IOWA STATE UNIVERSITY
Digital Repository

3-25-2014

Open Effects: A Hybrid Type-and-Effect System to Tackle Open World Assumption and its Application to Optimistic Concurrency

Yuheng Long
Iowa State University, csgzlong@iastate.edu

Mehdi Bagherzadeh
Iowa State University, mbagherz@iastate.edu

Hridesh Rajan
Iowa State University

Follow this and additional works at: http://lib.dr.iastate.edu/cs_techreports

Part of the Programming Languages and Compilers Commons

Open Effects: A Hybrid Type-and-Effect System to Tackle Open World Assumption and its Application to Optimistic Concurrency

Abstract

This work tackles the challenge of applying a type-and-effect system to reason about object-oriented programs with an open world assumption. An open world assumption challenges the design of a type-and-effect system when (1) all subclasses of a class are not known, and (2) an upper bound on effects of all subclasses is not available, e.g. when an effect specification is not available for that class – a common phenomenon in modern OO programs. The main problem is in the computation of the effects of a dynamically dispatched method invocation, because all possible dynamic types of its receiver are not known statically, and no static upper bound in the form of effect specification is available. Our new concept open effects solves these problems. The basic idea is to take a programmer-guided hybrid approach. Instead of using a predefined upper bound, our type-and-effect system takes the effects of a programmer-selected dynamically dispatched method call as open effects that encapsulate statically known information about the call, e.g. static type of receiver. The static part of our type-and-effect systems treats open effects as unknown, but the dynamic part of our type-and-effect systems reifies open effects. We also apply open effects to create a sound trust-but-verify type-and-effect system, to better enable concurrent execution of dynamically dispatched method invocations. If a programmer annotates the receiver of a certain method invocation as open, then the type system trusts the programmer and assigns an open effect to the method. The open effect is, optimistically, not supposed to conflict with other effects. Such optimistic assumptions are verified statically, if possible, or at runtime otherwise. Performance evaluations of an open effects-based type system for concurrency, on various benchmarks, show that it incurs negligible annotation and runtime overheads.

Keywords

type-and-effect, open effects, optimistic concurrency

Disciplines

Programming Languages and Compilers

Comments

Open Effects: A Hybrid Type-and-Effect System to Tackle Open World Assumption and its Application to Optimistic Concurrency

Yuheng Long and Mehdi Bagherzadeh and Hridesh Rajan
TR #13-04a
Initial Submission: October 15, 2013
Revised: March 25, 2014

Department of Computer Science
226 Atanasoff Hall
Iowa State University
Ames, Iowa 50011-1041, USA

Open Effects: A Hybrid Type-and-Effect System to Tackle Open World Assumption and its Application to Optimistic Concurrency

Yuheng Long and Mehdi Bagherzadeh and Hridesh Rajan

Dept. of Computer Science, Iowa State University
{csgzlong,mbagherz,hridesh}@iastate.edu

Abstract

This work tackles the challenge of applying a type-and-effect system to reason about object-oriented programs with an open world assumption. An open world assumption challenges the design of a type-and-effect system when (1) all subclasses of a class are not known, and (2) an upper bound on effects of all subclasses is not available, e.g. when an effect specification is not available for that class – a common phenomenon in modern OO programs. The main problem is in the computation of the effects of a dynamically dispatched method invocation, because all possible dynamic types of its receiver are not known statically, and no static upper bound in the form of effect specification is available. Our new concept *open effects* solves these problems. The basic idea is to take a programmer-guided hybrid approach. Instead of using a predefined upper bound, our type-and-effect system takes the effects of a programmer-selected dynamically dispatched method call as open effects that encapsulate statically known information about the call, e.g. static type of receiver. The static part of our type-and-effect systems treats open effects as unknown, but the dynamic part of our type-and-effect systems reifies open effects. We also apply open effects to create a sound trust-but-verify type-and-effect system, to better enable concurrent execution of dynamically dispatched method invocations. If a programmer annotates the receiver of a certain method invocation as open, then the type system *trusts* the programmer and assigns an open effect to the method. The open effect is, optimistically, not supposed to conflict with other effects. Such optimistic assumptions are *verified* statically, if possible, or at runtime otherwise. Performance evaluations of an open effects-based type system for concurrency, on various benchmarks, show that it incurs negligible annotation and runtime overheads.

1. Introduction

A type-and-effect system is a very important reasoning aid [23, 36]. It has been shown to help programmers in analyzing locking disciplines [4], dynamic updating mechanisms [38], checked exceptions [6, 32], detecting race conditions [12], etc. Basic idea behind a type-and-effect system is to add an encoding of computational effects into semantic objects of a language and a discipline for controlling these effects into its type system [49]. These effects describe how the state of a program will be modified by expressions in the language, e.g. a field expression may have a read or write effect to represent reading from or writing into memory [30, 49].

The open world assumption [39], which says class hierarchies can be extended even after static type checking, is an important property of modern object-oriented languages that enables modularity and reuse, and is key to creating libraries and frameworks. For example, it allows programmers to extend abstract classes defined in libraries and continue to use algorithms defined in libraries without having to type check those libraries again. However, this assumption makes the design of a type-and-effect system challenging. The main problem is in the computation of the effects of a dynamically dispatched method invocation, because all possible dynamic types of its receiver may not be known statically.

One solution is to use static effect annotations that provide an upper bound on the effects of a dynamically dispatched method [10, 25]. However, such effect annotations are not yet broadly available. Furthermore, the upper bound given by effect annotations should be broad enough, to cover the effect of all methods that could be possibly executed as the result of invocation of the dynamically dispatched method, but not too broad to lose effectiveness.

We put forth a new approach for handling an open world assumption in a type-and-effect system. Basic idea is to take a programmer-guided hybrid approach. We allow programmers to instruct the type-and-effect system to take the effects of certain dynamically dispatched method invocation as an unknown *open effect* statically (instead of using a conservative estimate), make static effect-guided decisions while assuming these unknowns, and if necessary verify those decisions when the receiver's dynamic type is known.

```
1  class Pair {
2    int fst, snd;
3    @open Op f; // Let static effects be open.
4               // If needed, defer effect-guided decisions.
5    Pair init() { fst = 1; snd = 2; this }
6    Op setOp(Op f) { this.f = f }
7    int apply() {
8      // fork{e1,e2} executes e1 and e2 concurrently if
9      // their effects don't conflict, otherwise sequentially
10     fork{
11       fst = f.op(fst),
12       snd = f.op(snd)
13     }
14   }
15 }

17 class Op {
18   int res;
19   Op init() { res = 0; this }
20   int op(int o) { 0 }
21 }
```

Prog1: a client program that uses Library

```
22 class Prefix extends Op {
23   int op(int o) { // Effects: writes res
24     res += o
25   }
26 }
27 Pair pr = new Pair().init();
28 Op pf = new Prefix().init();
29 pr.setOp(pf);
30 pr.apply()
```

Prog2: another independent client program that uses Library

```
32 class Hash extends Op {
33   int op(int o) { // Effects: None
34     int key = o;
35     // Hash computation
36     key = ..
37   }
38 }
39 Pair pr = new Pair().init();
40 Op ha = new Hash().init();
41 pr.setOp(ha);
42 pr.apply()
```

Figure 1. Library class `Pair` with open field `f` and clients *Prog1* and *Prog2* extending `Op`.

1.1 Motivating Example: Optimistic Concurrency

To motivate, consider computation of the effects of the dynamically dispatched method `op` in Figure 1 to decide whether to run two invocations of `op`, on lines 10–13, concurrently. The code contains the library classes `Pair` and `Op` which represent a pair of integers and operations carried out on pairs of elements, respectively. It also contains client programs *Prog1* and *Prog2* that extend the library class `Op` and override its method `op`, in the `Prefix` and `Hash` classes. `Prefix` computes a prefix sum in its effectful overriding of `op` with the effect of writing into the field `res`, shown as **wr**(*res*), whereas `Hash` computes a hash in its pure method `op` with no memory effects, i.e. ∅. To specify the effects of the method `op`, in a typical type-and-effect system, its effects should be broad enough to conservatively cover the effects of its overriding methods in all of its subtypes, i.e. `Prefix`, `Hash`, and possibly other *still unknown* subtypes. This results in the effect **wr**(*res*) for the dynamically dispatched method `op` which is the *union* of the effects of its overriding methods, i.e. union of **wr**(*res*) and ∅, in all of its subtypes.

For two dynamically dispatched methods, a typical type-and-effect system may disallow the concurrent execution of their invocations, because their broadly specified static effects may conflict. However, such a conflict may not actually happen at runtime, depending on the dynamic types of their receivers. To illustrate, consider the two invocations of the method `op`, on lines 10–13 of Figure 1, in the fork expression of the `apply` method. This, somewhat nontraditional, fork expression **fork**{e_1, e_2} [9] executes the expressions e_1 and e_2 concurrently if their effects do not conflict, and runs them sequentially otherwise. For a memory location, writing into the location conflicts with other reads and writes of the same location. A typical type-and-effect system would *serialize* the execution of these invocations of the method `op`, because their static effects **wr**(*res*) conflict with each other. Such serialization of these method invocations makes sense when their receiver `f` is of dynamic type `Prefix`, however, these invocations could run *concurrently*, for example when they have empty effects at runtime when their receiver `f` has the dynamic type of `Hash`. A typical type-and-effect system would miss such safe concurrency opportunities.

1.2 Open Effects for Optimistic Concurrency

We now explain the notion of *open effects* by showing its usage in creating a sound trust-but-verify type-and-effect system for optimistic concurrency. This system uses programmer's knowledge to better expose and enable safe concurrent execution in the presence of dynamically dispatched method invocations in open world object-oriented programs.

An effect system based on our work would have two kinds of effects: open and concrete effects. Concrete effects are standard. For example, a concrete effect may represent standard memory effects, which are reads and writes of memory locations [49]. An open effect represents the effects of a dynamically dispatched method invocation where the dynamic type of its receiver is not known statically, but the receiver is qualified by the programmer with an open annotation @**open**. We call such receiver references an *open reference*. An open effect is *concretized* at runtime when the dynamic type of open references is known. So, an effect system based on our work would be hybrid with two parts:

- *Static part*, that *(i)* computes the effects of the methods, one method at a time and independent of the dynamic types of the receivers of dynamically dispatched method

invocations; and *(ii)* possibly verifies the optimistic assumptions. Optimistic concurrency is enabled by this part of our type system as follows. If a programmer annotates the receiver of certain method invocations as **@open**, this static part trusts the programmer and assigns an open effect to the method invocation, which supposedly does not conflict with other effects, i.e. disjointness assumption. Such optimistic disjointness assumptions are verified statically, if enough static information is available, or at runtime otherwise.

- **Dynamic part**, that *(iii)* concretizes the statically computed open effects, on demand, and updates them by tracking open references and their values, and *(iv)* verifies, using runtime checks, optimistic assumptions made by the static part. For example, for optimistic concurrency this dynamic part verifies the disjointness assumptions that could not be verified statically.

We believe that this novel approach of using the programmer's knowledge of their program to selectively combine static and dynamic effect analyses [21, 41, 42] could help bring out the best of both worlds, and offers a complementary approach for using type-and-effect systems to reason about programs in object-oriented languages with an open world assumption.

1.3 Revisiting Motivating Example

To revisit our example, imagine that a programmer marked the field `f` of class `Pair` as open using an **@open** annotation, on line 3 of Figure 1. Doing so, the programmer is hinting the type-and-effect system that there may be parallelism opportunities when invoking dynamically dispatched methods on the receiver `f`. The static part of our system trusts the programmer and assigns the open effects **open**(*f op γ_1*) and **open**(*f op γ_2*) to the method invocations `f.op(fst)`, on line 11, and `f.op(snd)`, on line 12, respectively. The open effect **open**(*f op γ_1*) is the effect of the invocation of the dynamically dispatched method `op` on the open receiver `f` with the statically unknown effect of γ_1. The open effect **open**(*f op γ_2*) is similar. These open effects are assumed to not conflict with each other and other effects.

The static part continues by computing the effect of the expression `fst = f.op(fst)` to be writing and reading fields `fst` and `f` plus the effect of the invocation of the method `op` on the open receiver `f`, i.e. $\sigma_1 = \{\mathbf{wr}(\textit{fst}), \mathbf{rd}(f),$ **open**(*f op γ_1*)$\}$[1]. Similarly, the effects of `snd = f.op(snd)` is $\sigma_2 = \{\mathbf{wr}(\textit{snd}), \mathbf{rd}(f),$ **open**(*f op γ_2*)$\}$. These two expressions, of the fork expression on lines 10–13 could be executed concurrently if their effects σ_1 and σ_2 do not conflict, which in turn boils down to the verification of their open effects not conflicting, since $\mathbf{wr}(\textit{fst})$ and $\mathbf{wr}(\textit{snd})$ do not conflict. This could be verified statically if there is enough static information about the unknown effects of the method `op`

---

[1] Write effect of a field, e.g. $\mathbf{wr}(\textit{fst})$ covers its read effect, e.g. $\mathbf{rd}(\textit{fst})$ [9].

or otherwise dynamically. Different modular static analyses could be integrated into the open effects' type-and-effect system to boost its static analysis. In this work, we illustrate the integration of a modular alias analysis [21] as well as the integration of effect specifications [36**?**]. Our implementation integrates other static analyses such as purity analysis [41] and array effect analysis [42].

Open effects enables better exposure of safe concurrency opportunities, in the presence of dynamically dispatched method calls, as follows: for two subexpressions of a fork expression with their statically computed effects, which may contain open effects, there are three answers for the question of *do their effects conflict statically?*: yes (conflict), no (disjoint), and unknown (may or may not conflict). Using open effects and depending on the disjointness of the effects of its subexpressions, a fork expression is *soundly and statically* translated to:

(1) Yes (conflict): an unconditional sequential execution of subexpressions of the fork expression.

(2) No (disjoint): an unconditional parallel execution of subexpressions of the fork expression.

(3) Unknown (may or may not conflict): a conditional in which the unknown open effects of subexpressions are concretized and tested for conflicts. If the concretized effects conflict then run the subexpressions sequentially, i.e. *(1)*, else in parallel, i.e. *(2)*.

In Figure 1, there is not enough static information to decide if σ_1 and σ_2, especially their open effects, conflict and thus the case *(3)* above applies and the fork expression, on lines 10–13, translates to a conditional. This brings into the picture the dynamic part of our type-and-effect system that decides the disjointness of effects that could not be decided statically.

The dynamic part concretizes, or fills in, the unknown effects of open effects when the open references are known at runtime. For example, upon the execution of the expression `pr.setOp(pf)`, on line 29 of *Prog1*, the open reference `f`, of static type `Op`, in the open effects of **open**(*f op γ_1*) and **open**(*f op γ_2*) is set to object `pf`, of the dynamic type of `Prefix`. This causes these two open effects to be concretized to $\mathbf{wr}(\textit{res})$, because the method `op` of type `Prefix` has the effect of $\mathbf{wr}(\textit{res})$. With such concretization, the effects σ_1 and σ_2 conflict at runtime and thus the fork expression, and the invocations of the dynamically dispatched method `op`, lines 10–13, is sequentialized. Unlike *Prog1*, *Prog2* assigning the object `ha` of dynamic type `Hash` to the open reference `f`, via `pr.setOp(ha)` on line 41, and causes the open effects **open**(*f op γ_1*) and **open**(*f op γ_2*) to be concretized to empty set $\emptyset$, because the method `op` of type `Hash` is pure. This in turn allows the fork expression to be translated to concurrent execution of the invocations of method `op`, as their effects σ_1 and σ_2 do not conflict. *Depending*

on the dynamic type of the open field f, the fork expression could run sequentially or in parallel.

1.4 Contributions

In summary, the main contributions of this work are the following:

- Open effects and its usage in a trust-but-verify hybrid type-and-effect system to expose safe concurrency in open world concurrent programs with dynamic dispatch;

- Illustration of the integration of static alias analysis and effect specifications into open effects;

- Static semantics of open effects, in §2, and its dynamic semantics, in §3;

- Proof of soundness for open effects, in §3;

- *OpenEffectJ*, an OpenJDK prototype implementation of open effects; and

- Speedup and overhead evaluations along with the interplay of static and dynamic parts of *OpenEffectJ* in §4.

§5 compares open effects with previous work on reasoning about effects of programs in three categories of static, dynamic and hybrid techniques; and §6 concludes the paper after discussing some avenues for future work.

2. Open Effects: A Hybrid Type-and-Effect System

Our type-and-effect system has a static and a dynamic part. The static part encoded in the typing rules, *(i)* computes the effects of the methods, one method at a time and independent of dynamic dispatch; and *(ii)* verifies optimistic disjointness assumptions of open effects, using the available static information. This section presents select typing rules that form the novel basis of our static effect computation using open effects in the presence of dynamically dispatched methods. A complete list of the typing rules and their auxiliary functions can be found in our technical report [34].

2.1 Syntax

To encode open effects as a type-and-effect system, we use *OpenEffectJ*, a core expression language, shown in Figure 2, which is based on Classic Java [22]. The A-normal form syntax of *OpenEffectJ* is standard except for open annotations @**open** and the disjoint check expression $e_1\#e_2$. The disjointness expression statically checks if the effects of the expressions e_1 and e_2 are disjoint and evaluates to true if they are, and false otherwise. The disjointness expression $e_1\#e_2$ decides the disjointness of the effects of e_1 and e_2 but *does not* evaluate e_1 and e_2. Figure 2 shows fields annotated with @**open**, i.e. open fields, however, we also support open local variables and open parameters [34]. For simplicity, we assume unique field names, up to α-renaming, and no method overloading. The notations $\overline{term}$ and $[term]$ denote a finite possibly empty sequence and an optional *term*, respectively.

$$
\begin{array}{ll}
prog ::= \overline{decl}\ e \\
decl ::= \textbf{class}\ c\ \textbf{extends}\ d\ \{\ \overline{field}\ \overline{meth}\ \} \\
field ::= [\textbf{@open}]\ t\ f\ \textbf{;} \\
meth ::= t\ m\ (\ \overline{arg}\ )\ \{\ e\ \} \\
t ::= c\ |\ \textbf{int}\ |\ \textbf{bool} \\
arg ::= t\ var,\ \textbf{where}\ var \neq \textbf{this} \\
e ::= x\ |\ \textbf{null}\ |\ arg = e\ \textbf{;}\ e & \text{``Var, Null, Definition''} \\
\quad |\ x.m\ (\overline{x})\ |\ \textbf{new}\ c\ (\,) & \text{``Call, New''} \\
\quad |\ x \circ x\ |\ n\ |\ b\ |\ loc & \text{``Binary, Number, Boolean, Location''} \\
\quad |\ \textbf{if}\ x\ \textbf{then}\ e\ \textbf{else}\ e & \text{``Conditional''} \\
\quad |\ \textbf{this}.f\ |\ \textbf{this}.f = x & \text{``Get, Set Field''} \\
\quad |\ e\ \#\ e & \text{``Disjointness Check''}
\end{array}
$$

$$
\textbf{where}\quad
\begin{array}{rcl}
c & \in & \mathscr{C},\ \text{set of class names} \\
d & \in & \mathscr{C} \cup \{\textbf{Object}\} \\
f & \in & \mathscr{F},\ \text{set of field names} \\
m & \in & \mathscr{M},\ \text{set of method names} \\
n & \in & \mathscr{N},\ \text{set of natural numbers} \\
b & \in & \{true,\ false\},\ \text{boolean constants} \\
x,var & \in & \mathscr{V} \cup \{\textbf{this}\},\ \text{set of variable names} \\
\circ & \in & \{+,-,*,/\},\ \text{set of binary operations} \\
loc & \in & \mathscr{L},\ \text{set of locations}
\end{array}
$$

Figure 2. Syntax for *OpenEffectJ*.

Using the programmer's knowledge in annotating only certain receiver fields as open is to have as less overhead as possible compared to other alternatives. One alternative is to annotate types as open, however, it causes every reference of that type and its subtypes to be treated as open references which in turn could cause considerable concretization and verification overhead, especially when all references of a type have to pay the price for one reference being open. The same applies to another alternative in which every field of every object is considered open.

2.2 Type-and-Effect Attributes

Figure 3 shows *OpenEffectJ*'s type-and-effect attributes. The type of a program and its declarations are given as OK, whereas $(\overline{t} \rightarrow t, \sigma)$ **in** c, specifies the type of a method defined in class c with parameter types $\overline{t}$, return type t and a latent effect σ [49]. The latent effect of a method is the effects of the body of the method [49]. Finally, the attribute (t, σ) specifies an expression of type t with the effects σ.

$$
\begin{array}{ll}
\theta ::= \text{OK} & \text{``program/decl types''} \\
\quad |\ (\overline{t} \rightarrow t, \sigma)\ \textbf{in}\ c & \text{``method types''} \\
\quad |\ (t, \sigma) & \text{``expression types''} \\
\Pi ::= \{var_i \mapsto t_i\}_{i \in \mathbb{N}} & \text{``type environments''} \\
\sigma, \gamma ::= \emptyset\ |\ \top\ |\ \sigma \cup \sigma & \text{``effects''} \\
\quad |\ \textbf{rd}(f) & \text{``read effect''} \\
\quad |\ \textbf{wr}(f) & \text{``write effect''} \\
\quad |\ \textbf{open}(f\ m\ \gamma) & \text{``open effect''}
\end{array}
$$

Figure 3. Type-and-effect attributes for *OpenEffectJ*, based on [25, 49].

There are two kinds of effects in Figure 3: concrete effects and open effects. Concrete effects are standard read and write memory effects[2] $\mathbf{rd}(f)$ and $\mathbf{wr}(f)$,[3] that read and write a field $\mathtt{f}$, along with the empty effect $\emptyset$ and the top effect $\top$. The top effect allows read and write effects of any field [9]. An open effect $\mathbf{open}(f\ m\ \gamma)$ represents the effects of a dynamically dispatched method $\mathtt{m}$ invoked on an open receiver $\mathtt{f}$ marked with **@open**. The placeholder γ represents the unknown effect of the body of the method $\mathtt{m}$. We slightly misuse the set notation for presentation purposes.

Type checking rules use a standard implicit fixed class table *CT* which contains a list of program declarations [22]. Each method in the class table *CT* has its statically computed effects as part of its signature. The typing rules use a type environment Π, which maps a variable name *var* to its type *t*. The typing judgement $\Pi \vdash e \rightsquigarrow e' : (t, \sigma)$ says that with the typing environment Π the expression *e* is translated to the expression *e'* and has the type *t* and the effects σ. The semantic preserving translation $\rightsquigarrow$ does not change the type or the effects of expressions and in spirit is similar to elaboration in languages such as ML [37]. Subtyping is denoted using the relation $<:$ which is the standard reflexive-transitive closure of the declared subclass relationships [22] in *CT*.

2.3 Disjointness

Two expressions e_1 and e_2 can safely run in parallel, if their effects do not conflict, i.e. they are disjoint. Figure 4 shows the typing rules for the disjoint expression $e_1 \# e_2$. This expression statically checks if the effects of the expressions e_1 and e_2 conflict and depending on the answer translates into true, false or unknown.

Using open effects, the disjointness can be decided statically, provided enough static information is available, as in the rules (T-DISJOINT) and (T-CONFLICT), or otherwise it is deferred to runtime, as in (T-UNKNOWN). The availability of such static information is dependent on the static analyses integrated into the type system. First, we assume no extra static analysis to focus on the basic ideas behind open effects. Later we discuss adding a modular alias analysis and integrating effect specifications, as examples of static information that could be helpful in making some of the disjointness decisions statically.

In (T-DISJOINT), if there is no open effects in the effects of the expressions, i.e. $\downarrow_O (\sigma_1) = \downarrow_O (\sigma_2) = \emptyset$, then effects of the expressions are disjoint only if their concrete effects are disjoint, i.e. $\sigma_c^1 \mathbin{\textcircled{\#}} \sigma_c^2$. If the concrete effects of e_1 and e_2 are disjoint then $e_1 \# e_2$ statically translates to *true*, which in turn means they can run concurrently. Since $e_1 \# e_2$ does not execute any of the expressions e_1 or e_2, its effect is empty. The auxiliary functions $\downarrow_O (\sigma)$ and $\downarrow_C (\sigma)$, return the set

---

[2] Effects in our type-and-effect system are not accessible to programmers and thus exposure of the implementation details is not a concern [25].

[3] For simplicity, our formalism is not object sensitive, but our compiler implementation is both field and object sensitive [20].

(T-DISJOINT)
$$\Pi \vdash e_1 \rightsquigarrow e_1' : (t_1, \sigma_1) \qquad \Pi \vdash e_2 \rightsquigarrow e_2' : (t_2, \sigma_2)$$
$$\downarrow_O (\sigma_1) = \downarrow_O (\sigma_2) = \emptyset$$
$$\frac{\sigma_c^1 = \downarrow_C (\sigma_1) \qquad \sigma_c^2 = \downarrow_C (\sigma_2) \qquad \sigma_c^1 \mathbin{\textcircled{\#}} \sigma_c^2}{\Pi \vdash e_1 \# e_2 \rightsquigarrow true : (bool, \emptyset)}$$

(T-CONFLICT)
$$\Pi \vdash e_1 \rightsquigarrow e_1' : (t_1, \sigma_1) \qquad \Pi \vdash e_2 \rightsquigarrow e_2' : (t_2, \sigma_2)$$
$$\frac{\sigma_c^1 = \downarrow_C (\sigma_1) \qquad \sigma_c^2 = \downarrow_C (\sigma_2) \qquad !(\sigma_c^1 \mathbin{\textcircled{\#}} \sigma_c^2)}{\Pi \vdash e_1 \# e_2 \rightsquigarrow false : (bool, \emptyset)}$$

(T-UNKNOWN)
$$\Pi \vdash e_1 \rightsquigarrow e_1' : (t_1, \sigma_1) \qquad \Pi \vdash e_2 \rightsquigarrow e_2' : (t_2, \sigma_2)$$
$$\downarrow_O (\sigma_1) \neq \emptyset \ \vee \downarrow_O (\sigma_2) \neq \emptyset$$
$$\frac{\sigma_c^1 = \downarrow_C (\sigma_1) \qquad \sigma_c^2 = \downarrow_C (\sigma_2) \qquad \sigma_c^1 \mathbin{\textcircled{\#}} \sigma_c^2}{\Pi \vdash e_1 \# e_2 \rightsquigarrow e_1' \# e_2' : (bool, \emptyset)}$$

Auxiliary Functions:

$$\downarrow_C (\sigma) = \{\varepsilon | \varepsilon \in \sigma \wedge \varepsilon \in \{\mathbf{rd}(f), \mathbf{wr}(f), \top, \emptyset\}\}$$
$$\downarrow_O (\sigma) = \{\varepsilon | \varepsilon \in \sigma \wedge \varepsilon = \mathbf{open}(f\ m\ \gamma)\}$$

	(READ-WRITE)	(WRITE-WRITE)
(READ-READ)	$f \neq f'$	$f \neq f'$
$\mathbf{rd}(f) \mathbin{\textcircled{\#}} \mathbf{rd}(f')$	$\mathbf{rd}(f) \mathbin{\textcircled{\#}} \mathbf{wr}(f')$	$\mathbf{wr}(f) \mathbin{\textcircled{\#}} \mathbf{wr}(f')$

(CONFLICT)
$$\frac{\varepsilon = \mathbf{rd}(f) \vee \varepsilon = \mathbf{wr}(f)}{!(\mathbf{wr}(f) \mathbin{\textcircled{\#}} \varepsilon)}$$

(EMPTY)
$$\frac{\forall \varepsilon \in \sigma}{\varepsilon \mathbin{\textcircled{\#}} \emptyset}$$

(TOP)
$$\frac{\forall \varepsilon \in \sigma}{!(\varepsilon \mathbin{\textcircled{\#}} \top)}$$

Figure 4. Deciding disjointness of the effects of expressions e_1 and e_2.

---

of open and concrete effects of the effect set σ, respectively. The function $\textcircled{\#}$ simply checks for the disjointness of effects, in which a write and read of a field $\mathtt{f}$, i.e. $\mathbf{wr}(f)$ and $\mathbf{rd}(f)$, conflict and other effects are disjoint. The top effect $\top$ conflicts with every other effect. And ε is an effect element in the effect set σ.

Similar to (T-DISJOINT), the rule (T-CONFLICT) statically decides the disjointness of the effects of e_1 and e_2 and translates the expression $e_1 \# e_1$ to *false*, if their concrete effects conflict, i.e. $!(\sigma_c^1 \textcircled{\#} \sigma_c^2)$. In this rule there is no need to check the relation between open effects in σ_1 and σ_2 and such a check could be skipped. More importantly, the *concretization of these open effects, at runtime, could be skipped* which in turn results in less runtime checks and better performance.

Decision about disjointness in $e_1 \# e_2$ is deferred to runtime if it cannot be made statically using (T-DISJOINT) and (T-CONFLICT). The rule (T-UNKNOWN) defers such a decision by translating $e_1 \# e_2$ to $e_1' \# e_2'$. In (T-UNKNOWN), existence of open effects in either σ_1 or σ_2, i.e. $\downarrow_O (\sigma_1) \neq \emptyset \vee \downarrow_O (\sigma_2) \neq \emptyset$, prevents static decision making about dis-

(T-FORK-SEQUENTIAL)
$$\frac{\Pi \vdash e_1 \# e_2 \rightsquigarrow false : (bool, \emptyset) \qquad \Pi \vdash e_1 \rightsquigarrow e_1' : (t_1, \sigma_1) \qquad \Pi \vdash e_2 \rightsquigarrow e_2' : (t_2, \sigma_2)}{\Pi \vdash \texttt{fork}\{e_1, e_2\} \rightsquigarrow e_1'; e_2' : (t_2, \sigma_1 \cup \sigma_2)}$$

(T-FORK-PARALLEL)
$$\frac{\Pi \vdash e_1 \# e_2 \rightsquigarrow true : (bool, \emptyset) \qquad \Pi \vdash e_1 \rightsquigarrow e_1' : (t_1, \sigma_1) \qquad \Pi \vdash e_2 \rightsquigarrow e_2' : (t_2, \sigma_2)}{\Pi \vdash \texttt{fork}\{e_1, e_2\} \rightsquigarrow e_1' || e_2' : (t_2, \sigma_1 \cup \sigma_2)}$$

(T-FORK-UNKNOWN)
$$\frac{\begin{array}{c} \Pi \vdash e_1 \# e_2 \rightsquigarrow e_1' \# e_2' : (bool, \emptyset) \\ \Pi \vdash e_1 \rightsquigarrow e_1' : (t_1, \sigma_1) \qquad \Pi \vdash e_2 \rightsquigarrow e_2' : (t_2, \sigma_2) \qquad \sigma_c^1 = \downarrow_C (\sigma_I) \qquad \sigma_c^2 = \downarrow_C (\sigma_2) \qquad \sigma_o^1 = \downarrow_O (\sigma_I) \\ \sigma_o^2 = \downarrow_O (\sigma_2) \qquad cond = (concretize(\sigma_o^1) \circledast concretize(\sigma_o^2)) \wedge (concretize(\sigma_o^1) \circledast \sigma_c^2) \wedge (concretize(\sigma_o^2) \circledast \sigma_c^1) \end{array}}{\Pi \vdash \texttt{fork}\{e_1, e_2\} \rightsquigarrow \texttt{if}(cond) \texttt{ then } e_1' || e_2' \texttt{ else } e_1'; e_2' : (t_2, \sigma_1 \cup \sigma_2)}$$

Figure 5. Translation of $\texttt{fork}\{e_1, e_2\}$, to concurrent or sequential execution of e_1 and e_2.

jointness, as concretizations of these open effects may cause conflicts at runtime. More static information could help (T-UNKNOWN) to make some disjointness decisions statically, as discussed later.

Without open effects, expressions e_1 and e_2 may conflict if either e_1 or e_2 causes an invocation of a dynamically dispatched method especially if the method does not have any user-specified effect specifications. This is because a dynamically dispatched method with an unknown dynamic type and no effect specifications has the top effect that conflicts with any other effect [9].

2.3.1 Fork: An Example Use Case of Disjointness

An example use case of the disjoint expression $e_1 \# e_2$ is to combine static and runtime decision making about parallel or sequential execution of two expressions e_1 and e_2 in a fork expression. The fork expression $\texttt{fork}\{e_1, e_2\}$[4] executes e_1 and e_2 concurrently if their effects do not conflict, and sequentially otherwise. Figure 1 illustrates a fork expression on lines 10–13.

The rules (T-FORK-SEQUENTIAL) and (T-FORK-PARALLEL), in Figure 5, statically translate the fork expression to sequential or parallel executions of e_1 and e_2, respectively. The rule (T-FORK-UNKNOWN) defers such a decision to runtime because of the lack of the static information to decide disjointness of the effects of e_1 and e_2.

The rule (T-FORK-SEQUENTIAL) statically translates the fork expression to the sequential composition $e_1'; e_2'$ in which e_1' and e_2' run sequentially. The expressions e_1' and e_2' are translations of e_1 and e_2, respectively. This translation is sound because the effects of expressions e_1 and e_2 do conflict, i.e. $\Pi \vdash e_1 \# e_2 \rightsquigarrow false : (bool, \sigma)$. Similarly, the rule (T-FORK-PARALLEL) translates the fork expression into the parallel composition $e_1' || e_2'$, since the effects of e_1 and e_2 are disjoint, i.e. $\Pi \vdash e_1 \# e_2 \rightsquigarrow true : (bool, \sigma)$.

If (T-FORK-SEQUENTIAL) and (T-FORK-PARALLEL) cannot decide about sequential or parallel execution of the expressions in the fork, the decision is deferred to runtime using the rule (T-FORK-UNKNOWN). This rule translates a fork expression to an if expression $\texttt{if}(cond) \texttt{ then } e_1' || e_2' \texttt{ else } e_1'; e_2'$, which in its condition *cond* checks for disjointness of open effects σ_1^o and σ_2^o of the expressions in the fork, and their concrete effects σ_1^c and σ_2^c. The unknown open effects should be concretized, before being checked for disjointness. The auxiliary function *concretize* is used to concretize open effects at runtime. Concretization of open effects is discussed in §3. More static analyses, such as alias or purity analysis, or static effect specifications may help (T-FORK-UNKNOWN) to make some disjointness decisions statically, as discussed later in this section.

Without open effects, the fork expression $\texttt{fork}\{e_1, e_2\}$ will be translated to the sequential composition $e_1; e_2$, if either the expression e_1 or e_2 contains an invocation of a dynamically dispatched method with no effect specifications. This again is mainly because a dynamically dispatched method with an unknown dynamic type and no effect specifications has the top effect that conflicts with any other effect.

2.3.2 Alias Analysis Integration

Various modular static analyses could be integrated into our hybrid type-and-effect system to increase the precision of its static decision making about disjointness of the effects. In this section, we integrate a modular definite alias analysis [21] into open effects and illustrate its use.

For integration of alias analysis into open effects, we add an aliasing environment A to the typing judgement and change it to $\Pi, A \vdash e \rightsquigarrow e' : (t, \sigma, A')$. The aliasing environment A maps a variable to its aliases, i.e. $A ::= \{var_i \mapsto e_i\}_{i \in \mathbb{N}}$. The new typing judgement says that with the typing environment Π and aliasing environment A the expression e translates to e' and has the type t, the effects σ and the aliasing environment A'. For readability, some of the typing rules use the shorter typing judgement $\Pi, A \vdash x : t$

---

[4] Following previous work [9, 38], $\texttt{fork}\{e_1, e_2\}$ and $e_1 || e_2$ are used to illustrate a use case of the disjointness expression $e_1 \# e_2$ and are not part of the core syntax, in Figure 2.

for variables that do not cause any effect or changes in the aliasing. This shorter judgement stands for the judgement $\Pi, A \vdash x \rightsquigarrow x : (t, \emptyset, A)$.

The rules most concerned about aliasing information and keeping them updated are (T-DEFINE) and (T-SET). The rule (T-SET) assigns a variable x to a field f and creates the aliasing relation $x = \mathbf{this}.f$. This aliasing relation should be added to the aliasing environment A after discarding older aliasing relations for the field f via the kill operation $A \backslash f$. In (T-SET), the auxiliary function $typeOf$, takes a field f and returns the class d the field is defined in and its type t'.

(T-SET)
$$\frac{typeOf(f) = (d, t') \quad \Pi, A \vdash \mathbf{this} : c \quad \Pi, A \vdash x : t}{\Pi, A \vdash \mathbf{this}.f = x \rightsquigarrow \mathbf{this}.f = x : (t, \mathbf{wr}(f), A')}$$
$$c <: d \quad t <: t' \quad A' = A \backslash f \cup \{x = \mathbf{this}.f\}$$

In (T-DEFINE), a variable x is assigned the expression e'_1 in the scope of e'_2, creating the aliasing relation $x = e'_1$ that should be considered when evaluating e_2. The notation $A; x = e'_1$ stands for extending the aliasing environment A with the aliasing relation $x = e'_1$.

(T-DEFINE)
$$\frac{\Pi, A \vdash e_1 \rightsquigarrow e'_1 : (t_1, \sigma_1, A_1)}{\Pi; x : t, A_1; x = e'_1 \vdash e_2 \rightsquigarrow e'_2 : (t_2, \sigma_2, A_2) \quad t_1 <: t}{\Pi, A \vdash t\, x = e_1; e_2 \rightsquigarrow t\, x = e'_1; e'_2 : (t_2, \sigma_1 \cup \sigma_2, A_2)}$$

Use case (1): observational purity The static aliasing information maintained by the rules (T-SET) and (T-DEFINE) could be used in detecting observational purity [36, 41], as shown in the rule (T-CALL-PURE). The rule (T-CALL-PURE) says that in an invocation of $x_0.m(\overline{x})$ if both the receiver x_0 and the parameters $\overline{x}$ of the method m are newly created objects, then the effects of the invocation of m is empty [41], i.e. $\emptyset$. This means that we can statically decide to execute the method invocation $x_0.m(\overline{x})$ with any other expression concurrently without checking for their open effects. This is true because the effects of the method invocation is empty and does not conflict with any other effects of any other expression. The auxiliary function $findMeth$ looks up the class table CT and returns the declaration of the method m in the class c_0 or its supertypes.

(T-CALL-PURE)
$$\frac{A \vdash x_0 = \mathbf{new}\, c_0() \quad \forall x_i \in \overline{x}.\ A \vdash x_i = \mathbf{new}\, c_i()}{\Pi, A \vdash x_0.m(\overline{x}) \rightsquigarrow x_0.m(\overline{x}) : (t, \emptyset, \emptyset)}$$
$$findMeth(c_0, m) = (c', t, m\, (\overline{t\, var})\, \{e\}, \sigma)$$
$$\forall x_i \in \overline{x}.\ (\Pi, A \vdash x_i : t'_i) \wedge (t'_i <: t_i)$$

Use case (2): tracking open references Another use case of the static aliasing information is in statically tracking the open references, in the rule (T-CALL-OPEN). This rule assigns an open effect $\mathbf{open}(f\, m\, \gamma)$ to the method invocation $x_0.m(\overline{x})$ in which the receiver x_0 is an alias of the open field

f, i.e. $A \vdash x_0 = \mathbf{this}.f$. Without the aliasing information, and not knowing that the receiver x_0 is an alias of the open field f, the invocation $x_0.m(\overline{x})$ may get the top effect $\top$ because the runtime of its receiver is not known and it does not have any effect specifications.

(T-CALL-OPEN)
$$\frac{A \vdash x_0 = \mathbf{this}.f \quad typeOf(f) = (c, @\mathbf{open}\, c_0)}{\Pi, A \vdash x_0.m(\overline{x}) \rightsquigarrow x_0.m(\overline{x}) : (t, \mathbf{open}(f\, m\, \gamma), \emptyset)}$$
$$findMeth(c_0, m) = (c', t, m\, (\overline{t\, var})\, \{e\}, \sigma)$$
$$\forall x_i \in \overline{x}.\ (\Pi, A \vdash x_i : t'_i) \wedge (t'_i <: t_i)$$

For another example of the use of static aliasing information in deciding the disjointness of open effects, consider Figure 6 which shows a simplified example adapted from JavaGrande's RayTracer [47]. In this figure, the class `RayTracer` is responsible for rendering a display, and is extended by classes `RayTracer2D` and `RayTracer3D` to render two and three dimensional displays. Both of these subtypes override the method `run` in their supertype. In this example, using the purity analysis, one would conclude that the two expression of the fork, lines 5 and 6 can run concurrently, because the expression on line 5 has empty effects, the rule (T-CALL-PURE), and its pure effect does not conflict with the effects of the other expression.

```
1 @open RayTracer rt1;
2 rt1 = new RayTracer2D(new Display());
3 fork{
4   { RayTracer rt3D =
5     new RayTracer3D(new Display()); rt3D.run() },
6   { rt1.run() }
7 }
```

Figure 6. Concurrent execution of fork, because of observational purity [41].

2.3.3 More Static Analyses

In our prototype implementation of open effects, in addition to alias analysis, a few other static analyses such as purity analysis [41] and array effect analysis [42] are integrated into the type system. These analyses are not discussed here to focus on the basic ideas behind open effects.

2.4 Dynamic Dispatch and Open World Assumption

Two rules (T-CALL-OPEN) and (T-CALL) in our type-and-effect system type check dynamically dispatched method invocations. The differences between these rules highlight the contrast between handling of dynamic dispatch with and without open effects.

The rule (T-CALL-OPEN), discussed previously, uses an open effect to represent the unknown effects of a dynamically dispatched method invocation on an open receiver, and thus allowing it to run concurrently with other expressions if their effects do not conflict, or sequentially otherwise. In contrast, the rule (T-CALL) assigns the top effect $\top$ as the effects of the invocation of a dynamically dispatched

method on a non-open receiver, especially if there are no effect specifications for the method [9, 25], its receiver and its parameters are not newly created objects, or its receiver is not aliasing an open field. Assigning the top effect to a method invocation on a non-open receiver statically sequentializes its execution with any other expression, because the top effect conflicts with any other effects. However, assigning an open effect to invocation of a method with an open receiver, sequentializes execution of the method only if it cannot prove disjointness of its effects with other expressions. (T-CALL-OPEN) and (T-CALL) clearly show how open effects could be useful in exposing safe concurrency opportunities.

(T-CALL)
$$\Pi, A \vdash x_0 : c_0 \qquad A \vdash x_0 \neq \mathbf{new}\ c() \vee (\exists x_i \in \bar{x}.\ A \vdash x_i \neq \mathbf{new}\ c_i())$$
$$A \vdash x_0 \neq \mathbf{this}.f \vee (A \vdash x_0 = \mathbf{this}.f \wedge typeOf(f) \neq (d, @\mathbf{open}\ c''))$$
$$findMeth(c_0, m) = (c', t, m\,(\overline{t\ var}\,)\,\{e\}, \sigma)$$
$$\forall x_i \in \bar{x}.\ (\Pi, A \vdash x_i : t_i') \wedge (t_i' <: t_i)$$
$$\overline{\Pi, A \vdash x_0.m(\bar{x}) \rightsquigarrow x_0.m(\bar{x}) : (t, \top, \emptyset)}$$

Field set There are two rules (T-SET) and (T-SET-OPEN) for setting a field. The rule (T-SET), shown previously, sets a non-open field which results in a write effect and updating the aliasing information of the field. The rule (T-SET-OPEN) is similar to (T-SET) except that it generates a top effect $\top$, instead of a write effect. This is because setting an open field f results in concretizations of all open effects $\mathbf{open}(f\ m\ \gamma)$ with the open field f and any other open effect $\mathbf{open}(g\ m'\ \gamma')$ where g transitively points to the object containing f. The set expression does not know about all open effects $\mathbf{open}(g\ m'\ \gamma')$ or even $\mathbf{open}(f\ m\ \gamma)$ which are dependent on its field f and thus to be sound it has to assume the top effect $\top$ to cover all such effects. Concretization of open effects is discussed in more detail in §3.

(T-SET-OPEN)
$$typeOf(f) = (d, @\mathbf{open}\ t') \qquad \Pi, A \vdash \mathbf{this} : c \qquad \Pi, A \vdash x : t$$
$$c <: d \qquad t <: t' \qquad A' = A \backslash f \cup \{x = \mathbf{this}.f\}$$
$$\overline{\Pi, A \vdash \mathbf{this}.f = x \rightsquigarrow \mathbf{this}.f = x : (t, \top, A')}$$

2.4.1 Effect Specification Integration

Effect specifications statically specify an upper bound for the read and write memory effects of a dynamically dispatched method, independent of the dynamic type of its receiver. Effect specifications could be integrated into our hybrid type-and-effect system and used to discharge effect disjointness checks statically.

To integrate effect specifications in our type-and-effect system, similar to the class table CT, we assume an implicit specification table ST that maps a method to its effect specifications θ, if it has any. The effect specification θ is a set of concrete read and write effects $\mathbf{wr}(f)$ and $\mathbf{rd}(f)$ to memory locations f and does not contain any open effects.

Method declaration There are two rules (T-METHOD) and (T-METHOD-SPEC) in our type-and-effect system to type

check a method declaration depending on if effect specification for the method are available or not.

(T-METHOD)
$$(c, t\ m(\overline{t\ var})\{e'\}) \notin dom(ST)$$
$$override(m, c, (\bar{t} \rightarrow t)) \qquad \forall t_i \in \bar{t}.\ isType(t_i) \qquad isType(t)$$
$$(\overline{var : t}, \mathbf{this} : c), \emptyset \vdash e \rightsquigarrow e' : (t', \sigma, A) \qquad t' <: t$$
$$\overline{\vdash t\ m(\overline{t\ var})\{e\} \rightsquigarrow t\ m(\overline{t\ var})\{e'\} : (\bar{t} \rightarrow t, \sigma, A)\ \mathbf{in}\ c}$$

The rule (T-METHOD) type checks a method declaration that does not have any effect specifications. The rule (T-METHOD) says that latent effect of a method is the same as the effects of its body. In the absence of effect specifications for a method, the statically computed effects of the method and the methods that override it do not have any relation and can vary independently. This is represented in the rule (T-METHOD) by not requiring any relation between the effects σ of the method m and any other method with effects σ' it may override. The auxiliary function *override* only checks for compatibility of the argument and return types of the overridden and overriding methods with no effect specifications and allows their effects to be independent. The function *isType* checks for validity of a type.

(T-METHOD-SPEC)
$$\theta = ST(c, m)$$
$$override(m, c, (\bar{t} \rightarrow t)) \qquad \forall t_i \in \bar{t}.\ isType(t_i) \qquad isType(t)$$
$$(\overline{var : t}, \mathbf{this} : c), \emptyset \vdash e \rightsquigarrow e' : (t', \sigma, A) \qquad t' <: t$$
$$\overline{\vdash t\ m(\overline{t\ var})\{e\} \rightsquigarrow t\ m(\overline{t\ var})\{e'\} : (\bar{t} \rightarrow t, \theta, A)\ \mathbf{in}\ c}$$

The rule (T-CALL-SPEC) type checks the declaration of a method with effect specifications. The rule (T-CALL-SPEC) says the latent effect of a method with effect specification θ is θ instead of its statically computed effect σ. The auxiliary function $ST(c, m)$ returns the effect specifications of method m in class c. Effect specifications specify an upper bound for memory effects of a method independent of dynamic dispatch. Such independence is achieved by enforcing effect containment between the method and methods overriding it, such that the effect specifications of an overriding method in a subclass is contained in the effects of the superclass method it overrides. In (T-METHOD) the auxiliary function *override* should enforce effect containment, which boils down to checking if the set of effect specifications of an overriding method is the subset of the set of effect specifications of the method it overrides.

Method invocation The rules (T-CALL) and (T-OPEN-CALL) discussed previously assume no effect specifications for a method declaration. The rule (T-CALL-SPEC) type checks a method invocation where the method has effect specifications. The rule (T-CALL-SPEC) says that effects of the invocation of a method with effect specifications θ is the same as θ independent of if its receiver being an open receiver or not. Effect specifications, if available, improve the effect

precision of method invocation rules, especially compared to (T-Call-Open) in which the effect of the method invocation is the top effect $\top$. For each write effect of a field f in the specification of the method, its associated aliasing information in the aliasing environment must be killed.

(T-Call-Spec)
$$\frac{\begin{array}{cc} \Pi,A \vdash x_0 : c_0 & ST(c_0,m) = \theta \\ findMeth(c_0,m) = (c',t,m(\overline{t\ var})\ \{e\},\sigma) \\ \forall x_i \in \overline{x}.\ (\Pi,A \vdash x_i : t'_i) \wedge (t'_i <: t_i) \quad \forall f, \mathbf{wr}(f) \in \theta.A' = A \backslash f \end{array}}{\Pi,A \vdash x_0.m(\overline{x}) \rightsquigarrow x_0.m(\overline{x}) : (t,\theta,A')}$$

3. A Dynamic Semantics with Open Effects

The dynamic part of open effects which is encoded in *OpenEffectJ*'s dynamic semantics, *(i)* concretizes the statically computed open effects using the typing rules, and updates the open effects by tracking their open references and changes in their values; and *(ii)* verifies, using runtime checks, the disjointness assumptions that could not be verified statically.

3.1 Dynamic Semantics Objects

The dynamic semantics of open effects transitions from one configuration to another. A configuration $\Sigma = \langle e,\mu \rangle$, shown in Figure 7, consists of an expression e and a global store μ. The store maps a location *loc* to an object record of the form $o = [c.F.E]$, containing the concrete type c of the object *loc*, a field map F which maps field names of c to their values, and a new dynamic effect map E which *maps the method names of c to their runtime effects*. The effect map is necessary in tracking and updating of runtime effects of dynamically dispatched methods, for concretization of open effects and efficient verification of their disjointness, as the values of open references change during the program execution. Performance efficiency of these mechanisms is shown in §4. Dynamic semantics rules are presented using a one-step call-by-value reduction relation and a set of evaluation contexts $\mathbb{E}$ [22] which specify the evaluation order. Omitted semantics rules and auxiliary functions can be found in our technical report [34].

Evaluation relation:　$\hookrightarrow: \Sigma \dashrightarrow \Sigma$
Domains:

Σ	::=	$\langle e,\mu \rangle$	"Configurations"
μ	::=	$\{loc_i \mapsto o_i\}_{i\in\mathbb{N}}$	"Stores"
o	::=	$[c.F.E]$	"Object Records"
F	::=	$\{f_i \mapsto v_i\}_{i\in\mathbb{N}}$	"Field Maps"
v	::=	$\mathbf{null} \mid loc \mid n \mid b$	"Values"
E	::=	$\{m_i \mapsto \sigma_i\}_{i\in\mathbb{N}}$	"Effect Maps"

Evaluation contexts:

$$\mathbb{E} \quad ::= \quad - \mid t\ var\ =\ \mathbb{E}; e$$

Figure 7. Domains and evaluation contexts.

3.2 Tracking and Updating of Open References

There are several rules in *OpenEffectJ*'s dynamic semantics, including the rule for object creation and setting a field, that are key in tracking open references and updating the concretization of open effects dependent on these references.

(New)
$$\frac{\begin{array}{cc} loc \notin dom(\mu) & F = \{f \mapsto default(f) \mid f \in fields(c)\} \\ \mu' = \mu \oplus \{loc \mapsto [c.F.E]\} \\ E = \{m \mapsto \sigma \mid m \in methods(c).findMeth(c,m) = (c',t,m(\overline{t\ var}),\sigma)\} \end{array}}{\langle \mathbb{E}[\mathbf{new}\ c()],\mu \rangle \hookrightarrow \langle \mathbb{E}[loc],\mu' \rangle}$$

The rule (New), in addition to initializing a new object in memory, by assigning a fresh location *loc* to it, generates and initializes the effect map E for the newly created object. The effect map E maps the methods m of class c to their statically computed effects σ which have been computed using the typing rules as discussed in §2. The auxiliary function *findMeth* returns the definition of a method m of the class c in the class table *CT*. The function *default* returns the default value for each variable of a type. The operator $\oplus$ is an overriding operator such that if $\mu' = \mu \oplus \{loc \mapsto o\}$, then $\mu'(loc') = o$ if $loc' = loc$, otherwise $\mu'(loc') = \mu(loc')$. To illustrate, the object record for the newly created object pr of type Pair in Figure 1 is of the form $[Pair.\{fst \mapsto 0, snd \mapsto 0, f \mapsto \mathbf{null}\}.\{setOp \mapsto \{\top\}, apply \mapsto \{\mathbf{wr}(fst), \mathbf{wr}(snd), \mathbf{rd}(f), \mathbf{open}(f\ op\ \gamma)\}\}]$.

(Set)
$$\frac{[c.F.E] = \mu(loc) \quad \mu_0 = \mu \oplus (loc \mapsto [c.(F \oplus \{f \mapsto v\}).E]) \\ \mu' = update(\mu_0,loc,f,v)}{\langle \mathbb{E}[loc.f = v],\mu \rangle \hookrightarrow \langle \mathbb{E}[v],\mu' \rangle}$$

Setting the field f of the object *loc* updates the concretization of all open effects that are directly or transitively dependent on the open field f, until a fixpoint is reached. This is done using the auxiliary function *update*, in Figure 8, that first concretizes the open effects in the effect map E of the object *loc*. If the effect map E changes to E', i.e. $E \neq E'$, then it updates the concretization of all other transitively dependent open effects. The function *reverse* backward traverses the object graph, starting from *loc*, and finds all the open fields dependent on f, directly or transitively. An open field g is dependent on the open field f, if g, directly or transitively points to the object *loc* containing the field f. In practice, reverse pointers can be used to optimize this [8], as in our compiler's implementation.

Concretization of Open Effects There are two variations of concretization, shown in Figure 8: *(i)* concretization of an open effect when its open field is set, as in (Set), by *concretize$_{(1)}$* and *(ii)* concretization of an open effect in use cases such as translation of the fork in the rule (T-Fork-Unknown) in §2, by *concretize$_{(2)}$*. The function *concretize$_{(1)}$*, fills in the placeholder γ in open effects of the form $\mathbf{open}(f\ m\ \gamma)$, upon setting the field f of the object *loc*. Recall that $\mathbf{open}(f\ m\ \gamma)$ represents the effect of the invocation of the method m on the open field f. If f is set to $\mathbf{null}$,

$$update(\mu,loc,f,v) = \begin{cases} \mu & \text{if } E = E', \\ \mu_n & \text{if } E \neq E', \\ & \quad \mu_i = update(\mu_{i-1},loc_i,f_i,loc) \\ & \quad \{\langle loc_i,f_i\rangle\} = reverse(\mu,loc), 1 \leq i \leq n \\ & \quad \mu_0 = \mu \oplus \{loc \mapsto [c.F.E']\} \end{cases} \qquad \begin{aligned} &\textbf{where } E' = updateEff(\mu,f,v,E), \\ &\qquad \mu(loc) = [c.F.E] \end{aligned}$$

$$reverse(\mu,loc) = \{\langle loc',f\rangle \mid F(f) = loc \wedge loc' \in dom(\mu) \wedge \mu(loc') = [c.F.E]\}$$

$$updateEff(\mu,f,v,E) = \{ m \mapsto \{concretize_{(1)}(\mu,f,v,\varepsilon)\} \mid (m \mapsto \sigma) \in E \ \wedge \ \varepsilon \in \sigma\}$$

$$concretize_{(1)}(\mu,f,v,\varepsilon) = \begin{cases} \textbf{open}(f\ m\ \emptyset) & \text{if } \varepsilon = \textbf{open}(f\ m\ \gamma),\ v = \texttt{null} \\ \textbf{open}(f\ m\ \sigma) & \text{if } \varepsilon = \textbf{open}(f\ m\ \gamma),\ v = loc',\ [c.F.E] = \mu(loc'), \\ & \quad \sigma = \downarrow_C (E(m)) \ \cup \ \{\varepsilon' \in \sigma' \mid \varepsilon'' \in \downarrow_O (E(m)) \wedge \varepsilon'' = \textbf{open}(f'\ m'\ \sigma')\} \\ \varepsilon & \text{otherwise} \end{cases}$$

$$concretize_{(2)}(\mu,loc,\varepsilon) = \sigma \qquad \textbf{where } \varepsilon = \textbf{open}(f\ m\ \gamma),\ [c.F.E] = \mu(loc), \exists\,m' \in dom(E).\ \textbf{open}(f\ m\ \sigma) \in E(m')$$

$$concretize_{(2)}(\sigma) = (\textstyle\bigcup concretize_{(2)}(\mu,\texttt{this},\varepsilon)) \cup \downarrow_C (\sigma) \qquad \textbf{where } \varepsilon \in \downarrow_O (\sigma)$$

Figure 8. Auxiliary functions *update* and *concretize*

---

then invocation of m on f will not have any effects, thus replacing γ in $\textbf{open}(f\ m\ \gamma)$ by the concretized effect $\emptyset$, i.e. $\textbf{open}(f\ m\ \emptyset)$. If f is set to an object loc', then the invocation of m on f will be the invocation of m on the object loc', thus replacing γ in $\textbf{open}(f\ m\ \gamma)$ with the union of the concrete effects of the method m in the object loc', i.e. $\downarrow_C (E(m))$, and its concretized open effects, i.e. σ'. Note that, the open effect $\textbf{open}(f\ m\ \gamma)$ is concretized whenever its open field f is set. The function $concretize_{(2)}$ basically returns the effects concretized by the first variation, rather than directly concretizing them. This is because concretization of an open effect happens only when its open field is set. For a non-concretized open effect $\textbf{open}(f\ m\ \gamma)$ with the open field f in the object loc, its effect map E is searched till a concretized effect $\textbf{open}(f\ m\ \sigma)$ is found and σ is returned as the result. The current store μ in the configuration and variable $\texttt{this}$ are implicitly passed to $concretize_{(2)}$ in which $\texttt{this}$ is passed as the value for loc.

To illustrate, consider concretization of the open effect $\textbf{open}(f\ op\ \gamma)$ of the method $\texttt{apply}$ when its open field $\texttt{f}$ is set, by the expression $\texttt{pr.setOp(pf)}$ on line 29 of Figure 1. In $concretize_{(2)}$, the parameter v will be equal to $\texttt{pf}$ and thus the placeholder γ in the open effect will be replaced by $\textbf{wr}(res)$, which is the effect of of method $\texttt{op}$ of $\texttt{pf}$. Concretization of the open effect $\textbf{open}(f\ op\ \gamma)$ to $\textbf{open}(f\ op\ \textbf{wr}(res))$ in the effect of $\texttt{apply}$ causes the *update* to be invoked which updates all open effects which are dependent on $\texttt{pr}$, using *reverse*. In Figure 1, there is no object pointing to $\texttt{pr}$ and thus the fixpoint is reached and concretization stops.

3.3 Soundness of Open Effects

The type-and-effect encoding of open effects is proven sound using theorems that say: *(i)* statically computed effects are a sound approximation of concretized effects, Theorem 3.1; and *(ii)* concretized effects soundly approximate runtime effects, Theorem 3.2.

THEOREM 3.1. [Concretized effects refine static effects] *Given an expression e with the statically computed effects σ_s, which could contain open effects, and its dynamic concretization σ_c, i.e. $\sigma_c = concretize(\sigma_s)$, if $\Pi,A \vdash e \rightsquigarrow e' : (t,\sigma_s,A')$ holds statically and $(\mu,A) \vdash e' : (\sigma_c,A')$ holds dynamically for the runtime configuration $\langle e',\mu\rangle$, then: $\sigma_c \subseteq \sigma_s$.*

THEOREM 3.2. [Dynamic effects refine concretized effects] *For two configurations $\Sigma = \langle e,\mu\rangle$ and $\Sigma' = \langle e',\mu'\rangle$, if Σ transitions to Σ' producing runtime effect η, i.e. $\Sigma \xrightarrow{\eta} \Sigma'$, if concretized effects of e is σ_c, i.e. $(\mu,A) \vdash e : (\sigma_c,A')$, then there is a concretized effect σ_c' such that:*

(a) $(\mu',A_1) \vdash e' : (\sigma_c',A_1')$ and $\sigma_c' \subseteq \sigma_c$;
(b) $\eta \in \sigma_c$

Proof Sketch: Theorem 3.1 is proved by structural induction on derivations of $\Pi,A \vdash e \rightsquigarrow e' : (t,\sigma_s,A')$ and $(\mu,A) \vdash e' : (\sigma_c,A')$ whereas proof of Theorem 3.2 is by cases on transition steps for the transition relation $\Sigma \xrightarrow{\eta} \Sigma'$ [34].

4. Evaluation

In this section, we describe an evaluation of the open effects system. First, we describe benchmarks used in the experiments. Then, we present a detailed performance evaluation.

Finally, we analyze the distribution of static vs. dynamic checks for these programs.

4.1 Benchmark programs

We use the following frameworks, benchmarks and libraries in our experiments: a map-reduce framework (MapReduce), adapted from JSR [1], a pipeline framework (Pipeline) [10], Monte Carlo benchmark (MonteCarlo) [47], JDK's merge sort (MergeSort) and array list (ArrayList) libraries [1], depth first search graph traversal (DFS), a numerical integration application (Integrate) [1] and a sequence alignment application (Alignment) [5]. *In terms of annotation overhead, except MapReduce and DFS, with 2 @open annotations, open effects versions of other applications needed only 1 annotation.*

4.2 Performance Evaluation

We hypothesize that open effects is performance efficient while exposing safe concurrency opportunities in frameworks and libraries that could be *extended with possibly concurrency-unsafe code by clients*. To test our hypothesis we implemented open effects on top of OpenJDK[5] and parallelized a representative set of frameworks and libraries using open effects and the following widely used concurrency techniques: *(i)* Deuce [2], software transactional memory (STM); *(ii)* Multiverse [3], STM; *(iii)* RoadRunner (RR) [20], runtime race detector[6]; and *(iv)* Manually tuned concurrency; and compared their speedups and overheads. Results of our experiments show that: open effects almost does as well as manually tuned concurrency, with the negligible overhead of only 0.1% to at most 4.1% and less overhead compared to other techniques.

4.2.1 Setup

Client Code In MapReduce, the map phase computes the sum of the magnitudes formula $Math.sqrt(2*Math.pow(o,2))$ for each element o in a set of 100 million integers and the reduce step simply adds the results. Pipeline models Radix Sort in which the first stage generates a stream of 8 arrays of 1 million integers each and subsequent stages sort the arrays on different radixes. MergeSort sorts a list of 10 million randomly generated integers. For ArrayList, we apply the hash (Hash), prefix sum (Prefix) computations, illustrated in Figure 1, and a heavier computation (Heavy), which computes the same formula as in the MapReduce, on an array with 20 million elements. DFS solves an N-queens problem, with n equal to 11. Integration uses a recursive Gaussian quadrature of $(2*i-1) \cdot x^{2*i-1}$, summing over odd values of i from 1 to 12 and integrating from -5 to 6. Alignment uses a constant function returning -1 if two characters do not match for aligning two words of sizes 100 and 1 million.

Hardware All our experiments were run on a system with a total of 4 cores (Intel Core2 chips 2.40GHz) running Fedora GNU/Linux. For each experiment, an average of the results over 30 runs was taken and the default JVM parameters were used.

4.2.2 Overall Performance

DEFINITION 4.1. *(Runtime Overhead and Speedup) For a program p, with its sequential, open effects and manually tuned parallel versions, which respectively take T1, T2, and T3 seconds for their execution, the speed up is T1 / T2 and overhead is (T2 - T3) / T3.*

Figure 9 and Figure 10 show the performance results of running our experiments in terms of speedup and runtime overhead, as defined in Definition 4.1. Our results show that the open effects (*OpenEffectJ*) versions of the evaluation applications are almost as fast as the manually tuned concurrent versions and incur very small overhead, ranging from 0.1% to 4.1% at most, which is significantly less compared to other effect analysis techniques used in our experiments. In Figure 9, Multiverse or Deuce versions of the applications, run slower especially for ArrayList, because Multiverse creates a separate runnable object for each transaction which causes a slow down in applications with large number of transactions; Deuce, besides the transaction creation overhead, stores array access effects in a fine-grained manner, which could cause slow down for large arrays. *OpenEffectJ*, instead, uses an indexed array effect [9, 42] and a purity analysis, to decide about disjointness of effects. It is conceivable that using similar techniques, performance for Multiverse or Deuce's versions could be improved.

4.2.3 Effect Concretization Overhead

One may argue that concretization of open effects at runtime may not be efficient enough especially for deeply-nested data structures such as trees, mainly because setting an open field could cause updating various other open effects dependent on that field, as discussed in §3.2. However, this may not always be the case, especially when the benefits of the open effects outweigh its overheads. To investigate, we implemented a Fibonacci algorithm, adapted from OpenJDK's fork/join framework [1], using both open effects and techniques in Figure 9. To compute the n-th Fibonacci number $Fib(n)$, the algorithm uses a binary tree in which right and left subtrees represent $Fib(n-1)$ and $Fib(n-2)$, and the root adds them together to compute $Fib(n) = Fib(n-1) + Fib(n-2)$. This algorithm has two phases of *(i)* construction of the tree and *(ii)* computation of the Fibonacci numbers for the nodes. We ran experiments on trees of depths from 6 to 14, with up to 2^{14}-1 nodes, and the same number of open effects' concretizations, to compute $Fib(45)$.

Figure 11 shows the time each technique takes to compute the Fibonacci numbers, for the construction and computation phases as well as their total. It shows that for nested objects,

Application	Serial	Manual	RR[20]	Deuce[2]	Multiverse[3]	OpenEffectJ				Pattern
	time(s)	time	time	time	time	time	overhead %	speedup	annotation	
Hash	0.13	0.12	0.85	44.11	8.57	0.12	0.3	1.07	1	Forall
Heavy	1.31	0.39	1.12	34.87	12.43	0.39	0.1	3.39	1	Forall
Prefix	0.12	×	×	×	×	0.12	1.6	0.98	1	Forall
Alignment	2.44	1.86	21.50	14.91	8.34	1.93	4.1	1.26	1	Forall
MonteCarlo	3.87	1.22	2.04	10.52	1.33	1.25	2.7	3.10	1	Forall
Pipeline	2.25	2.11	3.48	↑	2.21	2.12	0.6	1.06	1	Pipeline
MergeSort	2.71	1.32	3.39	9.61	16.00	1.34	1.7	2.02	1	Recursive
DFS	18.83	9.20	17.88	9.79	12.23	9.23	0.3	2.04	2	Recursive
MapReduce	7.03	1.94	3.81	5.25	10.76	1.91	-1.5	3.68	2	Recursive
Integrate	2.13	0.59	1.53	1.46	2.42	0.61	2.4	3.50	1	Recursive

Figure 9. Performance Experiments. × indicates result discrepancies because of sequential inconsistencies and ↑ shows running out of memory after a considerably long time.
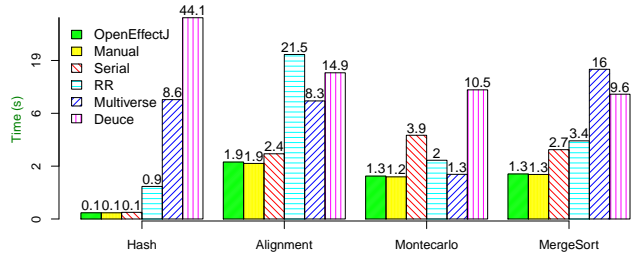
Figure 10. Open effects are almost as good as manually tuned concurrency.

Construction   Computation

Total
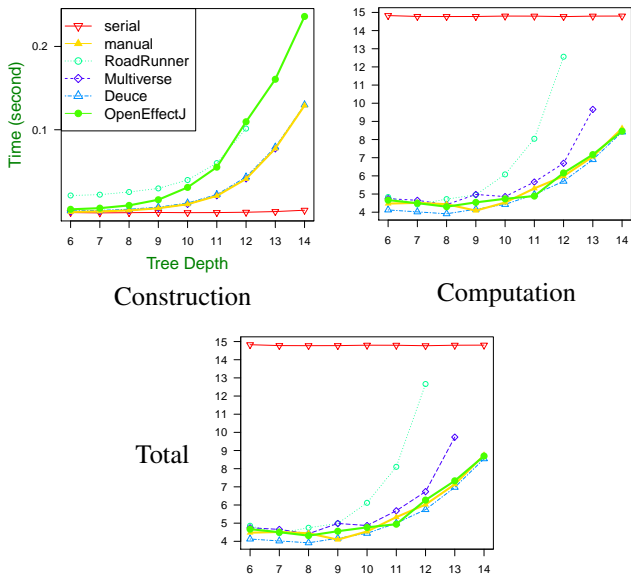
Figure 11. Building Fibonacci tree and computation of n-th Fibonacci number.

open effects does almost as well as the manually tuned concurrent version and better than other concurrent approaches, which is consistent with the results in Figure 9. This is despite the fact that the construction of the Fibonacci trees may take more time compared to other techniques, because of the concretization of open effects. However, benefits of open effects in the computation phase outweigh this overhead.

4.3 Disjointness Checks

Open effects decide the disjointness of effects dynamically only if it could not be decided statically. This divides the responsibility of deciding the disjointness between static and dynamic parts of our hybrid type-and-effect system.

Figure 12 shows the number of disjointness checks decided statically and dynamically for our evaluation applications. Unlike the performance evaluation where inputs of the evaluation applications are large, i.e. an array of 100 million for the MapReduce application, to measure the number of checks, the input of our evaluation applications are set in a way that they only decide the disjointness of two expressions. The number of checks are proportional to the size of the input in each application.
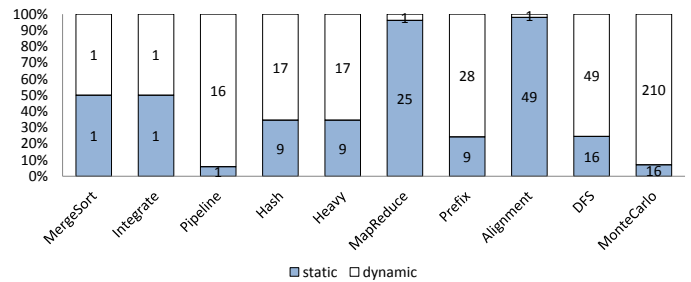
Figure 12. Statically vs dynamically decided disjointness checks.

In Figure 12 more than a quarter, i.e. about 28.5%, of the total number of disjointness checks across all evaluation applications are decided statically and the rest are decided dynamically. The number of statically versus dynamically decided checks varies from one application to another, with Alignment and MapReduce with the most number of statically decided checks and Pipeline and MonteCarlo with the least number of statically decided checks.

5. Related Work

In this section, we compare open effects with related works on reasoning about effects of programs, in three categories of hybrid, static and dynamic techniques.

Hybrid Open effects is closest in spirit to the ideas of gradual typing [45] and hybrid type checking [28] that blend advantages of static and dynamic type checking. Similarly, open effects blends the advantages of static and dynamic effect systems. Similar to Open Type Switch (Mach7) [48], which allows users to choose between type hierarchy openness and efficiency, open effects lets programmer choose the

openness of the effects of dynamically dispatched method invocations. Synchronization via scheduling (SVS) [8] computes effects of concurrent tasks as their reachable object graph, for programs written in a simple C-like language, with no dynamic dispatch. However, open effects support a full OO language with the support for overriding and dynamic dispatch, which makes accurate effect computation more challenging [49], and use smaller effect sets compared to reachability graphs for effect computations. Legion [51] and TWEJava [27] let the programmer specify the effect of each task and have a scheduler that coordinates these concurrent tasks. However, open effects require only open annotations, compared to task specifications. Chugh *et al.*[15] present a hybrid framework to check statically generated security constraints in JavaScript. In concurrent revisions [13], programmers annotate shared objects tasks could conflict on and provide their merge functions, and each task keeps a local copy of these objects to avoid data races, using copy-on-write. In contrast, open effects check for effect conflicts before the execution, either statically or at runtime.

Static Boyapati *et al.*[12] propose an ownership type system for deadlock and data race detection, Gordon *et al.*[24] use uniqueness and reference immutability to provide safe parallelism, Deterministic Parallel Java (DPJ) [10] provides determinism for parallel programs using effect parameters and effect constraints, such as effect containment. There are also other works on effect systems [25] for sequential programs such as data groups [31], ownership type systems [14] and heap representation techniques [14]. However, open effects is a hybrid technique that combines static and dynamic type-and-effect to better handle invocation of dynamically dispatched methods, without restrictions of effect containment between a type and its subtypes.

Dynamic FastTrack [18], Goldilocks [17], Pacer [11], IFRit [**?**], HAWKEYE [**?**], LEAN [**?**], and the work of Smaragdakis *et al.*[46] are data race detection techniques which monitor memory accesses. Transactional memory techniques [7, 16, 26, 33, 35, 43, 44, 52**? ?**] optimistically execute tasks concurrently, while also monitoring memory accesses, and rollback whenever effects of the tasks conflict. These techniques monitor memory footprints of a program, for all of its references. However, open effects only monitors open references and decide before execution of tasks, either statically or dynamically, if they need to be executed sequentially, because of conflicting effects, and thus do not roll back. In Galois [29], user provided commutativity specifications for methods are checked dynamically at runtime and the execution is rolled back if they are violated. However, open effects does not need commutativity specifications and does not roll back the execution.

6. Conclusion and Future Work

We proposed open effects, a trust-but-verify hybrid type-and-effect system to safely expose concurrency opportunities in invocations of dynamically dispatched methods in concurrent programs with an open world assumption and no effect specifications. We showed integration of a static alias analysis and effect specifications into open effects. We also showed that our prototype implementation of open effects could be quite efficient in exposing concurrency, by combining static and dynamic analyses, and performed as well as manually tuned concurrency, with negligible overhead of only *0.1 − 4.1%*. Since open effects is *complementary* to previous static and dynamic analyses, we believe this overhead could be decreased even more by integrating more sophisticated static and dynamic analyses, which is one venue for future work. Another direction for future work, is to explore a logical extreme, in which all references are implicitly open and a static analysis systematically eliminates ones causing unacceptable overheads.

References

[1] JSR-166y for Java 7. `http://gee.oswego.edu/dl/concurrency-interest/`.

[2] `https://sites.google.com/site/deucestm/`.

[3] `http://multiverse.codehaus.org/`.

[4] M. Abadi, C. Flanagan, and S. Freund. Types for safe locking: Static race detection for Java. *TOPLAS '06*, 28.

[5] S. Aluru, N. Futamura, and K. M. C. Parallel biological sequence comparison using prefix computations. *Journal of Parallel and Distributed Computing*, 2003.

[6] N. Benton and P. Buchlovsky. Semantics of an effect analysis for exceptions. In *TLDI '07*.

[7] E. D. Berger, T. Yang, T. Liu, and G. Novark. Grace: safe multithreaded programming for C/C++. In *OOPSLA '09*.

[8] M. J. Best, S. Mottishaw, C. Mustard, M. Roth, A. Fedorova, and A. Brownsword. Synchronization via scheduling: techniques for efficiently managing shared state. In *PLDI '11*.

[9] R. Bocchino, V. Adve, D. Dig, S. Adve, S. Heumann, R. Komuravelli, J. Overbey, P. Simmons, H. Sung, and M. Vakilian. A type and effect system for Deterministic Parallel Java. In *OOPSLA '09*.

[10] R. L. Bocchino and V. S. Adve. Types, regions, and effects for safe programming with object-oriented parallel frameworks. In *ECOOP '11*.

[11] M. Bond, K. Coons, and K. McKinley. Pacer: proportional detection of data races. In *PLDI'11*.

[12] C. Boyapati, R. Lee, and M. Rinard. Ownership types for safe programming: preventing data races and deadlocks. In *OOPSLA '02*.

[13] S. Burckhardt, A. Baldassin, and D. Leijen. Concurrent programming with revisions and isolation types. In *OOPSLA '10*.

[14] N. Cameron, S. Drossopoulou, J. Noble, and M. Smith. Multiple ownership. In *OOPSLA '07*.

[15] R. Chugh, J. A. Meister, R. Jhala, and S. Lerner. Staged information flow for JavaScript. In *PLDI '09*.

[16] C. Ding, X. Shen, K. Kelsey, C. Tice, R. Huang, and C. Zhang. Software behavior oriented parallelization. In *PLDI '07*.

[17] T. Elmas, S. Qadeer, and S. Tasiran. Goldilocks: a race and transaction-aware Java runtime. In *PLDI '07*.

[18] C. Flanagan and S. Freund. FastTrack: efficient and precise dynamic race detection. In *PLDI '09*, .

[19] C. Flanagan and S. Freund. Type-based race detection for Java. In *PLDI '00*, .

[20] C. Flanagan and S. Freund. The roadrunner dynamic analysis framework for concurrent programs. In *PASTE '10*, .

[21] C. Flanagan and S. Freund. Redcard: Redundant check elimination for dynamic race detectors. In *ECOOP' 13*, .

[22] M. Flatt, S. Krishnamurthi, and M. Felleisen. A Programmer's Reduction Semantics for Classes and Mixins. In *Formal Syntax and Semantics of Java*. Springer, 1999.

[23] D. Gifford and J. Lucassen. Integrating functional and imperative programming. In *LFP '86*.

[24] C. Gordon, M. Parkinson, J. Parsons, A. Bromfield, and J. Duffy. Uniqueness and reference immutability for safe parallelism. In *OOPSLA '12*.

[25] A. Greenhouse and J. Boyland. An object-oriented effects system. In *ECOOP '99*.

[26] M. Herlihy and J. E. B. Moss. Transactional memory: architectural support for lock-free data structures. In *ISCA '93*.

[27] S. Heumann, V. Adve, and S. Wang. The tasks with effects model for safe concurrency. In *PPoPP '13*.

[28] K. Knowles and C. Flanagan. Hybrid type checking. *TOPLAS '10*, 32.

[29] M. Kulkarni, K. Pingali, B. Walter, G. Ramanarayanan, K. Bala, and L. P. Chew. Optimistic parallelism requires abstractions. In *PLDI '07*.

[30] B. W. Lampson, J. J. Horning, R. L. London, J. G. Mitchell, and G. J. Popek. Report on the programming language Euclid. *SIGPLAN Not. '77*, 12(2).

[31] K. R. M. Leino. Data groups: specifying the modification of extended state. In *OOPSLA '98*.

[32] X. Leroy and F. Pessaux. Type-based analysis of uncaught exceptions. *TOPLAS '00*, 22(2).

[33] M. Lesani and J. Palsberg. Communicating memory transactions. In *POPL '11*.

[34] Y. Long, M. Bagherzadeh, and H. Rajan. Open Effects: Programmer-guided Effects for Open World Concurrent Programs. Technical report, Iowa State Univ., 2013.

[35] R. Lublinerman, J. Zhao, Z. Budimlić, S. Chaudhuri, and V. Sarkar. Delegated isolation. In *OOPSLA '11*.

[36] J. M. Lucassen and D. K. Gifford. Polymorphic effect systems. In *POPL '88*.

[37] R. Milner. A theory of type polymorphism in programming. *Journal of Computer And System Sciences '78*, (17).

[38] I. Neamtiu, M. Hicks, J. S. Foster, and P. Pratikakis. Contextual effects for version-consistent dynamic software updating and safe concurrent programming. In *POPL '08*.

[39] R. Reiter. *On closed world data bases*. Springer, 1978.

[40] M. C. Rinard and M. S. Lam. The design, implementation, and evaluation of Jade. *TOPLAS*, 20, 1998.

[41] S. Ru and M. Rinard. Purity and side effect analysis for Java programs. In *VMCAI '05*.

[42] R. Rugina and M. Rinard. Automatic parallelization of divide and conquer algorithms. In *PPoPP '99*.

[43] N. Shavit and D. Touitou. Software transactional memory. In *PODC '95*.

[44] T. Shpeisman, V. Menon, A.-R. Adl-Tabatabai, S. Balensiefer, D. Grossman, R. Hudson, K. Moore, and B. Saha. Enforcing isolation and ordering in STM. In *PLDI '07*.

[45] J. Siek and W. Taha. Gradual typing for objects. In *ECOOP '07*.

[46] Y. Smaragdakis, J. M. Evans, C. Sadowski, Y. Jaeheon, and C. Flanagan. Sound predictive race detection in polynomial time. In *POPL '12*.

[47] L. Smith, J. Bull, and J. Obdrizalek. A parallel Java Grande benchmark suite. In *SC '01*.

[48] Y. Solodkyy, G. Dos Reis, and B. Stroustrup. Open and efficient type switch for C++. In *OOPSLA '12*.

[49] J.-P. Talpin and P. Jouvelot. Polymorphic type, region and effect inference. *JFP '92*, 2(3), .

[50] J.-P. Talpin and P. Jouvelot. The type and effect discipline. *Inf. Comput. '94*, 111, .

[51] S. Treichler, M. Bauer, and A. Aiken. Language support for dynamic, hierarchical data partitioning. In *OOPSLA '13*.

[52] A. Welc, S. Jagannathan, and A. Hosking. Safe Futures for Java. In *OOPSLA '05*.

A. Open Parameters and Local Variables

So far, open effects have been discussed only in terms of open fields. However, open effects are not limited to open fields and can support open parameters and open local variables as well, especially when object fields flow into parameters or local variables, as illustrated in Figure 13. In this figure, the method Apply, has an open parameter g, used as the receiver for the invocation of the method op. This causes the effects of the method Apply to be **open**$(g\ op\ \gamma)$. Later this open effect is concretized to **open**$(f\ op\ \gamma)$, lines 5–7 where this.f is passed to Apply as the parameter g. For the open variable var, line 10, the open effect **open**$(var\ op\ \gamma)$ is generated for each method invocation on lines 12 and 13 which is concretized to **open**$(par\ op\ \gamma)$ later because the local variable var is set to the parameter par, on line 10.

```
1 private final int Apply (@open Op g, int x) {
2    g.op(x)
3 }
4 int apply() {
5    // f and g are aliases.
6    fst = Apply(this.f, fst);
7    snd = Apply(this.f, snd)
8 }

9 int apply(Op par) {
10   @open Op var = par;
11   fork {
12     var.op(1),
13     var.op(2)
14   }
15 }
```

Figure 13. Open parameter g, and open variable var.

B. *OpenEffectJ*'s Semantics

Figure 14 shows the rest of the *OpenEffectJ*'s typing rules omitted from §2. The rule (T-PROGRAM) says that a program type checks if all its declarations type check. The rule (T-CLASS) says that a class declaration type checks if all the newly declared fields are not fields of its super class, checked by the auxiliary function *validF*, its super class d is defined in the class table *CT*, checked by the auxiliary function *isClass*; and finally, all its declared methods type check. The typing rules in Figure 14 are mostly standard.

Figure 15 shows the auxiliary functions used in the typing rules. The auxiliary function *override*, used in (T-METHOD) requires that an overriding and overriden method in a subtype and its supertype have compatible types for their parameters and return values. The operation $\sqcap$ computes the intersection of two aliasing environments, i.e. $A_1 \sqcap A_2$ returns a map containing the aliasing information that exists in both A_1 and A_2.

Figure 16 shows the dynamic semantics rules that were omitted from §3 along with their auxiliary functions in Figure 17. The evaluation relation $\Sigma \overset{\eta}{\hookrightarrow} \Sigma'$ says that during the

(T-PROGRAM)
$$\frac{\forall decl \in \overline{decl}.\ \vdash decl \rightsquigarrow decl' : \text{OK} \qquad \vdash e \rightsquigarrow e' : (t, \sigma, A)}{\vdash \overline{decl}\ e \rightsquigarrow \overline{decl'}\ e' : (t, \sigma, A)}$$

(T-CLASS)
$$\frac{\forall [\text{@open}]\ t\ f \in \overline{field}.\ validF(f,d) \qquad isClass(d)}{\forall meth \in \overline{meth}.\ \vdash meth \rightsquigarrow meth' : (t', \sigma, A)\ \textbf{in}\ c}$$
$$\vdash \textbf{class}\ c\ \textbf{extends}\ d\ \{\ \overline{field}\ \overline{meth}\ \} \rightsquigarrow$$
$$\textbf{class}\ c\ \textbf{extends}\ d\ \{\ \overline{field}\ \overline{meth'}\ \} : \text{OK}$$

(T-GET)
$$\Pi, A \vdash \textbf{this} : c$$
$$\frac{typeOf(f) = (d, [\text{@open}]\ t) \qquad c <: d}{\Pi, A \vdash \textbf{this}.f \rightsquigarrow \textbf{this}.f : (t, \textbf{rd}(f), A)}$$

(T-NULL)
$$\frac{isClass(t)}{\Pi, A \vdash \textbf{null} : t}$$

(T-BINARY)
$$\Pi, A \vdash x_1 : int$$
$$\frac{\Pi, A \vdash x_2 : int}{\Pi, A \vdash x_1 \circ x_2 : int}$$

(T-VAR)
$$\frac{(x : t) \in \Pi}{\Pi, A \vdash x : t}$$

(T-BOOL)
$$\Pi, A \vdash b : int$$

(T-NUM)
$$\Pi, A \vdash n : int$$

(T-NEW)
$$\frac{isClass(c)}{\Pi, A \vdash \textbf{new}\ c() : c}$$

(T-LOC)
$$\frac{(loc : t) \in \Pi}{\Pi, A \vdash loc : t}$$

(T-CONDITION)
$$\Pi, A \vdash x : bool$$
$$\frac{\Pi, A \vdash e_1 \rightsquigarrow e_1' : (t, \sigma, A_1) \qquad \Pi, A \vdash e_2 \rightsquigarrow e_2' : (t, \sigma', A_2)}{\Pi, A \vdash \textbf{if}\ x\ \textbf{then}\ e_1\ \textbf{else}\ e_2 \rightsquigarrow}$$
$$\textbf{if}\ x\ \textbf{then}\ e_1'\ \textbf{else}\ e_2' : (t, \sigma \cup \sigma', A_1 \sqcap A_2)$$

Figure 14. *OpenEffectJ*'s omitted type-and-effect rules

evaluation, a configuration Σ transitions to another configuration Σ' producing the runtime memory read and write effects of η. The field and object sensitive runtime effect **read**(loc, f) repesents reading the field f of an object loc whereas **write**(loc, f) shows writing into the field. The transition relation $\Sigma \hookrightarrow \Sigma'$ represents a transition with no memory effects.

C. Soundness

C.1 Effect Refinement

To prove the type-and-effect system of *OpenEffectJ* sound, we should prove that the dynamic runtime effects of a program refine its static effects, that are computed by the typing rules. We prove this using two theorems which say that:

(i) concretized effects are a sound approximation of statically computed effects, Theorem 3.1; and

(ii) concretized effects soundly approximate runtime effects, Theorem 3.2.

$$\frac{\begin{array}{c} CT(c) = \textbf{class}\ c\ \textbf{extends}\ d\ \{\ \overline{field}\ \overline{meth}\ \} \\ \nexists meth \in \overline{meth}.\ meth = t\ \sigma\ m(\overline{v\ var})\{e\} \\ override(m,d,\overline{t} \to t) \\ (d,m) \notin dom(ST) \vee (\theta_c = ST(c,m) \wedge \theta_d = ST(d,m) \wedge \theta_c \subseteq \theta_d) \end{array}}{override(m,c,\overline{t} \to t)}$$

$$\frac{\begin{array}{c} (d,t,m\,(\overline{t\ var})\,\{e\},\sigma) = findMeth(c,m) \\ (d,m) \notin dom(ST) \vee (\theta_c = ST(c,m) \wedge \theta_d = ST(d,m) \wedge \theta_c \subseteq \theta_d) \end{array}}{override(m,c,\overline{t} \to t)}$$

$$override(m,Object,\overline{t} \to t)$$

$$\frac{\begin{array}{c} CT(c) = \textbf{class}\ c\ \textbf{extends}\ d\ \{\ \overline{field}\ \overline{meth}\ \} \\ \exists meth \in \overline{meth}.\ meth = (t,\sigma,m(\overline{t\ var})\{e\}) \end{array}}{findMeth(c,m) = (c,t,m(\overline{t\ var})\{e\},\sigma)}$$

$$\frac{\begin{array}{c} CT(c) = \textbf{class}\ c\ \textbf{extends}\ d\ \{\ \overline{field}\ \overline{meth}\ \} \\ \nexists meth \in \overline{meth}.\ meth = (t,\sigma,m(\overline{t\ var})\{e\}) \\ findMeth(d,m) = l \end{array}}{findMeth(c,m) = l}$$

$$\frac{\begin{array}{c} CT(c) = \textbf{class}\ c\ \textbf{extends}\ d\ \{\ \overline{field}\ \overline{meth}\ \} \\ \nexists field \in \overline{field}.\ field = [\textbf{@open}]\ t\ f \qquad validF(f,d) \end{array}}{validF(f,c)}$$

$$validF(f,Object)$$

$$\frac{\textbf{class}\ c\ \textbf{extends}\ d\ \{\ \overline{field}\ \overline{meth}\ \} \in CT}{isClass(c)}$$

$$\frac{isClass(t) \vee (t = int) \vee (t = bool)}{isType(t)}$$

$$\frac{\begin{array}{c} \textbf{class}\ c\ \textbf{extends}\ d\{\ \overline{field}\ \overline{meth}\ \} \in CT \\ \exists [\textbf{@open}]\ t\ f \in \overline{field} \end{array}}{typeOf(f) = (c,[\textbf{@open}]\ t)}$$

$$A_1 \sqcap A_2 = \{x = e \mid (A_1 \vdash x = e) \wedge (A_2 \vdash x = e)\}$$

Figure 15. Rest of *OpenEffectJ*'s auxiliary functions.

C.1.1 Preliminary Definitions

We first present some definitions used in the proofs of Theorem 3.1 and Theorem 3.2.

DEFINITION C.1. *(Dynamic trace) A dynamic trace $\overline{\eta}$ for an execution of a program is the sequence of dynamic effects η happening during its execution, where η can be a read effect **read**(loc,f) for field f of the object loc, or a write effect **write**(loc,f).*

(SET)
$$\frac{\begin{array}{c} [c.F.E] = \mu(loc) \qquad \mu_0 = \mu \oplus (loc \mapsto [c.(F \oplus (f \mapsto v)).E]) \\ \mu' = update(\mu_0,loc,f,v) \end{array}}{\langle \mathbb{E}[loc.f = v],\mu \rangle \xrightarrow{\textbf{write}(loc,f)} \langle \mathbb{E}[v],\mu' \rangle}$$

(GET)
$$\frac{\mu(loc) = [c.F.E] \qquad v = F(f)}{\langle \mathbb{E}[loc.f],\mu \rangle \xrightarrow{\textbf{read}(loc,f)} \langle \mathbb{E}[v],\mu \rangle}$$

(CALL)
$$\frac{\begin{array}{c} (c',t,m(\overline{t\ var})\{e\},\sigma) = findMeth(c,m) \\ [c.F.E] = \mu(loc) \qquad e' = [loc/\textbf{this},\overline{v/var}]e \end{array}}{\langle \mathbb{E}[loc.m(\overline{v})],\mu \rangle \hookrightarrow \langle \mathbb{E}[e'],\mu \rangle}$$

(BINARY)
$$\frac{v = v_1 \circ v_2}{\langle \mathbb{E}[v_1 \circ v_2],\mu \rangle \hookrightarrow \langle \mathbb{E}[v],\mu \rangle}$$

(DEFINE)
$$\langle \mathbb{E}[t\ var = v; e],\mu \rangle \hookrightarrow \langle \mathbb{E}[[v/var]e],\mu \rangle$$

(CONDITION-TRUE)
$$\langle \mathbb{E}[\textbf{if}\ true\ \textbf{then}\ e\ \textbf{else}\ e'],\mu \rangle \hookrightarrow \langle \mathbb{E}[e],\mu \rangle$$

(CONDITION-FALSE)
$$\langle \mathbb{E}[\textbf{if}\ false\ \textbf{then}\ e\ \textbf{else}\ e'],\mu \rangle \hookrightarrow \langle \mathbb{E}[e'],\mu \rangle$$

Figure 16. *OpenEffectJ*'s dynamic semantics rules.

$$\frac{CT(c) = \textbf{class}\ c\ \textbf{extends}\ d\ \{\ \overline{field}\ \overline{meth}\ \}}{methods(c) = methods(d) \cup \{m \mid (t,\sigma,m(\overline{t\ var})\{e\} \in \overline{meth}\}}$$

$$\frac{CT(c) = \textbf{class}\ c\ \textbf{extends}\ d\ \{\ \overline{field}\ \overline{meth}\ \}}{fields(c) = fields(d) \cup \{f \mid ([\textbf{@open}]t\ f) \in \overline{field}\}}$$

$$\frac{typeOf(f) = (d,[\textbf{@open}]\ int)}{default(f) = 0}$$

$$\frac{typeOf(f) = (d,[\textbf{@open}]\ bool)}{default(f) = false}$$

$$\frac{typeOf(f) = (d,[\textbf{@open}]\ c)}{default(f) = \textbf{null}}$$

Figure 17. Rest of *OpenEffectJ*'s auxiliary functions its dynamic semantics.

DEFINITION C.2. *(Static effect inclusion) A static effect ε is included in an effect set σ, which may contain open effects, written as $\varepsilon \in \sigma$, if:*

- *either $\varepsilon \in \downarrow_C (\sigma)$;*
- *or $\exists\, \mathbf{open}(f\ m\ \sigma') \in \downarrow_O (\sigma)\ \wedge\ \varepsilon \in \sigma'$.*

DEFINITION C.3. *(Dynamic effect refines static effect) A dynamic runtime effect η refines a static effect ε, written as $\eta \propto \varepsilon$, if:*

- *either $\eta = \mathbf{read}(loc, f) \wedge \varepsilon \in \{\mathbf{rd}(f), \mathbf{wr}(f)\}$;*
- *or $\eta = \mathbf{write}(loc, f) \wedge \varepsilon = \mathbf{wr}(f)$.*

In this definition, a write effect covers a read effect [9].

DEFINITION C.4. *(Static effect refinement) An effect set σ' refines another effect set σ if $\sigma' \subseteq \sigma$.*

DEFINITION C.5. *(Effect equivalent stores) Two stores μ and μ' are effect equivalent, written as $\mu \cong \mu'$, if:*

- *$dom(\mu) \subseteq dom(\mu')$; and*
- *$\forall loc \in \mu, \mu(loc) = [c.F.E] \Rightarrow \mu'(loc) = [c.F'.E]$, for some F'.*

DEFINITION C.6. *(Well-formed object) An object record $o = [c.F.E]$ is a well-formed in μ, written as $\mu \vdash o$, if for all open effects $\mathbf{open}(f\ m\ \sigma_0) \in \sigma \in rng(E)$:*

- *either $(F(f) = loc) \wedge (\mu(loc) = [c'.F'.E']) \wedge (E'(m) \subseteq \sigma_0)$;*
- *or $(F(f) = \mathbf{null}) \wedge (\sigma_0 = \emptyset)$;*
- *or $(typeOf(f) = (c, int))$;*
- *or $(typeOf(f) = (c, bool))$.*

DEFINITION C.7. *(Well-formed location) A location loc is well-formed in the store μ, written $\mu \vdash loc$, if:*

- *either $\mu(loc) = [c.F.E]$, $\forall m \in dom(E)$. $findMeth(c,m) = (c',t,m(\overline{t\ var})\{e\}, \sigma') \wedge (\mu, \emptyset) \vdash [loc/\mathbf{this}]e : (\sigma, A)$, then $\sigma \subseteq E(m)$;*
- *or $\mu(loc) = \mathbf{null}$.*

DEFINITION C.8. *(Well-formed store) A store μ is well-formed, written as $\mu \vdash \diamond$, if $\forall o \in rng(\mu)$. $\mu \vdash o$ and $\forall loc \in dom(\mu)$. $\mu \vdash loc$.*

Effect Concretization Figure 18 shows the rules for computation of concretized effects for *OpenEffectJ*'s expressions. In this figure, the effect judgement $(\mu, A) \vdash e : (\sigma, A')$ says that the expression e in a runtime configuration $\langle \mu, e \rangle$ with store μ and the aliasing environment A, has the concretized effect σ. The rule (E-CALL-OPEN) uses the *concretize* auxiliary function in Figure 8 for concretization of effects of a dynamically dispatched method invocation $x_0.m(\overline{x})$. The rules (E-GET) and (E-SET) assign a concretized effect $\mathbf{read}(loc, f)$ and $\mathbf{write}(loc, f)$ to the field read and write expressions. For other expressions, e.g. (T-DEFINE), their effects is the union of the concretized effects of their subexpressions.

Theorem 3.1: (Concretized effects refine static effects)
Given an expression e with statically computed effects σ_s,

which could contain open effects, and its dynamic concretization σ_c, i.e. $\sigma_c = concretize(\sigma_s)$, if $\Pi, A \vdash e \rightsquigarrow e' : (t, \sigma_s, A')$ holds statically and $(\mu, A) \vdash e' : (\sigma_c, A')$ dynamically for the runtime configuration $\langle e', \mu \rangle$, then $\sigma_c \subseteq \sigma_s$.

Proof: The proof is by a straightforward structural induction on the derivation of $\Pi, A \vdash e \rightsquigarrow e' : (t, \sigma, A')$ and $(\mu, A) \vdash e' : (\sigma', A')$.

1. For the base cases (GET), (SET), (SET-OPEN), (VAR), (NULL), (BOOL), (NUM), (NEW), (LOC), (BINARY), (CALL), (CALL-PURE), with no subexpressions, it is obvious that the effects are the same in the typing rules, §2 and Figure 14, and the effect judgment rules, Figure 18.

The remaining cases cover the induction step. The induction hypothesis (IH) is that the claim of the lemma holds for all sub-derivations of the derivation being considered.

2. (IF).

$$\Pi, A \vdash x \rightsquigarrow x : (bool, \emptyset, A)$$
$$\Pi, A \vdash e_0 \rightsquigarrow e'_0 : (t, \sigma_0, A_0)$$
$$\Pi, A \vdash e_1 \rightsquigarrow e'_1 : (t, \sigma_1, A_1)$$
$$\overline{\Pi, A \vdash \mathbf{if}\ x\ \mathbf{then}\ e_0\ \mathbf{else}\ e_1 \rightsquigarrow}$$
$$\mathbf{if}\ x\ \mathbf{then}\ e'_0\ \mathbf{else}\ e'_1 : (t, \sigma_0 \cup \sigma_1, A_0 \sqcap A_1)$$

$$(\mu, A) \vdash x : (\emptyset, A)$$
$$\frac{(\mu, A) \vdash e'_0 : (\sigma'_0, A_0) \qquad (\mu, A) \vdash e'_1 : (\sigma'_1, A_1)}{(\mu, A) \vdash \mathbf{if}\ x\ \mathbf{then}\ e'_0\ \mathbf{else}\ e'_1 : (\sigma'_0 \cup \sigma'_1, A_0 \sqcap A_1)}$$

By IH, $\sigma_0 \subseteq \sigma'_0$ and $\sigma_1 \subseteq \sigma'_1$. Therefore $(\sigma' = \sigma'_0 \cup \sigma'_1) \subseteq (\sigma_0 \cup \sigma_1 = \sigma)$.

3. (DEFINE).

$$\Pi, A \vdash e_1 \rightsquigarrow e'_1 : (t_1, \sigma_1, A_1)$$
$$\frac{\Pi; x : t, A_1; x = e'_1 \vdash e_2 \rightsquigarrow e'_2 : (t_2, \sigma_2, A_2) \qquad t_1 <: t}{\Pi, A \vdash t\ x = e_1; e_2 \rightsquigarrow t\ x = e'_1; e'_2 : (t_2, \sigma_1 \cup \sigma_2, A_2)}$$

$$\frac{(\mu, A) \vdash e'_1 : (\sigma'_1, A1) \qquad (\mu, A1; x = e'_1) \vdash e'_2 : (\sigma'_2, A2)}{(\mu, A) \vdash t\ x = e'_1; e'_2 : (\sigma'_1 \cup \sigma'_2, A2)}$$

By IH, $\sigma'_1 \subseteq \sigma_1$ and $\sigma'_2 \subseteq \sigma_2$. Therefore $(\sigma' = \sigma'_1 \cup \sigma'_2) \subseteq (\sigma_1 \cup \sigma_2 = \sigma)$. ∎

Theorem 3.2: (Dynamic effects refine concretized effects) [7] *For two configurations $\Sigma = \langle e, \mu \rangle$ and $\Sigma' = \langle e', \mu' \rangle$, if Σ transitions to Σ' producing runtime effect η, i.e. $\Sigma \xrightarrow{\eta} \Sigma'$, if the store μ is well-formed, i.e. $\mu \vdash \diamond$, and concretized effects of e is σ_c, i.e. $(\mu, A) \vdash e : (\sigma_c, A')$, then there is a concretized effect σ'_c such that:*

(a) $(\mu', A_1) \vdash e' : (\sigma'_c, A'_1)$ and $\sigma'_c \subseteq \sigma_c$;

(b) $\eta \propto \sigma_c$

We first state few lemmas which are used in the proof of the theorem.

---

[7] The theorem in §3 is the simplified version of the theorem presented here.

(E-CALL-OPEN)

$$\frac{A \vdash x = loc.f \qquad \mathbf{open}(f\ m\ \sigma) = concretize(\mu, loc, \mathbf{open}(f\ m\ \gamma)) \qquad typeOf(f) = (d, \mathbf{@open}\ c_0)}{(\mu, A) \vdash x.m(\bar{x}) : (\sigma, \emptyset)}$$

(E-CALL-LOC)
$$\frac{\mu(loc) = [c.F.E] \qquad E(m) = \sigma}{(\mu, A) \vdash loc.m(\bar{x}) : (\sigma, \emptyset)}$$

(E-CALL)
$$(\mu, A) \vdash x.m(\bar{x}) : (\top, \emptyset)$$

(E-GET)
$$(\mu, A) \vdash x.f : (\mathbf{rd}(f), A)$$

(E-GET-LOC)
$$(\mu, A) \vdash loc.f : (\mathbf{rd}(f), A)$$

(E-SET-OPEN)
$$\frac{typeOf(f) = (c, \mathbf{@open}\ c_0) \qquad A' = A \backslash f \cup \{x = \mathbf{this}.f\}}{(\mu, A) \vdash x.f = x' : (\top, A')}$$

(E-SET-OPEN-LOC)
$$\frac{typeOf(f) = (c, \mathbf{@open}\ c_0) \qquad A' = A \backslash f \cup \{x = \mathbf{this}.f\}}{(\mu, A) \vdash loc.f = loc' : (\top, A')}$$

(E-SET)
$$\frac{typeOf(f) = (c, t) \qquad A' = A \backslash f \cup \{x' = x.f\}}{(\mu, A) \vdash x.f = x' : (\mathbf{wr}(f), A')}$$

(E-SET-LOC)
$$\frac{typeOf(f) = (c, t) \qquad A' = A \backslash f \cup \{x = loc.f\}}{(\mu, A) \vdash loc.f = loc' : (\mathbf{wr}(f), A')}$$

(E-NEW)
$$(\mu, A) \vdash \mathbf{new}\ c() : (\emptyset, A)$$

(E-VAR)
$$(\mu, A) \vdash var : (\emptyset, A)$$

(E-NULL)
$$(\mu, A) \vdash \mathbf{null} : (\emptyset, A)$$

(E-LOC)
$$(\mu, A) \vdash loc : (\emptyset, A)$$

(E-DEFINE)
$$\frac{(\mu, A) \vdash e_1 : (\sigma_1, A1) \qquad (\mu, A1; x = e_1) \vdash e_2 : (\sigma_2, A_2)}{(\mu, A) \vdash t\ x = e_1; e_2 : (\sigma_1 \cup \sigma_2, A_2)}$$

(E-BINARY)
$$(\mu, A) \vdash x_1 \circ x_2 : (\emptyset, A)$$

(E-BINARY-LOC)
$$(\mu, A) \vdash v_1 \circ v_2 : (\emptyset, A)$$

(E-NUMBER)
$$(\mu, A) \vdash n : (\emptyset, A)$$

(E-BOOL)
$$(\mu, A) \vdash b : (\emptyset, A)$$

(E-CONDITION)
$$\frac{(\mu, A) \vdash e_0 : (\sigma_0, A_0) \qquad (\mu, A) \vdash e_1 : (\sigma_1, A_1)}{(\mu, A) \vdash \mathbf{if}\ x\ \mathbf{then}\ e_0\ \mathbf{else}\ e_1 : (\sigma_0 \cup \sigma_1, A_0 \sqcap A_1)}$$

Figure 18. *OpenEffectJ*'s effect concretization rules.

LEMMA C.9. *(Store preservation) Let the initial configuration of a program with a main expression e be $\Sigma_\star = \langle e, \bullet \rangle$. If $\langle e, \bullet \rangle \overset{\overline{\eta}}{\hookrightarrow}{}^* \langle e', \mu' \rangle$, then $\mu' \vdash \diamond$.*

Proof: The proof is by cases on the reduction step. In each case we show that $\mu \vdash \diamond$ implies that $\mu' \vdash \diamond$.

1. The cases (CONDITION-TRUE), (CONDITION-FALSE), (GET), (CALL), (BINARY) and (DEFINE), are trivial, because they do not change the store, i.e., $\mu' = \mu$.

   For all the remaining cases, to see $\mu' \vdash loc$, consider the definition of *initE*. It returns the effects computed by the static type-and-effect system, while the effect judgment is more accurate (Figure 18), i.e., by observation if $((\overline{var : t}, \mathbf{this} : c), \emptyset) \vdash e : (u, \sigma, A)$ and $(\mu, \emptyset) \vdash [loc/\mathbf{this}]e : (\sigma', A)$, then $\sigma' \subseteq \sigma$, therefore $\mu' \vdash loc$. Therefore, it suffices to show all the objects o are well-formed, i.e., $\mu' \vdash o$.

2. (NEW). Here $e = \mathbb{E}[new\ c()]$, $e' = \mathbb{E}[loc]$, where $loc \notin dom(\mu)$, $\mu' = \mu \oplus \{loc \mapsto [c.\{f \mapsto default(f) \mid f \in fields(c)\}.\{m \mapsto \sigma \in initE(c)\}]\}$. The only change to the store μ is the new object o created: $[c.\{f \mapsto default(f) \mid f \in fields(c)\}.\{m \mapsto \sigma \in initE(c)\}]$. All the fields are initiated to the default values, i.e., $\{f \mapsto default(f) \mid f \in$

fields(c)}. By the definition of *initE* (§3.2), all the open effects are initiated to **null**. Therefore, $\mu' \vdash o$.

3. (SET). Here $e = \mathbb{E}[loc.f = v]$, $e' = \mathbb{E}[v]$, $\mu' = \mu \oplus (loc \mapsto o)$, and $o = [u.F \oplus (f \mapsto v).E]$, where $\mu(loc) = [u.F.E]$ and $typeOf(f) = (c, t)$ for some c and t. The field f is not an open field, and by the function *update*, it does not update any effect, and $\mu' \vdash o$.

4. (SET OPEN), Here $e = \mathbb{E}[loc.f = v]$, $e' = \mathbb{E}[v]$, where $\mu_0 = \mu \oplus (loc \mapsto [c.(F \oplus (f \mapsto v)).E])$, and $\mu' = update(\mu_0, loc, f, v)$. The proof is by observation/construction of the *update* function. Each time it updates an object, it copied the corresponding effects of updated object and put it in the open effect (see the *concretize* function Figure 8). ∎

LEMMA C.10. *(Stationary effect) Let e be an expression, and μ and μ' two effect equivalent stores, i.e. $\mu \cong \mu'$, then e has the same effects in the two stores μ and μ'. In other words if $(\mu, A) \vdash e : (\sigma, A')$, then $(\mu', A) \vdash e : (\sigma, A')$.*

Proof: The proof is by induction on the structure of the expression e.

1. Cases of (NEW), (NULL), (LOC), (NUMBER), (BOOL), (BINARY) and (VAR) are trivial, since in these cases, $\sigma' = \sigma = \emptyset$.

For the remaining steps, the induction hypothesis (IH) says that the claim of the lemma holds for all sub-derivations of the derivation being considered.

2. The cases for (CONDITION), (DEFINE), (GET) and (SET) follow directly from IH.

3. (IF).

$$\frac{(\mu,A) \vdash x : (\emptyset,A) \quad (\mu,A) \vdash e_0 : (\sigma_0,A_0) \quad (\mu,A) \vdash e_1 : (\sigma_1,A_1)}{(\mu,A) \vdash \mathbf{if}\ x\ \mathbf{then}\ e_0\ \mathbf{else}\ e_1 : (\sigma_0 \cup \sigma_1, A_0 \sqcap A_1)}$$

$$\frac{(\mu',A) \vdash x : (\emptyset,A) \quad (\mu',A) \vdash e_0 : (\sigma_0',A_0) \quad (\mu',A) \vdash e_1 : (\sigma_1',A_1)}{(\mu',A) \vdash \mathbf{if}\ x\ \mathbf{then}\ e_0\ \mathbf{else}\ e_1 : (\sigma_0' \cup \sigma_1', A_0 \sqcap A_1)}$$

By IH, $\sigma_0' = \sigma_0$ and $\sigma_1' = \sigma_1$. Therefore $(\sigma' = \sigma_0' \cup \sigma_1') = (\sigma = \sigma_0 \cup \sigma_1)$.

4. (DEFINE)

$$\frac{(\mu,A) \vdash e_1 : (\sigma_1,A1) \quad (\mu,A1;x = e_1) \vdash e_2 : (\sigma_2,A2)}{(\mu,A) \vdash t\ var = e_1; e_2 : (\sigma_1 \cup \sigma_2, A2)}$$

$$\frac{(\mu',A) \vdash e_1 : (\sigma_1',A1) \quad (\mu',A1;x = e_1) \vdash e_2 : (\sigma_2',A2)}{(\mu',A) \vdash t\ var = e_1; e_2 : (\sigma_1' \cup \sigma_2', A2)}$$

By IH, $\sigma_1' = \sigma_1$ and $\sigma_2' = \sigma_2$. Therefore $(\sigma' = \sigma_1' \cup \sigma_2') = (\sigma = \sigma_1 \cup \sigma_2)$.

5. (GET)

$$(\mu,A) \vdash loc.f : (\mathbf{rd}(f),A) \qquad (\mu',A) \vdash loc.f : (\mathbf{rd}(f),A)$$

Therefore $\sigma' = \sigma = \mathbf{rd}(f)$.

6. (SET)

$$\frac{(\mu,A) \vdash x : (\emptyset,A) \quad typeOf(f) = (c,t) \quad A' = A \backslash f \cup \{x = loc.f\}}{(\mu,A) \vdash loc.f = x : (\mathbf{wr}(f),A')}$$

$$\frac{(\mu',A) \vdash x : (\emptyset,A) \quad typeOf(f) = (c,t) \quad A' = A \backslash f \cup \{x = loc.f\}}{(\mu',A) \vdash loc.f = x : (\mathbf{wr}(f),A')}$$

Thus $\sigma' = \sigma = \mathbf{wr}(f)$.

7. (SET-OPEN)

$$\frac{typeOf(f) = (c, @\mathbf{open}\ t) \quad A' = A \backslash f \cup \{x = loc.f\}}{(\mu,A) \vdash e_0.f = e_1 : (\top,A')} \qquad \frac{typeOf(f) = (c, @\mathbf{open}\ t) \quad A' = A \backslash f \cup \{x = loc.f\}}{(\mu',A) \vdash e_0.f = e_1 : (\top,A')}$$

Thus $\sigma' = \sigma = \top$.

8. (CALL-OPEN)

$$\frac{A \vdash x = loc.f \quad \mathbf{open}(f\ m\ \sigma_0') = concretize(\mu,loc,\mathbf{open}(f\ m\ \gamma)) \quad typeOf(f) = (d, @\mathbf{open}\ c_0)}{(\mu,A) \vdash x.m(\bar{x}) : (\sigma_0',\emptyset)}$$

$$\frac{A \vdash x = loc.f \quad \mathbf{open}(f\ m\ \sigma_1') = concretize(\mu',loc,\mathbf{open}(f\ m\ \gamma)) \quad typeOf(f) = (d, @\mathbf{open}\ c_0)}{(\mu',A) \vdash x.m(\bar{x}) : (\sigma_1',\emptyset)}$$

By IH, $(\sigma' = \sigma_0') = (\sigma = \sigma_1')$.

9. (CALL-LOC)

$$\frac{\mu(loc) = [c.F.E] \quad E(m) = \sigma_0}{(\mu,A) \vdash loc.m(\bar{x}) : (\sigma_0,\emptyset)} \qquad \frac{\mu'(loc) = [c.F.E] \quad E(m) = \sigma_0'}{(\mu',A) \vdash loc.m(\bar{x}) : (\sigma_0',\emptyset)}$$

Since $\mu \cong \mu'$, the effect maps E are the same and $\sigma_0 = \sigma_0'$. Thus $(\sigma' = \sigma_0') = (\sigma = \sigma_0)$.

10. (CALL)

$$(\mu,A) \vdash x.m(\bar{x}) : (\top,\emptyset) \qquad (\mu',A) \vdash x.m(\bar{x}) : (\top,\emptyset)$$

Thus $\sigma' = \sigma = \top$. $\blacksquare$

LEMMA C.11. *(Replacement with subeffect)*

If $\mu \vdash \diamond$, $\Sigma \xrightarrow{\eta} \Sigma'$, $\Sigma = \langle \mathbb{E}[e], \mu \rangle$, $\Sigma' = \langle \mathbb{E}[e'], \mu' \rangle$, $(\mu,A) \vdash \mathbb{E}[e] : (\sigma,A')$, $(\mu,A) \vdash e : (\sigma_0,A_0')$, $(\mu,A) \vdash e' : (\sigma_1,A_0')$, $\mu \cong \mu'$, and $\sigma_1 \subseteq \sigma_0$, then $(\mu,A) \vdash \mathbb{E}[e'] : (\sigma',A') \wedge \sigma' \subseteq \sigma$.

For two expression e and e', in the configurations $\Sigma = \langle \mathbb{E}[e], \mu \rangle$ and $\Sigma' = \langle \mathbb{E}[e'], \mu' \rangle$ such that $\Sigma \xrightarrow{\eta} \Sigma'$, if the store μ is well-formed, i.e. $\mu \vdash \diamond$, and the expression e has the effects σ_0, i.e. $(\mu,A) \vdash e : (\sigma_0,A_0')$ and $\mathbb{E}[e]$ has the concrete effects σ, i.e. $(\mu,A) \vdash \mathbb{E}[e] : (\sigma,A')$, and $\mu \cong \mu'$, and $\sigma_1 \subseteq \sigma_0$, then $(\mu,A) \vdash \mathbb{E}[e'] : (\sigma',A') \wedge \sigma' \subseteq \sigma$.

Lemma C.11 says that given two effect equivalent stores, and the same evaluation context, if the effect of the subsequent expression e' refines the original expression e, then the effect of the entire subsequent expression $\mathbb{E}[e']$ refines the entire original expression $\mathbb{E}[e]$.

Proof: The proof is by induction on the size of the evaluation context $\mathbb{E}$. The size of the $\mathbb{E}$ is the number of recursive applications of the syntactic rules necessary to create $\mathbb{E}$.

1. For the base case $\mathbb{E} = -$, the size of $\mathbb{E}$ is zero , and $(\sigma' = \sigma_1) \subseteq (\sigma = \sigma_0)$.

For the induction step we divide the evaluation context into two parts such that $\mathbb{E}[e_1] = \mathbb{E}_1[\mathbb{E}_2[e_2]]$, and $\mathbb{E}_2$ has the size one. The induction hypothesis (IH) says that the lemma holds for all evaluation contexts, which their sizes are smaller than the one ($\mathbb{E}_1$) considered in the induction step. We prove it case by case on the rule used to generate

$\mathbb{E}_2$. In each case we show that $(\mu, A) \vdash \mathbb{E}_2[e] : (\sigma, A')$ implies that $(\mu', A) \vdash \mathbb{E}_2[e'] : (\sigma', A')$, for some $\sigma' \subseteq \sigma$, and thus the claim holds by the IH.

2. For (E-DEFINE), (E-GET) and (E-SET) the proof follows directly from the IH.

3. (E-SET-OPEN) holds because in this case $\sigma = \top$. ∎

LEMMA C.12. *(Substitution effect) If* $(\mu, A) \vdash e : (\sigma, A')$, *then there is some* σ', *such that* $(\mu, A) \vdash [\overline{v/var}]e : (\sigma', A')$, *for all values v in* $\overline{v}$ *and free variables var in* $\overline{var}$, *and* $\sigma' \subseteq \sigma$.

Proof: The proof is by structural induction on the derivation of $(\mu, A) \vdash e : (\sigma, A')$ and by cases, based on the last step in that derivation.

1. Proof for (E-NEW), (E-NULL), (E-LOC) is trivial, since e has no variables, $\sigma' = \sigma = \emptyset$.

2. For (E-VAR), $(\mu, A) \vdash v : (\emptyset, A)$ and $(\mu, A) \vdash var = (\emptyset, A)$.

The remaining cases cover the induction step. The induction hypothesis (IH) is that the claim of the lemma holds for all sub-derivations of the derivation being considered.

3. For (E-CONDITION), (E-BINARY), (E-GET) and (E-DEFINE) the proof follows directly from the IH.

4. The case for (E-SET-OPEN) and (E-CALL) hold because in these cases $\top \in \sigma$. The effect of e is $\top$ and every effect refines $\top$.

5. (E-CALL-OPEN)

$$\frac{\begin{array}{c} A \vdash x = loc.f \\ \mathbf{open}(f\ m\ \sigma) = concretize(\mu, loc, \mathbf{open}(f\ m\ \gamma)) \\ typeOf(f) = (d, \mathbf{@open}\ c_0) \end{array}}{(\mu, A) \vdash x.m(\overline{x}) : (\sigma, \emptyset)}$$

Let $e_i' = [\overline{v/var}]x_i$ for $i \in \{1..n\}$, $[\overline{v/var}]e = x.m(\overline{e'})$. They result in the same effect by (E-CALL-OPEN).

6. (E-CALL-LOC)

$$\frac{\mu(loc) = [c.F.E] \qquad E(m) = \sigma}{(\mu, A) \vdash loc.m(\overline{x}) : (\sigma, \emptyset)}$$

Let $e_i' = [\overline{v/var}]x_i$ for $i \in \{1..n\}$, then $[\overline{v/var}]e = loc.m(\overline{v})$. Clearly $(\mu, A) \vdash [\overline{v/var}]e : (\sigma, A)$.

7. (E-SET)

$$\frac{typeOf(f) = (c, t) \qquad A' = A \backslash f \cup \{x = loc.f\}}{(\mu, A) \vdash loc.f = x : (\mathbf{wr}(f), A')}$$

Now $[\overline{v}/\overline{var}]e = (loc.f = [v/x]x)$. $(\mu, A) \vdash [v/x]x : (\emptyset, A)$. By the definition of $typeOf$, the result of $typeOf(f)$ remains unchanged, i.e. $typeOf(f) = (c, t)$. ∎

LEMMA C.13. *(Subexpression effect containment) If* $(\mu, A) \vdash e : (\sigma, A_0)$ *and* $(\mu, A) \vdash \mathbb{E}[e] : (\sigma', A_0')$, *then* $\sigma \subseteq \sigma'$.

Proof: By the effect rule for each expression, the effect of any direct subexpression is a subset of the entire expression.

C.2 Proof of Theorem 3.2

Using Lemma C.10 and Lemma C.11 To prove Theorem 3.2, in each reduction case, let $e = \mathbb{E}[e_0]$, $e' = \mathbb{E}[e_1]$, $(\mu, A) \vdash e_0 : (\sigma_0, A')$ and $(\mu', A) \vdash e_1 : (\sigma_1, A')$. Given that (a) $\mu \cong \mu'$, by Lemma C.11 and Lemma C.10, to prove (b), it suffices to prove $\sigma_1 \subseteq \sigma_0$. We divide the cases into 3 categories: in the first category, some variables (*var*) will be replaced by actual values (*v*); the cases, in the second category, access the store; and the other cases are listed right below. Here the rule leaves no dynamic trace, and (c) holds.

- (NEW) Here $e = \mathbb{E}[new\ c()]$, $e' = \mathbb{E}[loc]$, where $loc \notin dom(\mu)$, $\mu' = \mu \oplus \{loc \mapsto [c.\{f \mapsto default(f) \mid f \in fields(c)\}.\{m \mapsto \sigma \in initE(c)\}]\}$. Because this rule does not change any object, $\mu \cong \mu'$. Also $(\mu, A) \vdash \mathbf{new}\ c() : (\emptyset, A)$ and $(\mu, A) \vdash loc : (\emptyset, A)$, and (b) holds.

- (BINARY) Here $e = \mathbb{E}[v_1 \circ v_2]$, $e' = \mathbb{E}[v]$, where $v = v_1 \circ v_2$, $\mu' = \mu$. It is trivial to see that (b) holds.

Using Lemma C.12 We now present the case for method call and local declaration.

- (CALL) Here $e = \mathbb{E}[loc.m(\overline{v})]$, $(u', t_m, m(\overline{t\ var})\{e_2\}, \sigma_m) = findMeth(u, m)$, $e' = \mathbb{E}[e_1]$, $e_1 = [loc/\mathbf{this}, \overline{v/var}]e_2$, $\mu(loc) = [u.F.E]$. Let $(\mu, A) \vdash loc.m(\overline{v}) : (\sigma_0, A)$, i.e., $E(m) = \sigma_0$. Let $e_3 = [loc/\mathbf{this}]e_2$, $(\mu, A) \vdash e_3 : (\sigma_3, A_3)$ and $(\mu, A) \vdash e_1 : (\sigma_1, A_1)$. By Lemma C.12, $\sigma_1 \subseteq \sigma_3$. By $\mu \vdash \diamond$, Definition C.7 and Definition C.8, $\sigma_3 \subseteq \sigma_0$, thus $\sigma_1 \subseteq \sigma_0$.

- (DEFINE) Here $e = \mathbb{E}[t\ var = v; e_1]$, and $e' = \mathbb{E}[e_1']$, where $e_1' = [v/var]e_1$. Let $(\mu, A) \vdash e_1 : (\sigma_0, A_0)$, by (E-DEFINE), $(\mu, A) \vdash t\ var = v; e_1 : (\sigma_0, A_0)$. $(\mu, A) \vdash [v/var]e_1 : (\sigma_1, A_0)$, for some $\sigma_1 \subseteq \sigma_0$, by Lemma C.12.

Using Lemma C.13 We prove cases for field accesses:

- (GET) Here $e = \mathbb{E}[loc.f]$, $e' = \mathbb{E}[v]$, where $\mu(loc) = [u.F.E]$, $F(f) = v$, $\mu' = \mu$ and $\mu \cong \mu'$. Because $(\mu, A) \vdash loc.f : (\mathbf{rd}(f), A)$, and $(\mu', A) \vdash v : (\emptyset, A)$, (b) holds. Finally, $\eta = (\mathbf{read}(loc, f))$, and $\eta \propto \mathbf{rd}(f) \subseteq \sigma$, by Lemma C.13.

- (SET) Here $e = \mathbb{E}[loc.f = v]$, $e' = \mathbb{E}[v]$, $\mu' = \mu \oplus (loc \mapsto o)$, and $o = [u.F \oplus (f \mapsto v).E]$, where $\mu(loc) = [u.F.E]$ and $typeOf(f) = (c, t)$ for some t and c. The field is not an open field, and by the function *update*, it does not update any effect, and $\mu \cong \mu'$. To see $(\mu, A) \vdash \mathbb{E}[v] : (\sigma', A')$ and $\sigma' \subseteq \sigma$, we have $(\mu, A) \vdash loc.f = v : (\mathbf{wr}(f), A)$, and $(\mu, A) \vdash v : (\emptyset, A)$, thus $\sigma' \subseteq \sigma$. Finally, $\eta = (\mathbf{write}(loc, f))$, and $\eta \propto \mathbf{wr}(f) \subseteq \sigma$, by Lemma C.13.

- (SET OPEN) Here $e = \mathbb{E}[loc.f = v]$, $e' = \mathbb{E}[v]$, where $\mu_0 = \mu \oplus (loc \mapsto [c.(F \oplus (f \mapsto v)).E])$, and $\mu' = update(\mu_0, loc, f, v)$. The effect of e is $\top$ and every effect refines $\top$. ∎