward induction on type derivations. The instantiation Lemma also follows because instantiation preserves subtyping, as we show in Lemma 5.10, a result that follows immediately by Definition 5.2.

LEMMA 5.9 (Constraint Weakening). *If $\Gamma, C, \mathbf{H} \vdash \mathbf{e} : \mathbf{T}$ and $D \Vdash C$ for solvable D then $\Gamma, D, \mathbf{H} \vdash \mathbf{e} : \mathbf{T}$.*

LEMMA 5.10. *If $C \Vdash \mathbf{S} <: \mathbf{T}$ then $C[\bar{\mathbf{T}}/\bar{\mathbf{X}}] \Vdash \mathbf{S}[\bar{\mathbf{T}}/\bar{\mathbf{X}}] <: \mathbf{T}[\bar{\mathbf{T}}/\bar{\mathbf{X}}]$.*

LEMMA 5.11 (Instantiation). *If $\Gamma, C, \mathbf{H} \vdash \mathbf{e} : \mathbf{T}$ is derivable, then so is $\Gamma[\bar{\mathbf{T}}/\bar{\mathbf{X}}], C[\bar{\mathbf{T}}/\bar{\mathbf{X}}], \mathbf{H}[\bar{\mathbf{T}}/\bar{\mathbf{X}}] \vdash \mathbf{e} : \mathbf{T}[\bar{\mathbf{T}}/\bar{\mathbf{X}}]$.*

LEMMA 5.12 (Environment Weakening). *If $\Gamma, C, \mathbf{H} \vdash \mathbf{e} : \mathbf{T}$ is derivable and $\Gamma(\mathbf{C}) = \Gamma'(\mathbf{C})$ for all $\mathbf{C} \in \text{dom}(CT)$ and $\Gamma(\mathbf{x}) = \Gamma'(\mathbf{x})$ for all $\mathbf{x} \in \text{fv}(\mathbf{e})$, then $\Gamma', C, \mathbf{H} \vdash \mathbf{e} : \mathbf{T}$ is derivable.*

Since the T-WEAKEN rule is not syntax directed, it introduces non-determinism in derivations. The following result allows us to treat specified normal forms of derivations in proofs, without loss of generality.

LEMMA 5.13 (Normalization). *If $\Gamma, C, \mathbf{H} \vdash \mathbf{e} : \mathbf{T}$, then it follows by one instance of T-WEAKEN preceded by an instance of a syntax-directed rule corresponding to the form of \mathbf{e} .*

We also show that substitution preserves typings, to underpin the method invocation case of subject reduction. After that, we give a canonical forms Lemma, specifying the form of values inhabiting particular types. We also observe that values can always be assigned top-level ϵ effects—since values do not reduce, they can have no trace history.

LEMMA 5.14 (Substitution). *If $\Gamma; \mathbf{x} : \mathbf{S}, C, \mathbf{H} \vdash \mathbf{e} : \mathbf{T}$ and $\Gamma, C, \epsilon \vdash \mathbf{v} : \mathbf{S}'$ with $C \Vdash \mathbf{S}' <: \mathbf{S}$, then $\Gamma, C, \mathbf{H}' \vdash [\mathbf{v}/\mathbf{x}]\mathbf{e} : \mathbf{T}'$ for some \mathbf{T}' and \mathbf{H}' with $C \Vdash \mathbf{T}' <: \mathbf{T} \wedge \mathbf{H}' <: \mathbf{H}$.*

LEMMA 5.15 (Canonical Forms). *If $\Gamma, C, \mathbf{H} \vdash \mathbf{v} : [\mathbf{TC}]$ is derivable for closed \mathbf{v} and well formed Γ , then $\mathbf{v} = \mathbf{new} \ C(\bar{\mathbf{v}})$ for some $\bar{\mathbf{v}}$.*

LEMMA 5.16. *If $\Gamma, C, \mathbf{H} \vdash \mathbf{v} : [\mathbf{TC}]$ is derivable, then $C \Vdash \epsilon <: \mathbf{H}$ and $\Gamma, C, \epsilon \vdash \mathbf{v} : [\mathbf{TC}]$ is also derivable.*

The next two utility results allow the typings of method bodies to be retrieved by inversion of method invocation typings. Like Lemma 5.14, they will be instrumental for the method invocation case of subject reduction.

LEMMA 5.17. *Given $\Gamma(\mathbf{C}) = \forall \bar{\mathbf{X}}[\mathbf{C}].[\mathbf{TC}]$ and $\mathbf{mbody}(\mathbf{m}, \mathbf{C}) = \bar{\mathbf{x}}.\mathbf{e}$ for well formed Γ . Then there exists $(\mathbf{m} : \bar{\mathbf{S}} \xrightarrow{\mathbf{H}} \mathbf{S}) \in \mathbf{T}, \bar{\mathbf{R}}$ and \mathbf{R} such that the following are valid:*

$$C \Vdash \mathbf{R} <: \mathbf{S} \wedge \bar{\mathbf{S}} <: \bar{\mathbf{R}} \quad \Gamma; \mathbf{this} : [\mathbf{TC}]; \bar{\mathbf{x}} : \bar{\mathbf{R}}, C, \mathbf{H} \vdash \mathbf{e} : \mathbf{R}$$

Proof. By inversion of the T-CLASS and T-METH rules, and Lemma 5.12. \square

LEMMA 5.18. *Given well formed Γ and:*

$$\Gamma, C, \mathbf{H}_0 \vdash \mathbf{new} \ C(\bar{\mathbf{v}}) : [\mathbf{TC}] \quad \mathbf{mbody}(\mathbf{m}, \mathbf{C}) = \bar{\mathbf{x}}.\mathbf{e}$$

$$C \Vdash \mathbf{T} <: (\mathbf{m} : \bar{\mathbf{S}} \xrightarrow{\mathbf{H}_1} \mathbf{S})$$

Then:

$$\Gamma; \mathbf{this} : [\mathbf{TC}]; \bar{\mathbf{x}} : \bar{\mathbf{R}}, C, \mathbf{H}_2 \vdash \mathbf{e} : \mathbf{R}$$

is derivable, where:

$$C \Vdash \mathbf{R} <: \mathbf{S} \wedge \bar{\mathbf{S}} <: \bar{\mathbf{R}} \wedge \mathbf{H}_2 <: \mathbf{H}_1$$

Proof. By Lemma 5.13 and inversion of the T-NEW rule, we have that $[\mathbf{TC}] = [\mathbf{T}_0 \mathbf{C}][\bar{\mathbf{S}}'/\bar{\mathbf{X}}]$ where $\Gamma(\mathbf{C}) = \forall \bar{\mathbf{X}}[\mathbf{D}].[\mathbf{T}_0 \mathbf{C}]$ and $C \Vdash \mathbf{D}[\bar{\mathbf{S}}'/\bar{\mathbf{X}}]$. Then by Lemma 5.17, there exists $(\mathbf{m} : \bar{\mathbf{R}}_0 \xrightarrow{\mathbf{H}} \mathbf{R}_0) \in \mathbf{T}_0, \bar{\mathbf{S}}_0$ and \mathbf{S}_0 such that:

$$C \Vdash \mathbf{S}_0 <: \mathbf{R}_0 \wedge \bar{\mathbf{R}}_0 <: \bar{\mathbf{S}}_0 \quad \Gamma; \mathbf{this} : [\mathbf{T}_0 \mathbf{C}]; \bar{\mathbf{x}} : \bar{\mathbf{S}}_0, \mathbf{D}, \mathbf{H} \vdash \mathbf{e} : \mathbf{S}_0$$

are both valid. Letting:

$$\bar{\mathbf{R}} \triangleq \bar{\mathbf{S}}_0[\bar{\mathbf{S}}'/\bar{\mathbf{X}}] \quad \mathbf{R} \triangleq \mathbf{S}_0[\bar{\mathbf{S}}'/\bar{\mathbf{X}}] \quad \mathbf{H}_2 \triangleq \mathbf{H}[\bar{\mathbf{S}}'/\bar{\mathbf{X}}]$$

Then by Lemma 5.10:

$$C \Vdash R <: R_0[\bar{S}'/\bar{X}] \wedge \bar{R}_0[\bar{S}'/\bar{X}] <: \bar{R}$$

so by Definition 5.2 and Lemma 5.1:

$$C \Vdash R <: S \wedge \bar{S} <: \bar{R} \wedge H_2 <: H_1$$

and by Lemma 5.9 and Lemma 5.11:

$$\Gamma; \mathbf{this} : [\mathbf{TC}]; \bar{x} : \bar{R}, C, H_2 \vdash e : R$$

which was to be demonstrated. \square

We can now prove subject reduction. Note that the statement of the proof allows for an increase in precision of type and effect in reduced terms.

LEMMA 5.3 (Subject Reduction). *If $\Gamma, C, H \vdash e : [\mathbf{TC}]$ is derivable for closed e and well formed Γ , and $\eta, e \rightarrow \eta', e'$, then $\Gamma, C, H' \vdash e' : [\mathbf{T}'\mathbf{C}']$ is derivable with $C \Vdash [\mathbf{T}'\mathbf{C}'] <: [\mathbf{TC}]$ and $C \Vdash \eta'; H' <: \eta; H$.*

Proof. By induction on $\Gamma, C, H \vdash e : [\mathbf{TC}]$, case analysis on the last step in the derivation, and associated subcase analysis defined by possible last steps in the derivation of $\eta, e \rightarrow \eta', e'$, given the form of e as required by the case. We consider the most interesting and crucial cases, treating method invocation, casting and soft subtyping, events, and effect weakening.

Case T-INVK. In this case, by inversion of the rule we have:

$$e = e_0.m(\bar{e}) \quad H = H_1; H_2; H_3 \quad \Gamma, C, H_1 \vdash e_0 : [\mathbf{SD}] \quad \Gamma, C, H_2 \vdash \bar{e} : \bar{T}$$

$$C \Vdash S <: (m : \bar{T} \xrightarrow{H_3} [\mathbf{TC}])$$

A subcase analysis for the form of e in this case comprises R-INVK, RC-INVK-RECV, and RC-INVK-ARG, as follows.

Subcase R-INVK. In this subcase by rule R-INVK and Lemma 5.15 we have:

$$e_0 = \mathbf{new} D(\bar{v}) \quad \bar{e} = \bar{u} \quad e' = [\bar{u}/\bar{x}, \mathbf{new} D(\bar{v})/\mathbf{this}]e_1$$

$$mbody(m, D) = \bar{x}.e_1 \quad \eta' = \eta$$

Thus by Lemma 5.18:

$$\Gamma; \mathbf{this} : [\mathbf{SD}]; \bar{x} : \bar{S}, C, H_4 \vdash e_1 : [\mathbf{T}_1 \mathbf{C}_1]$$

is derivable, where:

$$C \Vdash [\mathbf{T}_1 \mathbf{C}_1] <: [\mathbf{TC}] \wedge \bar{T} <: \bar{S} \wedge H_4 <: H_3$$

Therefore by Lemma 5.16 and Lemma 5.14 the judgement:

$$\Gamma, C, H'_4 \vdash [\bar{u}/\bar{x}, \text{new } D(\bar{v})/\text{this}]e_1 : [T'_1 C'_1]$$

is derivable, where $C \Vdash [T'_1 C'_1] <: [T_1 C_1]$ and $C \Vdash H'_4 <: H_4$. Thus $C \Vdash [T'_1 C'_1] <: [TC]$ and $C \Vdash H'_4 <: H_3$ by Lemma 5.1, so it only remains to be shown that $C \Vdash \eta; H'_4 <: \eta; H_3$. But $C \Vdash \epsilon <: H_1$ and $C \Vdash \epsilon <: H_2$ by Lemma 5.16, hence $C \Vdash H_3 <: H_1; H_2; H_3$ by Lemma 4.2, so the result follows by Lemma 4.2 and Lemma 5.1.

Subcase RC-INVK-RECV. In this subcase $e' = e_1.m(\bar{e})$ for some e_1 such that $\eta, e_1 \rightarrow \eta', e_1$. But then by the induction hypothesis we have $\Gamma, C, H'_1 \vdash e_1 : [S' D']$ with $C \Vdash \eta'; H'_1 <: \eta; H_1 \wedge [S' D'] <: [SD]$. These facts, Definition 5.2, and Lemma 5.1 imply $C \Vdash S' <: (m : [\bar{R}\bar{B}] \xrightarrow{H_3} [TC])$, hence $\Gamma, C, H'_1; H_2; H_3 \vdash e_1.m(\bar{e}) : [TC]$ is derivable by an instance of T-INVK, with $C \Vdash \eta'; H'_1; H_2; H_3 <: \eta; H_1; H_2; H_3$ by Lemma 4.2, so this case holds. Subcase RC-INVK-ARG follows in a similar manner.

Case T-CAST. In this case, by inversion of the rule we have:

$$e = (C)e_0 \quad \Gamma, C, H \vdash e_0 : [SD] \quad C \Vdash [SD] <: [TC]$$

A subcase analysis for the form of e in this case comprises R-CAST and RC-CAST, as follows.

Subcase R-CAST. In this subcase we have:

$$e_0 = \text{new } D(\bar{v}) \quad D <: C \quad e' = e_0 \quad \eta = \eta'$$

by inversion of R-CAST and Lemma 5.15. But then $C \Vdash [SD] <: [TC]$ by Lemma 5.2; the result follows. Subcase RC-CAST follows in a manner similar to the RC-INVK-RECV subcase of the T-INVK case above.

Case T-EVENT. In this case by inversion of the rule we have:

$$e = \text{ev}[i] \quad H = \text{ev}[i] \quad [TC] = \text{Unit}$$

The only reduction rule that applies to the form of e in this case is R-EVENT, inversion of which obtains:

$$e' = () \quad \eta' = \eta; \text{ev}[i]$$

But $\Gamma, C, \epsilon \vdash () : \text{Unit}$ by assumption and Lemma 5.16, so this case holds.

Case T-WEAKEN. In this case we have $\Gamma, C, H' \vdash e : [TC]$ with $C \Vdash H' <: H$ by inversion of rule T-WEAKEN. But then $\Gamma, C, H'' \vdash e' : [T' C']$ with $C \Vdash \eta'; H'' <: \eta; H' \wedge [T' C'] <: [TC]$ by the induction hypothesis, and $C \Vdash \eta'; H'' <: \eta; H$ Lemma 4.2 and Lemma 5.1, so this case holds. \square

We then prove one auxiliary lemma followed by our main type safety result, demonstrating that run-time checks in well-typed programs are guaranteed to succeed:

LEMMA 5.19. *If $\Gamma, C, \mathbb{H} \vdash \mathbf{E}[\mathbf{chk}[x]] : \mathbf{T}$ then $C \Vdash \mathbf{chk}[x]; \mathbb{H}' <: \mathbb{H}$.*

LEMMA 5.4 (Static Enforcement of Checks). *Given closed e and well-formed Γ , if the judgement $\Gamma, C, \mathbb{H} \vdash \mathbf{e} : \mathbf{T}$ is valid, and $\epsilon, \mathbf{e} \rightarrow^* \eta, \mathbf{E}[\mathbf{chk}[x]]$, then $\eta \vdash \mathbf{chk}[x]$.*

Proof. By Lemma 5.3 and Lemma 5.19, $\Gamma, C, \mathbb{H}' \vdash \mathbf{E}[\mathbf{chk}[x]] : \mathbf{T}$ is derivable with $C \Vdash \mathbf{chk}[x]; \mathbb{H}'' <: \mathbb{H}'$ and $C \Vdash \eta; \mathbb{H}' <: \mathbb{H}$. Now, by assumption there exists a solution ρ of C such that $\rho(\mathbb{H})$ is valid; but since $\llbracket \rho(\mathbf{chk}[x]; \mathbb{H}'') \rrbracket \subseteq \llbracket \rho(\mathbb{H}') \rrbracket$ and $\llbracket \rho(\eta; \mathbb{H}') \rrbracket \subseteq \llbracket \rho(\mathbb{H}) \rrbracket$ by previous facts and Definition 5.2, therefore $\llbracket \rho(\eta; \mathbf{chk}[x]; \mathbb{H}'') \rrbracket \subseteq \llbracket \rho(\mathbb{H}) \rrbracket$. By Lemma 4.2 it follows that the effect $\rho(\eta; \mathbf{chk}[x]; \mathbb{H}')$ is valid, thus $\eta \vdash \mathbf{chk}[x]$ by Definition 4.1. \square

6. Type Inference for $\mathbf{FJ}_{\text{trace}}$

We now develop an implementation of the $\mathbf{FJ}_{\text{trace}}$ type system. This includes a type inference algorithm for reconstructing type judgements. In addition, it is necessary to check satisfiability of constraints, to ensure that these judgements are coherent. For this purpose we adapt standard *closure* techniques for subtyping constraints (Palsberg and O’Keefe, 1995), extended to accommodate effect constraints. Closure distills inferred constraints into their basic component elements, upon which a simple structural consistency check can be used to ensure satisfiability. Inference and closure serve as preliminary phases for statically verifying trace based program assertions, which can finally be accomplished by model checking trace effects, as in (Skalka and Smith, 2004). Model checking is applicable, since we endow trace effects with an LTS semantics (Sect. 4), for which a variety of model checking techniques exist (Burkart et al., 2001). However, standard techniques expect term, rather than constraint, representation of LTSs. Therefore, it is also necessary to define a means of extracting a unified trace effect representation from inferred typing judgements. For this purpose we define an algorithm, called *hextract*, to obtain a unified representation of trace effects from the inferred constraint representation. The composition of type inference, closure, and hextraction obtain an algorithmic technique for enforcing type safety, as observed in Theorem 6.1.

6.1. THE TYPE LANGUAGE

The type and constraint language used by inference is the same as that presented in Sect. 5, albeit specialized to integrate neatly with closure.

$$\begin{array}{c}
\frac{
\begin{array}{l}
mtype(\mathbf{m}, \mathbf{C}) = \bar{\mathbf{D}} \rightarrow \mathbf{D} \quad CT(\mathbf{C}) = \text{class } \mathbf{C} \text{ extends } \mathbf{B} \{ \dots \} \\
\text{if } mtype(\mathbf{m}, \mathbf{B}) = \bar{\mathbf{E}} \rightarrow \mathbf{E} \text{ then } \bar{\mathbf{E}} = \bar{\mathbf{D}} \text{ and } \mathbf{E} = \mathbf{D}
\end{array}
}{
hexpand(\mathbf{m}, \mathbf{C}) = [\bar{\mathbf{t}} \bar{\mathbf{D}}] \xrightarrow{h} [\bar{\mathbf{t}} \mathbf{D}]
} \\
\\
\frac{
\begin{array}{l}
fields(\mathbf{C}) = \bar{\mathbf{D}} \bar{\mathbf{f}} \quad meths(\mathbf{C}) = \bar{\mathbf{m}} \quad \forall \mathbf{m}_i \in \bar{\mathbf{m}}. hexpand(\mathbf{m}_i, \mathbf{C}) = \mathbf{T}_i
\end{array}
}{
hexpand(\mathbf{C}) = [\bar{\mathbf{f}} : [\bar{\mathbf{t}} \bar{\mathbf{D}}] \bar{\mathbf{m}} : \bar{\mathbf{T}} \mathbf{C}]
} \\
\\
hexpand(\bar{\mathbf{C}}) = hexpand(\mathbf{C}_1), \dots, hexpand(\mathbf{C}_n)
\end{array}$$

Figure 12. *hexpand* type construction

There are two forms of object types used by inference, a form $[\bar{\mathbf{t}} \mathbf{C}]$ that we call *abstract*, and an *expanded* form defined in Fig. 12. This form is similar to the expanded form of object types defined in Sect. 5.4, where field and method types of a given class are explicitly listed, except method types are assigned abstract effects, and *canonical* hexpansion will always choose fresh variables with which to construct the type.

Inference always assigns expanded type forms to expressions (see Lemma 6.6 for a precise statement of this), which invariant provides a uniform representation for proofs. More importantly, the invariant preserves soundness when imposing constraints due to method or field selection. That is, the selection of a method or field \mathbf{x} from an object of type $[\mathbf{T} \mathbf{C}]$ will impose a constraint of the form $\mathbf{T} <: (\mathbf{x} : \mathbf{S})$. If $[\mathbf{T} \mathbf{C}]$ is in expanded form, a lower bound constraint on the width of \mathbf{T} is hard-coded by the form of \mathbf{T} , whereas if \mathbf{T} was a type variable $\bar{\mathbf{t}}$, a lower bound width constraint on $\bar{\mathbf{t}}$ would have to be explicitly imposed to ensure soundness.

Since expanded types do have abstract forms as subterms, the latter occur in types and constraints in inference judgements, but only serve as conduits for transitive flow during closure. Furthermore, type variables $\bar{\mathbf{t}}$ never appear “bare” during inference, but only within an abstract object type form $[\bar{\mathbf{t}} \mathbf{C}]$, so that object type variables always have an implicit width constraint associated with \mathbf{C} .

6.2. TYPE INFERENCE JUDGEMENTS

Type inference rules are given in Fig. 13; the W subscripting the relation \vdash_W distinguishes type inference from logical typing judgements, and is named after the polymorphic type reconstruction algorithm W (Damas and Milner, 1982). The type inference rules are deterministic except for the choice of type variables; we call *canonical* those derivations that al-

T-VAR $\Gamma, \mathbf{true}, \epsilon \vdash_W \mathbf{x} : \Gamma(\mathbf{x})$	
T-FIELD $\frac{\Gamma, C, H \vdash_W \mathbf{e} : [\mathbf{TC}] \quad D \mathbf{f} \in \mathit{fields}(C) \quad \mathit{expand}(D) = [\mathbf{SD}]}{\Gamma, C \wedge \mathbf{T} <: (\mathbf{f} : [\mathbf{SD}]), H \vdash_W \mathbf{e}.\mathbf{f} : [\mathbf{SD}]}$	
T-SEQNIL $\Gamma, \mathbf{true}, \epsilon \vdash_W \emptyset : \emptyset$	T-SEQCONS $\frac{\Gamma, C_1, H_1 \vdash_W \bar{\mathbf{e}} : \bar{\mathbf{T}} \quad \Gamma, C_2, H_2 \vdash_W \mathbf{e} : \mathbf{T}}{\Gamma, C_1 \wedge C_2, H_1; H_2 \vdash_W \bar{\mathbf{e}}, \mathbf{e} : \bar{\mathbf{T}}, \mathbf{T}}$
T-INVK $\frac{\Gamma, D, H' \vdash_W \bar{\mathbf{e}} : \bar{\mathbf{S}} \quad \Gamma, C, H \vdash_W \mathbf{e} : [\mathbf{TC}] \quad \mathit{mtype}(\mathbf{m}, C) = \bar{D} \rightarrow D \quad \mathit{expand}(D) = [\mathbf{SD}]}{\Gamma, C \wedge D \wedge \mathbf{T} <: (\mathbf{m} : \bar{\mathbf{S}} \xrightarrow{h} [\mathbf{SD}]), H; H'; h \vdash_W \mathbf{e}.\mathbf{m}(\bar{\mathbf{e}}) : [\mathbf{SD}]}$	
T-NEW $\frac{\Gamma(C) = \forall \bar{X}[D].[\mathbf{TC}] \quad \mathit{fieldsig}[\mathbf{TC}] = \bar{\mathbf{T}} \quad \Gamma, C, H \vdash_W \bar{\mathbf{e}} : \bar{\mathbf{S}}}{\Gamma, C \wedge D[\bar{X}'/\bar{X}] \wedge \bar{\mathbf{S}} <: \bar{\mathbf{T}}[\bar{X}'/\bar{X}], H \vdash_W \mathbf{new} C(\bar{\mathbf{e}}) : [\mathbf{T}[\bar{X}'/\bar{X}]} C}$	
T-EVENT $\Gamma, \mathbf{true}, \mathit{ev}[i] \vdash_W \mathit{ev}[i] : \mathbf{Unit}$	T-CHECK $\Gamma, \mathbf{true}, \mathit{chk}[i] \vdash_W \mathit{chk}[i] : \mathbf{Unit}$
T-CAST $\frac{\Gamma, C, H \vdash_W \mathbf{e} : \mathbf{T} \quad \mathit{expand}(D) = [\mathbf{SD}]}{\Gamma, C \wedge \mathbf{T} <: [\mathbf{SD}], H \vdash_W (D)\mathbf{e} : [\mathbf{SD}]}$	
T-METH $\frac{\Gamma; \bar{\mathbf{x}} : \bar{\mathbf{S}}, C, H \vdash_W \mathbf{e} : \mathbf{S} \quad \mathit{expand}(\bar{D}) = \bar{\mathbf{S}} \quad \Gamma(\mathbf{this}).\mathbf{m} = [\bar{\mathbf{T}}D] \xrightarrow{h} \mathbf{T} \quad \mathit{mbody}(\mathbf{m}, C) = \bar{\mathbf{x}}.\mathbf{e}}{\Gamma, C \wedge \mathbf{S} <: \mathbf{T} \wedge [\bar{\mathbf{T}}D] <: \bar{\mathbf{S}} \wedge H <: h \vdash_W \mathbf{m}, C : [\bar{\mathbf{T}}D] \xrightarrow{h} \mathbf{T}}$	
T-CLASSCONSTRAINT $\frac{\mathit{meths}(C) = \bar{\mathbf{m}} \quad \Gamma; \mathbf{this} : \Gamma(C), D_i \vdash_W \mathbf{m}_i, C : \mathbf{T}_i \text{ for all } \mathbf{m}_i \in \bar{\mathbf{m}}}{\Gamma \vdash_W C : D_1 \wedge \dots \wedge D_n, \Gamma(C)}$	
T-CLASS $\frac{\Gamma; \bar{C} : \mathit{expand}(\bar{C}) \vdash_W C_i : D_i, \mathbf{T}_i \text{ and } \mathit{fv}(D, \mathbf{T}_i) = \bar{X}_i \text{ for all } i \in 1..n \quad D = D_1 \wedge \dots \wedge D_n}{\Gamma \vdash_W C_1 : \forall \bar{X}_1[D].\mathbf{T}_1, \dots, C_n : \forall \bar{X}_n[D].\mathbf{T}_n}$	

Figure 13. FJ_{trace} type inference rules

ways choose fresh type variables, and hereafter restrict our consideration to canonical derivations without loss of generality.

Type inference rules are mostly analogous to their logical counterparts. The T-CLASS rule is slightly different, and the T-CONSTRAINT rule interposed, to deal with mutually recursive class definitions. If \bar{C} are all mutually recursive, then the T-CLASS rule allows typings to be assigned to them “in parallel”, as a group— a group of one, in case a given class is not mutually recursive with any others. We specify a normal form for environments in inference derivations, and also we define the inference analog of well formed environments, which formalizes the notion of complete typing for a full class table.

DEFINITION 6.1. *Any Γ is in inference form iff $\text{hexpand}(C) = T$ for all $C : \forall \bar{X}[C]. T \in \Gamma$, and for all $x : S \in \Gamma$ there exists D such that $\text{hexpand}(D) = T$.*

Hereafter, we restrict inference judgements to the use of environments in inference form, and observe that derivations preserve this property.

We obtain soundness for type inference via the following result for expression inference; generalization to method and class inference are obtained on this basis. The result follows by induction on inference derivations. The Lemma states that inference derivations can be reconstructed as logical derivations of less general typings; this formulation is necessary to allow the induction to go through, since logical judgements are given complete constraints a priori, whereas they are reconstructed from the leaves towards the root in inference derivations. The proof is given in Sect. 6.4.

LEMMA 6.1 (Soundness of Inference). *If $\Gamma, C, H \vdash_W e : T$ is derivable with Γ and $C \wedge D, T$ realizable, then $\Gamma, C \wedge D, H \vdash e : T$ is derivable.*

6.3. CLOSURE AND EXTRACTION OF EFFECTS

To automatically check satisfiability of constraints, hence realizability of types, the type implementation comprises a constraint closure algorithm and consistency check. We say that a constraint C is *consistent* iff $\vdash C : ok$ is derivable given the deterministic rules in Fig. 15. For brevity in the definition of closure, we introduce the following notation:

DEFINITION 6.2. *Let \hat{C} range over atomic constraints, i.e.:*

$$\hat{C} ::= \mathbf{true} \mid T <: T \mid T \triangleleft T$$

and for all $C = \hat{C}_1 \wedge \dots \wedge \hat{C}_n$, let $\text{set}(C) = \{C_1, \dots, C_n\}$. Then define:

$$\hat{C} \in C \iff \hat{C} \in \text{set}(C) \quad D \subseteq C \iff \text{set}(D) \subseteq \text{set}(C)$$

$$\begin{array}{c}
\text{C-FN} \\
(\bar{T} \xrightarrow{H} T <: \bar{S} \xrightarrow{H'} S) \rightsquigarrow_{close} (\bar{S} <: \bar{T} \wedge T <: S \wedge H <: H') \\
\\
\text{C-TRANS} \qquad \qquad \qquad \text{C-STTRANS} \\
(R <: S \wedge S <: T) \rightsquigarrow_{close} R <: T \qquad (R <: S \wedge S \triangleleft T) \rightsquigarrow_{close} R \triangleleft T \\
\\
\text{C-OBJ} \\
[(\bar{x} : \bar{T}) C] <: [(\bar{y} : \bar{S}) D] \rightsquigarrow_{close} (\bar{x} : \bar{T}) <: (\bar{y} : \bar{S}) \\
\\
\text{C-ROW} \qquad \qquad \qquad \text{C-CAST} \\
(\bar{x} : \bar{R}, \bar{y} : \bar{S}) <: (\bar{x} : \bar{T}) \rightsquigarrow_{close} \bar{R} <: \bar{T} \qquad \frac{C <: D}{[TC] \triangleleft [SD] \rightsquigarrow_{close} [TC] <: [SD]} \\
\\
\text{C-CONTEXT} \\
\frac{C' \subseteq C \quad C' \rightsquigarrow_{close} D \quad D \not\subseteq C}{C \rightarrow_{close} C \wedge D}
\end{array}$$

Figure 14. Constraint closure rules

Constraint closure is then defined via the rewrite rules given in Fig. 14 and Definition 6.3. The closure rules are mostly standard, except for those that treat soft subtyping constraints, C-CAST and C-STTRANS. These rules implement selective pruning of the constraint graph described previously; note that they effectively discard unsound flow along soft subtyping edges.

DEFINITION 6.3 (Closure). *The rewrite relations \rightsquigarrow_{close} and \rightarrow_{close} are defined in Fig. 14. C is closed iff there does not exist D such that $C \rightarrow_{close} D$. The relation \rightarrow_{close}^* is the reflexive, transitive closure of \rightarrow_{close} . We define $close(C)$ as a closed constraint such that $C \rightarrow_{close}^* close(C)$.*

Correctness of closure for inferred types and constraints is stated as follows. While our interpretation has slight modifications for specialized object type forms, our approach to type constraint closure is fundamentally standard, and correctness of standard closure techniques has been thoroughly treated in previous work, notably (Palsberg and O’Keefe, 1995). Solvability of closed, consistent effect constraints is established constructively in Lemma 6.3, and solutions of type and effect constraints can be composed to obtain a solution for whole constraints generated by inference. This argument is developed in more detail in (Skalka et al., 2006).

$$\boxed{
\begin{array}{c}
\vdash \mathbf{true} : ok \qquad \frac{\vdash C : ok \quad \vdash D : ok}{\vdash C \wedge D : ok} \qquad \vdash H <: H' : ok \\
\\
\frac{C <: D}{\vdash [TC] <: [SD] : ok} \qquad \vdash [TC] < [SD] : ok \qquad \vdash (\bar{T} \xrightarrow{H} T <: \bar{S} \xrightarrow{H'} S) : ok \\
\\
\vdash (\bar{x} : \bar{R}, \bar{y} : \bar{S}) <: (\bar{x} : \bar{T}) : ok
\end{array}
}$$

Figure 15. Constraint consistency rules

LEMMA 6.2. *If $\Gamma, C, H \vdash_W e : T$ is derivable, then C is satisfiable iff $close(C)$ is consistent.*

A key property of effect constraints that ensures their satisfiability, is that they define a system of lower bounds on effect variables. That is, if $\Gamma, C, H \vdash_W e : T$ is derivable and $H <: H' \in close(C)$, then H' is an effect variable h . It is easy to demonstrate this property by observing that it is established by inference, and preserved by closure. This form of effect constraint implies that, given $\Gamma, C, H \vdash_W e : T$, a solution for H can be obtained, more or less, by recursively joining the lower bounds of type variable components of H in $close(C)$. The algorithm for extracting term representation of effects from a constraint representation, called *hextract*, is defined in Fig. 16. Since extraction does not alter the given constraint in any way, the fixed parameter C is written as a subscript. Extraction returns a closed trace effect that is a sound term representation of the top-level effect of given closed expressions, in the following sense. The result constructively proves satisfiability of closed, consistent effect constraints.

LEMMA 6.3 (Extraction Correctness). *If $\Gamma, C, H \vdash_W e : T$ is derivable for closed e and $close(C)$ is consistent, then $hextract_{close(C)}(H)$ is defined and there exists a solution ρ of C such that $hextract_{close(C)}(H) = \rho(H)$.*

Correctness of the *hextract* algorithm is rigorously established in (Skalka et al., 2006) by a fixpoint construction argument. Since the form of constraints treated in that result is the same as that generated by inference in this paper, the result is entirely applicable.

To make the final connection to logical typing judgements, we must define the analogue of well-formed environments, wherein all type schemes bindings are realizable and logically derivable. In the inference system, this requires that environments contain bindings that are logically inferable, and whose constraint closures are consistent.

$$\begin{aligned}
\text{hextract}_C(\epsilon, hs) &= \epsilon \\
\text{hextract}_C(\text{ev}[i], hs) &= \text{ev}[i] \\
\text{hextract}_C(\text{chk}[i], hs) &= \text{chk}[i] \\
\text{hextract}_C(\mathbf{H}_1; \mathbf{H}_2, hs) &= (\text{hextract}_C(\mathbf{H}_1, hs)); (\text{hextract}_C(\mathbf{H}_2, hs)) \\
\text{hextract}_C(\mathbf{H}_1 | \mathbf{H}_2, hs) &= (\text{hextract}_C(\mathbf{H}_1, hs)) | (\text{hextract}_C(\mathbf{H}_2, hs)) \\
\text{hextract}_C(\mathbf{h}, \{\mathbf{h}\} \cup hs) &= \mathbf{h} \\
\text{hextract}_C(\mathbf{h}, hs) &= \mu \mathbf{h}. \text{hextract}_C(\text{bounds}_C(\mathbf{h}), hs \cup \{\mathbf{h}\}) \\
\text{bounds}_C(\mathbf{h}) &= \mathbf{H}_1 | \dots | \mathbf{H}_n \quad \text{where } \{\mathbf{H}_1, \dots, \mathbf{H}_n\} = \{\mathbf{H} \mid \mathbf{H} <: \mathbf{h} \in C\}
\end{aligned}$$

Figure 16. *hextract* and *bounds* functions

DEFINITION 6.4. *Letting σ range over constrained type schemes, and given:*

$$\Gamma = \mathbf{C}_1 : \sigma_1; \dots; \mathbf{C}_n : \sigma_n$$

Then Γ is inferable iff for all $0 < i \leq n$ there exists $0 < k \leq i$ such that:

$$\mathbf{C}_1 : \sigma_1; \dots; \mathbf{C}_{k-1} : \sigma_{k-1} \vdash_W \mathbf{C}_k : \sigma_k, \dots, \mathbf{C}_i : \sigma_i$$

and for all $\forall \bar{\mathbf{X}}[C]. \mathbf{T} \in \Gamma$, $\text{close}(C)$ is consistent.

Putting the pieces together, a complete analysis is obtained by composition of inference, closure, consistency, and *hextraction*. Validity of effects can be implemented by techniques described in (Skalka and Smith, 2004).

THEOREM 6.1 (Soundness of Analysis). *Suppose $\Gamma, C, \mathbf{H} \vdash_W \mathbf{e} : \mathbf{T}$ is derivable for inferable Γ and closed \mathbf{e} , $\text{close}(C)$ is consistent, and $\text{hextract}_{\text{close}(C)}(\mathbf{H})$ is valid. Then $\epsilon, \mathbf{e} \rightarrow^* \eta, \mathbf{E}[\text{chk}[\mathbf{x}]]$ implies $\eta \vdash \text{chk}[\mathbf{x}]$.*

The proof is given in Sect. 6.4, following by soundness of inference as established by Lemma 6.1, Lemma 6.2, Lemma 6.3, and type safety as established by Lemma 5.4.

6.4. PROPERTIES

We now investigate formal properties of type inference in more depth. First, we observe a fairly obvious characteristic of constraints, that any solution of a compound constraint is also a solution of its parts.

LEMMA 6.4. $C \wedge D \Vdash D$.

Now, we characterize the essential property allowing the innards of an hexpanded object type to be unwrapped, while still retaining a width constraint associated with the object class. This is essential for ensuring soundness of inference. The result is a straightforward consequence of the definition of *hexpand*.

LEMMA 6.5. *Let $\text{hexpand}(\mathbf{C}) = [\mathbf{T}\mathbf{C}]$. Then if $\rho(\mathbf{T})$ is well kinded, so is $\rho([\mathbf{T}\mathbf{C}])$.*

We also observe the invariant that all top-level types in inference are in hexpanded form, which are well-formed with respect to any inferred top-level constraint. The result follows by induction on type derivations.

LEMMA 6.6. *If $\Gamma, C, \mathbf{H} \vdash_W \mathbf{e} : \mathbf{T}$ is derivable for well formed Γ , then $\mathbf{T} = \text{hexpand}(\mathbf{C})$, and satisfiability of C implies realizability of C, \mathbf{T} .*

Soundness of inference follows by induction on inference derivations. In essence, soundness is proved by transforming an inference derivation into a logical derivation, by allowing the root constraint to be pushed towards the leaves in the induction, via appropriate statement of the result (i.e. the logical derivation permits weakening of the top level constraint with some constraint D).

LEMMA 6.1 (Soundness of Inference). *If $\Gamma, C, \mathbf{H} \vdash_W \mathbf{e} : \mathbf{T}$ is derivable with Γ and $C \wedge D, \mathbf{T}$ realizable, then $\Gamma, C \wedge D, \mathbf{H} \vdash \mathbf{e} : \mathbf{T}$ is derivable.*

Proof. By induction on the derivation of $\Gamma, C, \mathbf{H} \vdash_W \mathbf{e} : [\mathbf{T}\mathbf{C}]$ and case analysis on the last step.

Cases T-VAR, T-SEQNIL, T-EVENT, and T-CHECK are immediate.

Case T-FIELD. In this case, by inversion of the rule we have:

$$\begin{array}{l} \mathbf{e} = \mathbf{e}_0.\mathbf{f} \quad \text{hexpand}(\mathbf{C}) = [\mathbf{T}\mathbf{C}] \quad C = C_0 \wedge \mathbf{S} <: (\mathbf{f} : [\mathbf{T}\mathbf{C}]) \\ \Gamma, C_0, \mathbf{H} \vdash_W \mathbf{e} : [\mathbf{S}\mathbf{D}] \quad C \mathbf{f} \in \text{fields}(\mathbf{D}) \end{array}$$

Since by assumption C has a solution ρ , therefore $\rho(\mathbf{S})$ is well kinded, so $C, [\mathbf{S}\mathbf{D}]$ is realizable by Lemma 6.6 and Lemma 6.5. Hence $\Gamma, C, \mathbf{H} \vdash \mathbf{e} : [\mathbf{S}\mathbf{D}]$ is derivable by the induction hypothesis, and $C \Vdash \mathbf{S} <: (\mathbf{f} : [\mathbf{T}\mathbf{C}])$ by Lemma 6.4, so the result follows by an instance of T-FIELD.

Case T-INVK. By inversion of the rule in this case we have:

$$\begin{array}{l} C = C_1 \wedge C_2 \wedge \mathbf{S} <: (\mathbf{m} : [\bar{\mathbf{S}}\mathbf{B}] \xrightarrow{\mathbf{h}} [\mathbf{T}\mathbf{C}]) \quad \mathbf{e} = \mathbf{e}_0.\mathbf{m}(\bar{\mathbf{e}}) \quad \mathbf{H} = \mathbf{H}_1; \mathbf{H}_2; \mathbf{h} \\ \Gamma, C_1, \mathbf{H}_1 \vdash_W \mathbf{e}_0 : [\mathbf{S}\mathbf{D}] \quad \Gamma, C_2, \mathbf{H}_2 \vdash_W \bar{\mathbf{e}} : [\bar{\mathbf{S}}\mathbf{B}] \\ \text{mtype}(\mathbf{m}, \mathbf{D}) = \bar{\mathbf{C}} \rightarrow \mathbf{C} \quad \text{hexpand}(\mathbf{C}) = [\mathbf{T}\mathbf{C}] \end{array}$$

Since C has a solution ρ by assumption, therefore $\rho(\mathbf{S})$ and $\rho([\bar{\mathbf{S}}\bar{\mathbf{B}}])$ are well kinded, so $C, [\mathbf{S}\bar{\mathbf{B}}]$ is realizable, and $C, [\mathbf{S}\bar{\mathbf{D}}]$ is realizable by Lemma 6.6 and Lemma 6.5. Thus, $\Gamma, C, \mathbf{H}_1 \vdash \mathbf{e}_0 : [\mathbf{S}\bar{\mathbf{D}}]$ and $\Gamma, C, \mathbf{H}_2 \vdash \bar{\mathbf{e}} : [\bar{\mathbf{S}}\bar{\mathbf{B}}]$ by the induction hypothesis, and $C \Vdash \mathbf{S} < : (\mathbf{m} : [\bar{\mathbf{S}}\bar{\mathbf{B}}] \xrightarrow{\mathbf{h}} [\mathbf{T}\bar{\mathbf{C}}])$ by Lemma 6.4, so the result follows by an instance of T-INVK.

Case T-NEW. In this case, by inversion of the rule we have:

$$\begin{aligned} \mathbf{e} = \mathbf{new} \ \mathbf{C}(\bar{\mathbf{e}}) \quad \mathbf{T} = [\mathbf{S}[\bar{\mathbf{X}}'/\bar{\mathbf{X}}] \ \mathbf{C}] \quad C = C_0 \wedge D_0[\bar{\mathbf{X}}'/\bar{\mathbf{X}}] \wedge \bar{\mathbf{S}} < : \bar{\mathbf{T}}[\bar{\mathbf{X}}'/\bar{\mathbf{X}}] \\ \Gamma(\mathbf{C}) = \forall \bar{\mathbf{X}}[D_0].[\mathbf{S}\bar{\mathbf{C}}] \quad \mathit{fieldsig}[\mathbf{S}\bar{\mathbf{C}}] = \bar{\mathbf{T}} \quad \Gamma, C_0, \mathbf{H} \vdash_W \bar{\mathbf{e}} : \bar{\mathbf{S}} \end{aligned}$$

But $C \wedge D$ is satisfiable by assumption, so $C \wedge D, \bar{\mathbf{S}}$ is realizable, therefore $\Gamma, C \wedge D, \mathbf{H} \vdash \bar{\mathbf{e}} : \bar{\mathbf{S}}$ is derivable by the induction hypothesis. And by Lemma 6.4:

$$C \wedge D \Vdash D_0[\bar{\mathbf{X}}'/\bar{\mathbf{X}}] \quad C \wedge D \Vdash \bar{\mathbf{S}} < : \bar{\mathbf{T}}[\bar{\mathbf{X}}'/\bar{\mathbf{X}}]$$

so the result follows in this case by an instance of T-NEW. The T-CAST and T-SEQ cases follow in a straightforward manner by the induction hypothesis. \square

A final technical hurdle is to show that inferable environments are well-formed. The result shows how to obtain logical class typing judgements from inferred ones with consistent closures.

LEMMA 6.7. *If Γ is inferable, then it is well-formed.*

Proof. Assume that the following judgement is derivable in the inference system, with $\mathit{close}(D)$ consistent:

$$\Gamma \vdash_W \mathbf{C}_1 : \forall \bar{\mathbf{X}}_1[D].\mathbf{T}_1, \dots, \mathbf{C}_n : \forall \bar{\mathbf{X}}_n[D].\mathbf{T}_n$$

Then it suffices to show that for any $i \in [1..n]$ the judgement:

$$\Gamma; \mathbf{C}_1 : \forall \bar{\mathbf{X}}_1[D].\mathbf{T}_1; \dots; \mathbf{C}_n : \forall \bar{\mathbf{X}}_n[D].\mathbf{T}_n \vdash \mathbf{D} : \forall \bar{\mathbf{X}}_i[D].\mathbf{T}_i$$

is derivable in the logical system. Inverting the T-CLASS rule we can reconstruct:

$$\frac{\Gamma; \bar{\mathbf{C}} : \mathit{expand}(\bar{\mathbf{C}}) \vdash_W \mathbf{C}_i : D_i, \mathbf{T}_i \text{ and } \mathit{fv}(D, \mathbf{T}_i) = \bar{\mathbf{X}}_i \text{ for all } i \in 1..n \quad D = D_1 \wedge \dots \wedge D_n}{\Gamma \vdash_W \mathbf{C}_1 : \forall \bar{\mathbf{X}}_1[D].\mathbf{T}_1, \dots, \mathbf{C}_n : \forall \bar{\mathbf{X}}_n[D].\mathbf{T}_n}$$

where for each $i \in 1..n$, the constraint D_i is of the form $C_1 \wedge \dots \wedge C_j$, and for each $\mathbf{m}_k \in \mathit{meths}(\mathbf{C}_i)$ a judgement:

$$\Gamma; \bar{\mathbf{C}} : \mathit{expand}(\bar{\mathbf{C}}); \mathbf{this} : \Gamma(\mathbf{C}), C_k \vdash_W \mathbf{m}_k, \mathbf{C}_i : \mathbf{T}_k$$

is derivable, by inversion of the T-CLASSCONSTRAINT rule. Let:

$$\Gamma' \triangleq (\Gamma; \bar{c} : \text{hexpand}(\bar{c}); \text{this} : \Gamma(\mathbf{C}))$$

Then by inversion of the T-METH rule we have that \mathbf{T}_k is of the form $[\bar{\mathbf{T}}\bar{\mathbf{D}}] \xrightarrow{\mathbf{h}} \mathbf{T}$, and C_k is of the form:

$$C \wedge \mathbf{S} <: \mathbf{T} \wedge [\bar{\mathbf{T}}\bar{\mathbf{D}}] <: \bar{\mathbf{S}} \wedge \mathbf{H} <: \mathbf{h}$$

and we can reconstruct:

$$\frac{\Gamma'; \bar{x} : \bar{\mathbf{S}}, C, \mathbf{H} \vdash_W \mathbf{e} : \mathbf{S} \quad \text{hexpand}(\bar{\mathbf{D}}) = \bar{\mathbf{S}} \quad \Gamma'(\text{this}).\mathbf{m}_k = [\bar{\mathbf{T}}\bar{\mathbf{D}}] \xrightarrow{\mathbf{h}} \mathbf{T} \quad \text{mbody}(\mathbf{m}_k, \mathbf{C}_i) = \bar{x}.\mathbf{e}}{\Gamma', C \wedge \mathbf{S} <: \mathbf{T} \wedge [\bar{\mathbf{T}}\bar{\mathbf{D}}] <: \bar{\mathbf{S}} \wedge \mathbf{H} <: \mathbf{h} \vdash_W \mathbf{m}_k, \mathbf{C}_i : [\bar{\mathbf{T}}\bar{\mathbf{D}}] \xrightarrow{\mathbf{h}} \mathbf{T}}$$

Since $\text{close}(D)$ is consistent by assumption, therefore D has a solution by Lemma 6.2. And since \mathbf{S} occurs in D , therefore D, \mathbf{S} is realizable, hence by Lemma 6.1 the judgement $\Gamma'; \bar{x} : \bar{\mathbf{S}}, D, \mathbf{H} \vdash \mathbf{e} : \mathbf{S}$ is derivable in the logical system, from which can be derived $\Gamma'; \bar{x} : \bar{\mathbf{S}}, D, \mathbf{h} \vdash \mathbf{e} : \mathbf{S}$ by Lemma 6.4 and an instance of T-WEAKEN. But then the form of C_k as noted above, Lemma 6.4, and an instance of T-METH in the logical system allows us to derive $\Gamma', D \vdash \mathbf{m}_k, \mathbf{C}_i : [\bar{\mathbf{T}}\bar{\mathbf{D}}] \xrightarrow{\mathbf{h}} \mathbf{T}$. But since $[\bar{\mathbf{T}}\bar{\mathbf{D}}]$ and \mathbf{h} and \mathbf{T} all occur in D , therefore $D, [\bar{\mathbf{T}}\bar{\mathbf{D}}] \xrightarrow{\mathbf{h}} \mathbf{T}$ is realizable. The result follows by an instance of T-CLASS in the logical system. \square

Now we can demonstrate our main result, showing how the different components of inference can be combined to generate a sound automated type analysis. The result follows by soundness of inference and closure, correctness of *hextraction*, and type safety.

THEOREM 6.1 (Soundness of Analysis). *Suppose $\Gamma, C, \mathbf{H} \vdash_W \mathbf{e} : \mathbf{T}$ is derivable for inferable Γ and closed \mathbf{e} , $\text{close}(C)$ is consistent, and $\text{hextract}_{\text{close}(C)}(\mathbf{H})$ is valid. Then $\epsilon, \mathbf{e} \rightarrow^* \eta, \mathbf{E}[\text{chk}[\mathbf{x}]]$ implies $\eta \vdash \text{chk}[\mathbf{x}]$.*

Proof. If Γ is inferable, then it is well-formed by Lemma 6.7. If $\text{close}(C)$ is consistent, then C has a solution by Lemma 6.2. Therefore C, \mathbf{T} is realizable by Lemma 6.6, and $\Gamma, C, \mathbf{H} \vdash \mathbf{e} : \mathbf{T}$ is derivable in the logical system by Lemma 6.1. By Lemma 6.3, there exists a solution ρ of C such that $\text{hextract}_{\text{close}(C)}(\mathbf{H}) = \rho(\mathbf{H})$, the validity of which establishes validity of $\Gamma, C, \mathbf{H} \vdash \mathbf{e} : \mathbf{T}$. The result follows by Lemma 5.4. \square

R-THROW $\eta, \mathbf{E}[\mathbf{throw}] \rightarrow \eta, \mathbf{throw}$	
RC-TRY $\frac{\eta, e_1 \rightarrow \eta', e'_1}{\eta, \mathbf{try}\{e_1\}\mathbf{catch}\{e_2\} \rightarrow \eta, \mathbf{try}\{e'_1\}\mathbf{catch}\{e_2\}}$	
R-TRY $\frac{v \neq \mathbf{throw}}{\eta, \mathbf{try}\{v\}\mathbf{catch}\{e_2\} \rightarrow \eta, v}$	R-CATCH $\eta, \mathbf{try}\{\mathbf{throw}\}\mathbf{catch}\{e_2\} \rightarrow \eta, e_2$

Figure 17. Semantics of exceptions

7. Control Flow and Effect Transformations

In this section we show that a benefit of our approach is that trace effect representations are amenable to transformational techniques, for flexibility of analysis. These transformations can be used to post-process inferred effects, without requiring any reworking of the inference component of analysis. This means that certain extensions of the language can be treated statically in a *modular* fashion, including extensions that have a fundamental effect on program semantics and control flow.

We consider *exnization* and *stackification* transformations. Exnization implements the impact of exceptions on effect representations. Stackification is useful in a stack-based safety context— that is, where events associated with function activations are “forgotten” when that activation returns, as in e.g. Java stack inspection. Thus, security contexts are established by the current calling context.

7.1. A TRANSFORMATION FOR EXCEPTIONS

Exceptions exist in Java, so a realistic application of our approach must account for them. In this section we consider a first approximation of the full exception feature set, where we assume there exists only one anonymous exception in the language. The FJ_{trace} language of expressions is extended to include the form \mathbf{throw} for throwing this anonymous exception, and to include the form $\mathbf{try}\{e_1\}\mathbf{catch}\{e_2\}$ for handling thrown exceptions, yielding the language FJ_{exn} . The semantics of FJ_{exn} are the semantics of FJ_{trace} extended with the rules in Fig. 17, where evaluation contexts \mathbf{E} are as specified in Definition 5.5.

$$\begin{array}{c}
\text{T-THROW} \\
\Gamma, \mathbf{true}, \mathbf{throw} \vdash \mathbf{throw} : [\mathbf{TC}] \\
\\
\text{T-TRYCATCH} \\
\frac{\Gamma, C, H_1 \vdash e_1 : [\mathbf{TC}] \quad \Gamma, C, H_2 \vdash e_2 : [\mathbf{TC}]}{\Gamma, C, H_1 \dot{\vdash} H_2 \vdash \mathbf{try}\{e_1\}\mathbf{catch}\{e_2\} : [\mathbf{TC}]}
\end{array}$$

Figure 18. Logical typing rules for exceptions

7.1.1. Logical Type System

To treat these new language forms in type analysis, we introduce two new forms to the language of effects: **throw** to identify control flow points where an exception is thrown, and $H_1 \dot{\vdash} H_2$ to represent the effect of handlers, where H_1 represents the effect of the **try** clause, and H_2 represents the effect of the **catch** clause. We endow these forms with an LTS semantics appropriate to the behavior of exceptions, as follows:

$$\begin{array}{c}
\mathbf{throw}; H \xrightarrow{\epsilon} \mathbf{throw} \quad \mathbf{throw} \dot{\vdash} H \xrightarrow{\epsilon} H \quad \epsilon \dot{\vdash} H \xrightarrow{\epsilon} \epsilon \\
\\
H_1 \dot{\vdash} H_2 \xrightarrow{a} H'_1 \dot{\vdash} H_2 \quad \text{if } H_1 \xrightarrow{a} H'_1
\end{array}$$

The transition rules ensure that for any sequence of events $H_1; H_2$, if H_1 encounters a **throw**, then none of the events in H_2 are reached. Also, if H_1 in the effect $H_1 \dot{\vdash} H_2$ of some handler encounters a **throw** event, then the transition rules ensure that the effect H_2 of the handling clause is encountered. Otherwise, if H_1 transitions safely to ϵ without encountering a throw then the events in H_2 are never encountered.

We also redefine the interpretation of trace effects as follows, to account for the fact that effect computation may be terminated by an uncaught exception:

$$\llbracket H \rrbracket_{\text{exn}} = \frac{\{a_1 \cdots a_n \mid H \xrightarrow{a_1} \cdots \xrightarrow{a_n} H'\}}{\cup} \{a_1 \cdots a_n \downarrow \mid H \xrightarrow{a_1} \cdots \xrightarrow{a_n} H' \text{ where } H' \in \{\epsilon, \mathbf{throw}\}\}$$

We refer to this interpretation as the *exn interpretation* of effects, as opposed to the previous *exn-free interpretation* of effects. The definition of primitive subtyping is modified to accommodate this interpretation of effects:

$$\frac{\llbracket H \rrbracket_{\text{exn}} \subseteq \llbracket H' \rrbracket_{\text{exn}}}{H \preceq_{\text{fin}} H'}$$

Finally, we extend the logical type system to account for exceptions and handlers, as in Fig. 18.

$$\begin{aligned}
XZ f \epsilon &= \{\epsilon\}, \emptyset \\
XZ f \text{ ev}[i] &= \{\text{ev}[i]\}, \emptyset \\
XZ f \text{ throw} &= \emptyset, \{\epsilon\} \\
XZ f H_1 | H_2 &= \text{let } s_1, t_1 = f H_1 \text{ in} \\
&\quad \text{let } s_2, t_2 = f H_2 \text{ in} \\
&\quad s_1 \cup s_2, t_1 \cup t_2 \\
XZ f H_1 ; H_2 &= \text{let } s_1, t_1 = f H_1 \text{ in} \\
&\quad \text{let } s_2, t_2 = f H_2 \text{ in} \\
&\quad s_1 ; s_2, t_1 \cup (s_1 ; t_2) \\
XZ f H_1 \uparrow H_2 &= \text{let } s_1, t_1 = f H_1 \text{ in} \\
&\quad \text{let } s_2, t_2 = f H_2 \text{ in} \\
&\quad s_1 \cup (t_1 ; s_2), t_1 ; t_2 \\
XZ f \mu h. H &= f H[\mu h. H/h]
\end{aligned}$$

Figure 19. The XZ combinator

While the exn interpretation of effects is a sufficient and appealing approach to obtaining type safety, it has the drawback that throw and $H_1 \uparrow H_2$ forms are non-standard and not treated by existing model-checking techniques. Our solution to this is to transform FJ_{exn} effects into FJ_{trace} effects, such that the latter conservatively approximate the former. In particular, we are interested in a class of functions we call *exn transformers*, that take an effect H and return a pair of sets s, t of exn -free trace effects, where s generates the *safe* traces in $\llbracket H \rrbracket$ that are not terminated by a raised exception, and t generates the remaining *throw* traces in $\llbracket H \rrbracket$. To manipulate these sets in subsequent development, we define the following notation:

$$s_1 ; s_2 = \{H_1 ; H_2 \mid H_1 \in s_1 \text{ and } H_2 \in s_2\} \quad s[H/h] = \{H'[H/h] \mid H' \in s\}$$

$$\begin{aligned}
\text{join}(\{H\}) &= H \\
\text{join}(\{H\} \cup s) &= H \text{ join } (s)
\end{aligned}$$

We equate sets of effects up to their join interpretation, i.e. $s = t$ iff $\text{join}(s) = \text{join}(t)$. Formally, we define exn transformers as follows:

DEFINITION 7.1. *An exn transformer is a total function f on trace effects such that for all H , $f(H)$ is a pair of trace effect sets (s, t) such that $\llbracket H \rrbracket_{\text{exn}} \subseteq \llbracket \text{join}(s \cup t) \rrbracket$.*

We precisely characterize the behavior of exn transformers via the combinator XZ defined in Fig. 19, clearly illustrating the relation between input and output of transformation of closed effects. For example,

T-THROW $\Gamma, \mathbf{true}, \mathbf{throw} \vdash_W \mathbf{throw} : \mathbf{t}$
T-TRYCATCH $\frac{\Gamma, C, \mathbf{H}_1 \vdash_W e_1 : [\mathbf{T}C] \quad \Gamma, D, \mathbf{H}_2 \vdash_W e_2 : [\mathbf{S}C]}{\Gamma, C \wedge D \wedge \mathbf{S} <: \mathbf{t} \wedge \mathbf{T} <: \mathbf{t}, \mathbf{H}_1 \uparrow \mathbf{H}_2 \vdash_W \mathbf{try}\{e_1\}\mathbf{catch}\{e_2\} : [\mathbf{t}C]}$

Figure 20. Type inference rules for exceptions

all the throw paths of \mathbf{H}_1 are prefixes of the safe and throw paths of \mathbf{H}_2 in the transform interpretation of $\mathbf{H}_1 \uparrow \mathbf{H}_2$, reflecting its *exn* interpretation. An important consequence of the definition of XZ and the *exn* interpretation of effects is:

LEMMA 7.1. *Any fixpoint of XZ is an *exn* transformer.*

It would be possible to show that XZ is monotonic, guaranteeing the existence of an *exn* transformer. However, we will instead define an algorithm that is shown to be a fixpoint of XZ , providing a technique for applying standard model checking techniques to FJ_{exn} type safety enforcement.

7.1.2. Type Inference System

To obtain type inference for FJ_{exn} , we extend the FJ_{trace} type inference system with the rules specified in Fig. 20 for language forms related to exceptions. Also necessary is an extension of *hextract* to accommodate new effect forms:

$$\begin{aligned} \text{hextract}(\mathbf{throw}, hs) &= \mathbf{throw} \\ \text{hextract}(\mathbf{H}_1 \uparrow \mathbf{H}_2, hs) &= \text{hextract}(\mathbf{H}_1, \mathbf{hs}) \uparrow \text{hextract}(\mathbf{H}_2, \mathbf{hs}) \end{aligned}$$

This completes the logical specification of the system, for which we can demonstrate a subject reduction result as in Lemma 7.3 below. But to apply standard model checking techniques to FJ_{exn} as discussed above, we now define an *exn* transformer. The core of it is a function called *exnize*, defined in Fig. 21. Since termination in the algorithm precludes unrolling of μ -bound effects, it must deal with free variables. Thus, in addition to returning safe and throw paths, *exnize* returns a set of effects called *precursors*. Precursors are safe paths that end in a variable (tail recursive paths); they may be “pruned back” from other paths (safe, throw, or recursive), since if a recursive call causes a throw, everything after the call will be short-circuited until the first enclosing handler. The idea is that precursors should be joined with throw paths

within a μ -binding, since any number of recursive calls can be made before a throw is encountered; this yields μ -bound throw paths. For example, given:

$$\mathbf{H} \triangleq \mu \mathbf{h}. \mathbf{ev}[1] | (\mathbf{ev}[2]; \mathbf{h}; \mathbf{ev}[3]) | (\mathbf{ev}[4]; \mathbf{throw}; \mathbf{ev}[5])$$

we yield the following safe and throw paths obtained from \mathbf{H} :

$$\mathit{safe}: \mu \mathbf{h}. \mathbf{ev}[1] | (\mathbf{ev}[2]; \mathbf{h}; \mathbf{ev}[3]) \quad \mathit{throw}: \mu \mathbf{h}. \mathbf{ev}[4] | (\mathbf{ev}[2]; \mathbf{h})$$

An added complication is that recursive calls \mathbf{h} may precede other tail recursions or throws in recursive or throw paths; the analysis replaces these “inner” recursions with the safe recursive call paths of \mathbf{h} , to obtain safe preceding ground paths.

In more detail, the algorithm *exnize*, defined in Fig. 21, returns a triple s, t, r , where s are the safe paths, t are the throw paths, and r are the precursors, each represented as history effect sets. The exception transformation is defined via some auxiliary functions, including *map* and *filter* defined as usual (where the latter accepts only values that match the given predicate). We also adapt a pattern-matching syntax, so that functions $(\lambda(\mathbf{H}; \mathbf{h}). \dots)$ match effect arguments of the form $\mathbf{H}; \mathbf{h}$; we observe that precursors are guaranteed to be of the this form, by definition of *exnize* and type inference.

At the top-level, the ultimate transformation is the join of the deduced safe and throw paths; note also that at the top-level, the set of precursors should be empty. Thus, the exception transformation of an effect \mathbf{H} is implemented as:

$$\text{let } s, t, \emptyset = \mathit{exnize}(\mathbf{H}) \text{ in } \mathit{join}(s \cup t)$$

For brevity, we have excluded a special subcase of $\mathit{exnize}(\mu \mathbf{h}. \mathbf{H})$, where the recursive call $\mathit{exnize}(\mathbf{H})$ returns s, t, r such that $s = \emptyset$. However, this is the case where every program control path through a function throws an exception, which we believe will be rare, and can be easily dealt with by modifying the case where s is not empty.

7.1.3. Properties

We now establish relevant results for the theory developed in this section, in particular we demonstrate type soundness, and show that the *exnize* transformation obtains a sound approximation of the *exn* interpretation in the *exn*-free interpretation of effects.

Type soundness is proved by extending Lemma 5.3, in light of the definitions above. In order to prove the R-THROW case, we observe that the effect of an expression $\mathbf{E}[\mathbf{throw}]$ will reflect that \mathbf{throw} is the “first thing that happens”.

$$\begin{aligned}
\text{exnize}(\epsilon) &= \{\epsilon\}, \emptyset, \emptyset \\
\text{exnize}(\mathbf{ev}[\mathbf{i}]) &= \{\mathbf{ev}[\mathbf{i}]\}, \emptyset, \emptyset \\
\text{exnize}(\mathbf{throw}) &= \emptyset, \{\epsilon\}, \emptyset \\
\text{exnize}(\mathbf{h}) &= \{\mathbf{h}\}, \emptyset, \{\mathbf{h}\} \\
\text{exnize}(\mathbf{H}_1|\mathbf{H}_2) &= \text{let } s_1, t_1, r_1 = \text{exnize}(\mathbf{H}_1) \text{ in} \\
&\quad \text{let } s_2, t_2, r_2 = \text{exnize}(\mathbf{H}_2) \text{ in} \\
&\quad s_1 \cup s_2, t_1 \cup t_2, r_1 \cup r_2 \\
\text{exnize}(\mathbf{H}_1;\mathbf{H}_2) &= \text{let } s_1, t_1, r_1 = \text{exnize}(\mathbf{H}_1) \text{ in} \\
&\quad \text{let } s_2, t_2, r_2 = \text{exnize}(\mathbf{H}_2) \text{ in} \\
&\quad s_1; s_2, t_1 \cup (s_1; t_2), r_1 \cup (s_1; r_2) \\
\text{exnize}(\mathbf{H}_1 \uparrow \mathbf{H}_2) &= \text{let } s_1, t_1, r_1 = \text{exnize}(\mathbf{H}_1) \text{ in} \\
&\quad \text{let } s_2, t_2, r_2 = \text{exnize}(\mathbf{H}_2) \text{ in} \\
&\quad s_1 \cup (t_1; s_2), t_1; t_2, r_1 \cup (t_1; r_2) \\
\text{exnize}(\mu\mathbf{h}.\mathbf{H}) &= \text{let } s, t, r = \text{exnize}(\mathbf{H}) \text{ in} \\
&\quad \text{let } \mathbf{H}_s = \mu\mathbf{h}.\text{join}(s) \text{ in} \\
&\quad \text{let } r' = \text{map } (\lambda(\mathbf{H}; \mathbf{h}'). \mathbf{H}[\mathbf{H}_s/\mathbf{h}]; \mathbf{h}') r \text{ in} \\
&\quad \text{let } r_{\mathbf{h}} = \text{filter } (\lambda(\mathbf{H}; \mathbf{h}'). \mathbf{h}' = \mathbf{h}) r' \text{ in} \\
&\quad \text{let } t' = \text{map } (\lambda\mathbf{H}. \mu\mathbf{h}.\text{join}(\{\mathbf{H}[\mathbf{H}_s/\mathbf{h}]\} \cup r_{\mathbf{h}})) t \text{ in} \\
&\quad \{\mathbf{H}_s\}, t', r' - r_{\mathbf{h}}
\end{aligned}$$

Figure 21. The exception transformation function *exnize*

LEMMA 7.2. *If $\Gamma, C, \mathbf{H} \vdash \mathbf{E}[\mathbf{throw}] : \mathbf{T}$ then $C \Vdash \mathbf{throw}; \mathbf{H}' <: \mathbf{H}$.*

This Lemma is proved in the identical manner as Lemma 5.19, since evaluation context forms are conserved. Now, we can easily prove subject reduction for FJ_{exn} .

LEMMA 7.3 (FJ_{exn} Subject Reduction). *If $\Gamma, C, \mathbf{H} \vdash \mathbf{e} : [\mathbf{TC}]$ is derivable for closed e and well-formed Γ , and $\eta, \mathbf{e} \rightarrow \eta', \mathbf{e}'$, then $\Gamma, C, \mathbf{H}' \vdash \mathbf{e}' : [\mathbf{TC}]$ is derivable with $C \Vdash \eta'; \mathbf{H}' <: \eta; \mathbf{H}$.*

Proof. Since the semantics of FJ_{exn} are conservative with respect to the semantics of FJ_{trace} , Lemma 5.3 implies the result for the exn-free subset of FJ_{exn} . What remains is to prove the RC-TRY, R-TRY, and R-CATCH, and R-THROW cases. The first three follow in a straightforward manner by definition of logical typing for exception handlers, the interpretation of effects $\mathbf{H}_1 \uparrow \mathbf{H}_2$, and the induction hypothesis. The interesting case is R-THROW, which goes as follows. We have $\mathbf{e} = \mathbf{E}[\mathbf{throw}]$, $\mathbf{e}' = \mathbf{throw}$, and $\eta = \eta'$, and by Lemma 7.2 it is the case that $C \Vdash \mathbf{throw}; \mathbf{H}'' <: \mathbf{H}$ for some \mathbf{H}'' . But $\Gamma, C, \mathbf{throw} \vdash \mathbf{throw} : [\mathbf{TC}]$ by T-THROW, and it is easy to

show that $C \Vdash \mathbf{throw} <: \mathbf{throw}; \mathbf{H}''$, given the LTS semantics of \mathbf{throw} , therefore $C \Vdash \eta; \mathbf{throw} <: \eta; \mathbf{throw}; \mathbf{H}''$ by Lemma 4.1. \square

Finally, we establish that *exnize* provides an exn transformer, via the following Lemma.

LEMMA 7.4. *($\lambda x.let\ s, t, \emptyset = exnize(x)$ in s, t) is an exn transformer.*

We just need to show that the function is a fixpoint of XZ . Clearly, the main issue is to prove that *exnize* deals with variables and the $\mu\mathbf{h.H}$ case correctly. This is essentially established by the next auxiliary Lemma.

LEMMA 7.5. *Let all identifiers be as defined in the $\mu\mathbf{h.H}$ case of *exnize*, and let $s_{\circ}, r_{\circ}, t_{\circ} = exnize(H[\mu\mathbf{h.H}/\mathbf{h}])$. Then:*

1. $s_{\circ} = s[\mathbf{H}_s/\mathbf{h}]$
2. $r_{\circ} = r' - r_{\mathbf{h}}$
3. $t_{\circ} = (map\ (\lambda\mathbf{H.H}[\mathbf{H}_s/\mathbf{h}])\ t) \cup (\cup(map\ (\lambda(\mathbf{H};\mathbf{h}).\{\mathbf{H}\}; t')\ r_{\mathbf{h}}))$
4. $join(s_{\circ}) = \mathbf{H}_s$
5. $join(t_{\circ}) = join(t')$

Proof. (Sketch). Each property follows by a straightforward induction on \mathbf{H} , and (intuitively) since the computation of $exnize(\mathbf{H}[\mu\mathbf{h.H}/\mathbf{h}])$ will encounter $\mu\mathbf{h.H}$ in recursive calls, rather than \mathbf{h} , hence the latter will be replaced with appropriate parts of $exnize(\mu\mathbf{h.H})$ in the transformation. In essence, property (1) follows since the safe paths of $exnize(\mathbf{H}[\mu\mathbf{h.H}/\mathbf{h}])$ will be the same as the safe paths of $exnize(\mathbf{H})$, except with every instance of \mathbf{h} replaced with the safe paths of $exnize(\mu\mathbf{h.H})$ — that is, \mathbf{H}_s . Property (2) follows since the only precursors in r' not also returned by $exnize(\mathbf{H}[\mu\mathbf{h.H}/\mathbf{h}])$ will be those terminated by \mathbf{h} — that is, $r_{\mathbf{h}}$ — since \mathbf{h} is no longer free in $\mathbf{H}[\mu\mathbf{h.H}/\mathbf{h}]$. Property (3) follows by definition of the $\mathbf{H}_1; \mathbf{H}_2$ case of *exnize*, and since $exnize(\mathbf{H}[\mu\mathbf{h.H}/\mathbf{h}])$ will encounter the same \mathbf{throw} terminated paths as $exnize(\mathbf{H})$, except with the safe paths of $exnize(\mu\mathbf{h.H})$ substituted for internal occurrences of \mathbf{h} , since free \mathbf{h} encountered by $exnize(\mathbf{H})$ will instead be occurrences of $\mu\mathbf{h.H}$. For the same reason, the \mathbf{h} -terminated precursors of $exnize(\mathbf{H})$ will be treated as safe prefixes of throw paths of $exnize(\mu\mathbf{h.H})$ in $exnize(\mathbf{H}[\mu\mathbf{h.H}/\mathbf{h}])$. Properties (4) and (5) follow immediately by (1) and (3) and properties of effect equivalence as described in Lemma 4.1. \square

Now we can establish the precondition of the main result for *exnize*.

LEMMA 7.6. *($\lambda x.let\ s, t, \emptyset = exnize(x)$ in s, t) is a fixpoint of XZ .*

$\frac{\text{R-INVK} \quad mbody(m, C) = \bar{x}.e}{\varsigma, (\text{new } C(\bar{v})).m(\bar{u}) \rightarrow \varsigma :: \epsilon, \cdot[\bar{u}/\bar{x}, \text{new } C(\bar{v})/\text{this}]e.}$	
$\text{R-EVENT} \quad \varsigma :: \eta, \text{ev}[i] \rightarrow \varsigma :: \eta; \text{ev}[i], ()$	$\text{R-POP} \quad \varsigma :: \eta, E[\cdot v \cdot] \rightarrow \varsigma, E[v]$

Figure 22. The stack-based semantics of FJ_{stack}

Proof. All cases are trivial except the $\mu h.H$ case, where it remains to be shown that $exnize(\mu h.H) = exnize(H[\mu h.H])$ by definition of $XZ-$ but this case follows immediately by Lemma 7.5 (2), (4) and (5). \square

7.2. A TRANSFORMATION FOR STACK-BASED POLICIES

Rather than consistently accruing events in a trace, a stack-based model can be defined where events generated by method activations are associated with the activation call-stack frame; when the activation is popped, so are the associated events. This allows security decisions to be made with respect to events in the current calling context. The Java stack inspection (Gong et al., 1997) access control mechanism, for example, is based on sequences of events on the call stack. A stack-based access control model has additionally been combined with a history-based security mechanism in (Abadi and Fournet, 2003), and a static analysis for enforcing general stack-based properties via temporal logic is presented in (Besson et al., 2001). Direct type inference for a stack-based security model has been studied previously, e.g. in (Skalka and Smith, 2000); however, since security mechanisms such as that proposed in (Abadi and Fournet, 2003) require both a history- and stack-based perspective, we believe that our uniform approach is simpler than e.g. combining direct stack- and history-based inference in such a context.

In this section, we observe that a stack-based security model can be statically enforced, not by redefining the inference system discussed in Sect. 6, but by an effect post-processing technique called *stackification*. Stackification takes as input an effect that predicts the trace generated by a program, and returns the stack contexts generated by a program. The stack contexts generated by a program are formalized by refiguring the FJ_{trace} operational semantics with regard to stacks, rather than histories, yielding the language model we call FJ_{stack} . Stack contexts maintain a notion of ordering; hence stacks, which we denote ς , are LIFO sequences of histories, and are either **nil** or constructed with a

$$\begin{aligned}
\mathit{stackify}(\epsilon) &= \epsilon \\
\mathit{stackify}(\epsilon; \mathbf{H}) &= \mathit{stackify}(\mathbf{H}) \\
\mathit{stackify}(\mathbf{ev}[i]; \mathbf{H}) &= \mathbf{ev}[i]; \mathit{stackify}(\mathbf{H}) \\
\mathit{stackify}(\mathbf{h}; \mathbf{H}) &= \mathbf{h} | \mathit{stackify}(\mathbf{H}) \\
\mathit{stackify}((\mu \mathbf{h}. \mathbf{H}_1); \mathbf{H}_2) &= (\mu \mathbf{h}. \mathit{stackify}(\mathbf{H}_1)) | \mathit{stackify}(\mathbf{H}_2) \\
\mathit{stackify}((\mathbf{H}_1 | \mathbf{H}_2); \mathbf{H}) &= \mathit{stackify}(\mathbf{H}_1; \mathbf{H}) | \mathit{stackify}(\mathbf{H}_2; \mathbf{H}) \\
\mathit{stackify}((\mathbf{H}_1; \mathbf{H}_2); \mathbf{H}_3) &= \mathit{stackify}(\mathbf{H}_1; (\mathbf{H}_2; \mathbf{H}_3)) \\
\mathit{stackify}(\mathbf{H}) &= \mathit{stackify}(\mathbf{H}; \epsilon)
\end{aligned}$$

Figure 23. The *stackify* algorithm

cons operator ($::$):

$$\varsigma ::= \mathbf{nil} \mid \varsigma :: \eta \quad \text{trace stacks}$$

The FJ_{stack} source language is identical to FJ_{trace} , except that so-called framed expressions $\cdot \mathbf{e} \cdot$ are included, to delimit regions of code associated with a stack frame. Stack frames are associated with activations in keeping with the standard stack-based model. We also extend evaluation contexts (Definition 5.5) with the form $\cdot \mathbf{E} \cdot$. Thus, in the operational semantics defined in Fig. 22, the rule governing method invocation, R-INVK, will push a new frame on the stack, and delimit the code region associated with the activation. Frames are popped, as in R-POP, when activations return the result of evaluation. Events are accrued in order within an activation frame, as in R-EVENT.

Given this model, a static analysis in FJ_{stack} will approximate the stack contexts that can be generated during program execution. As mentioned above, we accomplish this by a *stackify* transformation, which takes as input trace effects output by FJ_{trace} type inference, which is also applicable to FJ_{stack} source programs (framed expressions are only generated at run-time). A phenomenon exploited by our transformation is that the scope of inferred method effects is always delimited by a μ -binding. This is because the *hextract* algorithm will resolve any trace effect variable h as a μ -bound effect, and every method is assigned a variable h as its effect during inference. In other words, stack “pushes” and “pops” are implicitly recorded during inference as the beginning and end of μ -scope.

This means that *stackify*, defined in Fig. 23, can use the syntax of effects to recognize corresponding pushes and pops. Note that in the transformation of μ -bound effects, any effects \mathbf{H}_2 following a μ -bound effect \mathbf{H}_1 will be considered as part of a different stack context, since \mathbf{H}_1

is associated with an activation that will be pushed and popped before any events predicted by H_2 can occur.

The *stackify* algorithm generally exploits a normal form representation of effects as a sequence $H_1; H_2$. The last three clauses use trace effect equalities to massage trace effects into this normal form. Observe that the range of *stackify* consists of trace effects that are all tail-recursive; stacks are therefore finite-state transition systems and more efficient model-checking algorithms are possible for stacks than for general histories (Esparza et al., 2001).

EXAMPLE 7.1. *With a , b and c representing arbitrary events, and results of stackification simplified via effect equivalences to increase readability:*

$$\begin{aligned} \text{stackify}(a; b) &= a; b & \text{stackify}(a; (\mu h. b)) &= a; b \\ \text{stackify}((\mu h. a); b) &= a|b & \text{stackify}(\mu h. a|(b; h)) &= (\mu h. a|(b; h))|\epsilon \\ \text{stackify}(\mu h. a|(h; b)) &= (\mu h. a|b|h)|\epsilon \end{aligned}$$

In the second example, since $\mu h. a$ precedes b , but $\mu h. a$ denotes the effect of a function call, stackification specifies that no events precede b since a will be popped before encountering b . In the last example, since b is preceded by h , which represents recursive μ -scope and hence a recursive call, any events preceding b will be popped before b is encountered, hence stackification specifies that no events precede b . The ϵ in the last two examples is an artifact of the transformation, that could be cleaned up with some minor alterations to *stackify*. Stackification is thoroughly discussed, and correctness of the approach rigorously established, in application to a core functional calculus in (Skalka et al., 2006). Those results can easily be adapted to the Object Oriented model presented here.

8. Conclusion

In this section we conclude with a discussion of related work and some final remarks.

8.1. RELATED WORK

Previous work relevant to the application of trace based security models has been noted in Sect. 1. A number of different systems have been developed to enforce trace based, or temporal, properties of program execution. Perhaps the principal division between them is run-time

(Schneider, 2000; Abadi and Fournet, 2003) *vs.* compile-time (Ball and Rajamani, 2000; Chen and Wagner, 2002; Besson et al., 2001; Schmidt, 1998) verification. The focus of this paper is on the latter, which have in common the idea of extracting an abstract interpretation of some form from a program and verifying properties of that abstraction. The MOPS system (Chen and Wagner, 2002) compiles C programs to Push-down Automata (PDAs) reflecting the program control flow, where transitions are program transitions and the automaton stack abstracts the program call stack. (Jensen et al., 1999; Besson et al., 2001) assume that some (undefined) algorithm has already converted a program to a control flow graph, expressed as a form of PDA.

These aforementioned abstractions work well for procedural programs, but are not powerful enough to fully address advanced language features such as higher order functions and objects. Our type and effect (Talpin and Jouvelot, 1992; Amtoft et al., 1999) approach, on the other hand, allows abstract interpretation of higher order (Skalka and Smith, 2004) and Object Oriented programs. Trace effects yielded by the analysis provide a conservative approximation of trace behavior via an LTS (Labelled Transition System) interpretation. This allows the expression of program assertions as temporal logical formulae and the automated verification of assertions via model-checking techniques (Steffen and Burkart, 1992).

Some of the aforementioned systems also automatically verify assertions at compile-time via model-checking, including (Ball and Rajamani, 2000; Chen and Wagner, 2002; Besson et al., 2001), though none of these define a rigorous process for extracting an LTS from higher order or Object Oriented programs. In these works, the specifications are temporal logics, regular languages, or finite automata, and the abstract control flow is extracted as an LTS in the form of a finite automaton, grammar, or PDA. These particular formats are chosen because these combinations of logics and abstract interpretations can be automatically model-checked.

Security automata (Schneider, 2000) use finite automata for the specification and run-time enforcement of language safety properties. Systems have also been developed for statically verifying correctness of security automata using dependent types (Walker, 2000) and in a more general form as refinement types (Mandelbaum et al., 2003). These systems do not extract any abstract interpretations, so are in a somewhat different category than the aforementioned (and our) work.

Logical assertions can be local, concerning a particular program point, or global, defining the whole behavior required. However, access control systems (Wallach and Felten, 1998; Abadi and Fournet, 2003; Edjlali et al., 1998), use local checks. Since we are interested in

the static enforcement of access control mechanisms, the focus in this paper is on local, compile-time checkable assertions, though in principle the verification of global properties is possible in our system. Related work has also modified our basic approach to enforce “policy framings” that support so-called local liveness properties (Bartoletti et al., 2005b).

Perhaps the most closely related work is (K. Marriott and Sulzmann, 2003), which proposes a similar type and effect system and type inference algorithm, but their “resource usage” abstraction is of a markedly different character, based on grammars rather than LTSs. Their system lacks parametric polymorphism, which restricts expressiveness in practice, and verifies global, rather than local, assertions. Furthermore, their system analyzes only history-based properties, not stack-based properties as in our system. The system of (Igarashi and Kobayashi, 2002) is based on linear types, not effect types. Their usages U are similar to our history effects H , but the usages have a much more complex grammar and appear to have no real gain in expressiveness. Their specification logic is left abstract, thus they provide no automated mechanism for expressing or deciding assertions. Also, both of these systems are developed in functional, not Object Oriented, language models.

The systems in (Colcombet and Fradet, 2000; Besson et al., 2002; Jensen et al., 1999; Besson et al., 2001) use LTSs extracted from control-flow graph abstractions to model-check program security properties expressed in temporal logic. Their approach is close in several respects, but we are primarily focused on the programming language as opposed to the model-checking side of the problem. Their analyses assume the pre-existence of a control-flow graph abstraction, which is in the format for a first-order program analysis only. Our type-based approach is defined directly at the language level, and type inference provides an explicit, scalable mechanism for extracting an abstract program interpretation, which is applicable to Object Oriented features. Furthermore, polymorphic effects are inferable in our system and events may be parameterized by constants so partial dataflow information can be included. We believe our results are critical to bringing this general approach to practical fruition for production programming languages such as ML and Java (Skalka et al., 2005; Skalka, 2005).

Another important related work is (Higuchi and Ohori, 2006), where a type system for static enforcement of stack inspection is developed for Java bytecode. In particular, they address issues related to dynamic dispatch and linking that are similar to those discussed in Sect. 2, although our system treats a more general class of program properties than stack inspection. Also, their system essentially relies on the “join of all effects” approach discussed in Sect. 2 for dynamically dispatched

method typings, whereas we develop a more accurate solution based on parametric polymorphism.

Some recent work (Logozzo, 2004; Alur et al., 2005) has focused on analyzing patterns of method invocations for model checking safety properties of Object Oriented programs, although traces in these works are represented as regular expressions, not LTSs. While there are some similarities in their approaches and applications, we believe that our system is the first to consider the extension of trace effects *per se* to Object Oriented programs, as a technique for statically verifying assertions in a general event trace program logic.

8.2. SUMMARY

We have defined the language FJ_{trace} , a version of Featherweight Java (FJ) extended with event traces and checks, which are local assertions imposing well-formedness properties on traces. This provides a foundation for a general logic of trace based program properties in Object Oriented languages such as Java. We have defined a static type and effect analysis that automatically generates conservative approximations of FJ_{trace} program trace behavior, called trace effects. Trace effects are endowed with a label transition system semantics, and are therefore amenable to model checking for static verification of asserted trace based properties. The analysis is sound, in that static verification of program trace effects guarantees success of dynamic checks.

The Object Oriented paradigm presents several challenges to trace effect analysis, including complications due to inheritance, method override, and dynamic dispatch. In particular, we have observed that different versions of methods in a given class hierarchy should not be required to agree in their trace effects, since this requirement would be overly restrictive. We have proposed a particular application of parametric polymorphism to promote flexibility in the presence of dynamic dispatch. We have also shown that a novel definition of subtyping constraints in a regular tree model can be used for flexibility in application to Object Oriented program features, including dynamically checked downcasts.

In addition to a basic Object Oriented model based on FJ, we have considered extensions including exceptions and stack-based security contexts. Transformations of inferred trace effects were defined, that were demonstrated to faithfully reflect the behavior of these extensions. This provides additional evidence that trace effects are scalable static representations of program trace behavior, well suited to a general purpose Object Oriented model.

References

- Abadi, M. and C. Fournet: 2003, 'Access Control Based on Execution History'. In: *Proceedings of the 10th Annual Network and Distributed System Security Symposium (NDSS'03)*.
- Alur, R., P. Cerny, P. Madhusudan, and W. Nam: 2005, 'Synthesis of interface specifications for Java classes'. In: *POPL '05: Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of programming languages*. pp. 98–109, ACM Press.
- Amtoft, T., F. Nielson, and H. R. Nielson: 1999, *Type and Effect Systems*. Imperial College Press.
- Ball, T. and S. K. Rajamani: 2000, 'Bebop: A Symbolic Model Checker for Boolean Programs'. In: *SPIN*. pp. 113–130.
- Bartoletti, M., P. Degano, and G. L. Ferrari: 2005a, 'Enforcing Secure Service Composition.'. In: *CSFW*. pp. 211–223, IEEE Computer Society.
- Bartoletti, M., P. Degano, and G. L. Ferrari: 2005b, 'History-Based Access Control with Local Policies.'. In: V. Sassone (ed.): *FoSSaCS*, Vol. 3441 of *Lecture Notes in Computer Science*. pp. 316–332, Springer.
- Bartoletti, M., P. Degano, and G. L. Ferrari: 2005c, 'Policy framings for access control'. In: *WITS '05: Proceedings of the 2005 workshop on Issues in the theory of security*. pp. 5–11, ACM Press.
- Besson, F., T. de Grenier de Latour, and T. Jensen: 2002, 'Secure calling contexts for stack inspection'. In: *Proceedings of the Fourth ACM SIGPLAN Conference on Principles and Practice of Declarative Programming (PPDP'02)*. pp. 76–87, ACM Press.
- Besson, F., T. Jensen, D. L. Métayer, and T. Thorn: 2001, 'Model checking security properties of control flow graphs'. *J. Computer Security* **9**, 217–250.
- Bruce, K. B., L. Cardelli, G. Castagna, J. Eifrig, S. F. Smith, V. Trifonov, G. T. Leavens, and B. C. Pierce: 1995, 'On Binary Methods'. *Theory and Practice of Object Systems* **1**(3), 221–242.
- Burkart, O., D. Caucal, F. Moller, , and B. Steffen: 2001, 'Verification on Infinite Structures'. In: S. S. J. Bergstra, A. Pons (ed.): *Handbook on Process Algebra*. North-Holland.
- Cartwright, R. and M. Fagan: 1991, 'Soft typing'. In: *Proceedings of the ACM SIGPLAN 1991 conference on Programming language design and implementation*. pp. 278–292, ACM Press.
- Chen, H. and D. Wagner: 2002, 'MOPS: an Infrastructure for Examining Security Properties of Software'. In: *Proceedings of the 9th ACM Conference on Computer and Communications Security*. Washington, DC, pp. 235–244.
- Colcombet, T. and P. Fradet: 2000, 'Enforcing Trace Properties by Program Transformation'. In: *27th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. pp. 54–66.
- Damas, L. and R. Milner: 1982, 'Principal type-schemes for functional programs'. In: *ACM Symposium on Principles of Programming Languages (POPL)*. pp. 207–212.
- Edjlali, G., A. Acharya, and V. Chaudhary: 1998, 'History-Based Access Control for Mobile Code'. In: *ACM Conference on Computer and Communications Security*. pp. 38–48.
- Eifrig, J., S. Smith, and V. Trifonov: 1995, 'Type Inference for Recursively Constrained Types and its Application to OOP'. In: *Mathematical Foundations*

- of *Programming Semantics*, Vol. 1 of *Electronic Notes in Theoretical Computer Science*. Elsevier Science.
- Esparza, J., A. Kucera, and S. Schwoon: 2001, 'Model-Checking LTL with Regular Valuations for Pushdown Systems'. In: *TACS: 4th International Conference on Theoretical Aspects of Computer Software*.
- Foster, J. S., T. Terauchi, and A. Aiken: 2002, 'Flow-Sensitive Type Qualifiers'. In: *Proceedings of the 2002 ACM SIGPLAN Conference on Programming Language Design and Implementation*. Berlin, Germany, pp. 1–12.
- Gong, L., M. Mueller, H. Prafullchandra, and R. Schemers: 1997, 'Going Beyond the Sandbox: An Overview of the New Security Architecture in the Java Development Kit 1.2'. In: *USENIX Symposium on Internet Technologies and Systems*. Monterey, CA, pp. 103–112.
- Higuchi, T. and A. Ohori: 2006, 'A static type system for JVM access control'. *toplas*. To appear.
- Holzmann, G. J. and M. H. Smith: 2001, 'Software model checking: extracting verification models from source code'. *Software Testing, Verification & Reliability* **11**(2), 65–79.
- Igarashi, A. and N. Kobayashi: 2002, 'Resource Usage Analysis'. In: *Conference Record of POPL'02: The 29th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. Portland, Oregon, pp. 331–342.
- Igarashi, A., B. C. Pierce, and P. Wadler: 2001, 'Featherweight Java: a minimal core calculus for Java and GJ'. *ACM Trans. Program. Lang. Syst.* **23**(3), 396–450.
- Jensen, T., D. L. Métayer, and T. Thorn: 1999, 'Verification of Control Flow Based Security Properties'. In: *Proceedings of the 1999 IEEE Symposium on Security and Privacy*.
- K. Marriott, P. J. S. and M. Sulzmann: 2003, 'Resource Usage Verification'. In: *Proc. of First Asian Programming Languages Symposium, APLAS 2003*.
- Logozzo, F.: 2004, 'Separate Compositional Analysis of Class-based Object-oriented Languages'. In: *Proceedings of the 10th International Conference on Algebraic Methodology And Software Technology (AMAST'2004)*, Vol. 3116 of *Lectures Notes in Computer Science*. pp. 332–346, Springer-Verlag.
- Mandelbaum, Y., D. Walker, and R. Harper: 2003, 'An Effective Theory of Type Refinements'. In: *Proceedings of the the Eighth ACM SIGPLAN International Conference on Functional Programming (ICFP'03)*. Uppsala, Sweden.
- Palsberg, J. and P. O'Keefe: 1995, 'A type system equivalent to flow analysis'. In: *POPL '95: Proceedings of the 22nd ACM SIGPLAN-SIGACT symposium on Principles of programming languages*. New York, NY, USA, pp. 367–378, ACM Press.
- Palsberg, J. and S. Smith: 1996, 'Constrained types and their expressiveness'. *ACM Transactions on Programming Languages and Systems* **18**(5), 519–527.
- Schmidt, D. A.: 1998, 'Trace-Based Abstract Interpretation of Operational Semantics'. *Lisp and Symbolic Computation* **10**(3), 237–271.
- Schneider, F. B.: 2000, 'Enforceable security policies'. *Information and System Security* **3**(1), 30–50.
- Skalka, C.: 2005, 'Trace effects and object orientation'. In: *PPDP '05: Proceedings of the 7th ACM SIGPLAN international conference on Principles and practice of declarative programming*. New York, NY, USA, pp. 139–150, ACM Press.
- Skalka, C. and S. Smith: 2000, 'Static Enforcement of Security with Types'. In: *Proceedings of the the Fifth ACM SIGPLAN International Conference on Functional Programming (ICFP'00)*. Montréal, Canada, pp. 34–45.

- Skalka, C. and S. Smith: 2004, 'History Effects and Verification'. In: *Asian Programming Languages Symposium*. Springer.
- Skalka, C., S. Smith, and D. Van Horn: 2005, 'A Type and Effect System for Flexible Abstract Interpretation of Java'. In: *Proceedings of the ACM Workshop on Abstract Interpretation of Object Oriented Languages*.
- Skalka, C., S. Smith, and D. Van Horn: 2006, 'Types and Trace Effects for Higher Order Programs'. Technical Report CS-06-03, The University of Vermont.
- Steffen, B. and O. Burkart: 1992, 'Model Checking for Context-Free Processes'. In: *CONCUR'92, Stony Brook (NY)*, Vol. 630 of *Lecture Notes in Computer Science (LNCS)*. Heidelberg, Germany, pp. 123–137, Springer-Verlag.
- Talpin, J.-P. and P. Jouvelot: 1992, 'The Type and Effect Discipline'. In: *Seventh Annual IEEE Symposium on Logic in Computer Science, Santa Cruz, California*. Los Alamitos, California, pp. 162–173, IEEE Computer Society Press.
- Trifonov, V. and S. Smith: 1996, 'Subtyping Constrained Types'. In: *Proceedings of the Third International Static Analysis Symposium*, Vol. 1145. pp. 349–365, Springer Verlag.
- Walker, D.: 2000, 'A Type System for Expressive Security Policies'. In: *Conference Record of POPL'00: The 27th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. Boston, Massachusetts, pp. 254–267.
- Wallach, D. S. and E. Felten: 1998, 'Understanding Java Stack Inspection'. In: *Proceedings of the 1998 IEEE Symposium on Security and Privacy*.