

Types and trace effects of higher order programs

CHRISTIAN SKALKA

The University of Vermont, Burlington, Vermont, USA
(e-mail: skalka@cs.uvm.edu)

SCOTT SMITH

The Johns Hopkins University, Baltimore, Maryland, USA
(e-mail: scott@cs.jhu.edu)

DAVID VAN HORN

Brandeis University, Waltham, Massachusetts, USA
(e-mail: dvanhorn@cs.brandeis.edu)

Abstract

This paper shows how type effect systems can be combined with model-checking techniques to produce powerful, automatically verifiable program logics for higher order programs. The properties verified are based on the ordered sequence of events that occur during program execution, so-called *event traces*. Our type and effect systems infer conservative approximations of the event traces arising at run-time, and model-checking techniques are used to verify logical properties of these histories. Our language model is based on the λ -calculus. Technical results include a type inference algorithm for a polymorphic type effect system, and a method for applying known model-checking techniques to the *trace effects* inferred by the type inference algorithm, allowing static enforcement of history- and stack-based security mechanisms. A type safety result is proven for both unification and subtyping constraint versions of the type system, ensuring that statically well-typed programs do not contain trace event checks that can fail at run-time.

1 Introduction

Safe and secure program execution is crucial for modern information systems, but is difficult to attain in practice because of both programmer errors and intentional attacks. Various programming language-based techniques exist for increasing program safety, by verifying at compile- and/or run-time that programs possess certain safety properties. In addition to narrowly focused systems for verification of specific properties such as memory safety or stack inspection security, recent research has explored security paradigms for enforcing general classes of properties. One such paradigm comprises the class of properties of program *event traces* that can be expressed in temporal logics (Igarashi & Kobayashi 2002; K. Marriott & Sulzmann 2003; Skalka & Smith 2004; Skalka 2005; Bartoletti *et al.* 2005b). This paper establishes a foundation for automated verification of higher order programs in this

paradigm, using type and effect analysis. We develop a process for automatically predicting program event traces at compile-time and for statically verifying properties of traces.

Events are records of program actions, explicitly inserted into program code either manually (by the programmer) or automatically (by the compiler). Events are sufficiently abstract to represent a variety of program actions—e.g., opening a file, access control privilege activation, or entry to or exit from critical regions. Event traces maintain the ordered sequences of events that occur during program execution, and assertions enforce properties of event traces. While such properties can be anything in the abstract, in practice they must be expressed in automatically verifiable logics. Therefore, we use model-checkable linear temporal logics that express regular properties of traces Steffen & Burkart (1992).

For example, if a program is sending and receiving data over an SSL socket, the relevant events are opening and closing of sockets, and reading and writing of data packets. An example event trace produced by a program run could be

```
ssl_open("snork.cs.jhu.edu",4434); ssl_hs_begin(4434);
ssl_hs_success(4434); ssl_put(4434); ssl_get(4434);
ssl_open("moo.cs.uvm.edu",4435); ssl_hs_begin(4434);
ssl_put(4435); ssl_close(4434); ssl_close(4435)
```

Here, `ssl_open` is a sample event with two arguments, a url and a port. Event traces can then be used to detect logical flaws or security violations. For SSL, sockets must first be opened, handshake begun, handshake success, and only then can data be got/put over the socket. Thus, the above trace is illegal because data are put on socket 4435 before notification has been received that handshake was successful on that socket. Codifying this property of event traces as a local check in a decidable logic provides a rigorous definition of well-formedness, *and* allows mechanical verification of it.

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

1.1 A type-based approach

The focus of this paper is on the compile-time verification of higher order program trace properties expressible in linear temporal logic. The verification itself is based on known model-checking techniques for program abstractions, and we define a

novel type theory to infer program trace abstractions, yielding a fully automated analysis. These abstractions, called *trace effects*, are integrated directly into the type language, so that type reconstruction automates abstract interpretation of program trace behaviour. We develop both a *Hindley–Milner (HM) style* system where types are represented in term form and the implementation relies on unification, and a *subtyping constraint* system where the implementation relies on constraint closure. The flexibility of parametric polymorphism over types and effects is available in both systems. Type theory provides an especially rigorous and extensible foundation for the static analysis of trace-based program properties in a higher order language setting. The main contribution of this paper is the development and study of this foundation, and definition of a sound automated system covering both automatic extraction of program trace approximations and their verification.

To characterize the analysis and prove its correctness, the metatheory of types provides an appealing formalism. We demonstrate subject reduction, progress, and type safety results for constraint subtyping, which extend to the HM system by virtue of conservation of the latter in the former. Conservativity of both systems with respect to the underlying pure HM system is also demonstrated. Soundness and completeness of subtyping constraint inference are proved, as is soundness for unification-based inference. Completeness of constraint subtyping also entails a principal types result for the system. So-called *trace approximation* is demonstrated for both type systems, meaning that trace effects conservatively approximate run-time trace behavior.

To complete the analysis, we define and prove correct an automatic verification technique for trace effects, based on model checking in the linear μ -calculus. We also develop a variation on the language model so that stack-based trace policies can be defined. In a stack-based model, trace events are “popped” with their associated activations, allowing definition of policies such as stack inspection. The static analysis and type safety result are modified to accommodate this variation, requiring only a postprocessing transformation of inferred trace effects to implement.

An OCaml implementation of our analysis is presented and discussed in Van Horn (2006a), and is available for download Van Horn (2006b). The implementation includes both the HM style and constraint subtyping systems, and the stack-based postprocessing transformation of effects.

1.2 Organisation of the Paper

In Section 2, we introduce the language model λ_{trace} , which includes dynamic traces in configurations and dynamic checks on traces. In Section 3, we define a language of trace effects for static approximation of dynamic traces. In Section 4, we define a HM style polymorphic type system for assigning effects to programs. In Section 5, we give a unification-based inference algorithm for automatically reconstructing program effects. In Section 6, we define a more flexible constraint subtyping system, and in Section 7, we define an inference algorithm for reconstructing types in this system. In Section 8, we define an alternate stack-based version of the language and show how postprocessing inferred effects can approximate traces in this variation.

c	$\in \mathcal{C}$	atomic constants
b	$::= \text{true} \mid \text{false}$	boolean values
v	$::= x \mid \lambda_z x.e \mid c \mid b \mid \neg \mid \vee \mid \wedge \mid ()$	values
e	$::= v \mid e e \mid \mathcal{E}(e) \mid \text{if } e \text{ then } e \text{ else } e \mid \text{let } x = v \text{ in } e$	expressions
η	$::= \epsilon \mid \mathcal{E}(e) \mid \eta; \eta$	traces
E	$::= [] \mid v E \mid E e \mid \mathcal{E}(E) \mid \text{if } E \text{ then } e \text{ else } e$	evaluation contexts

Fig. 1. λ_{trace} language syntax.

In Section 9, we develop an automatic model-checking technique for verification of effects. In Section 10, we describe examples of how the system can be applied to language-based security. We conclude with a discussion of related work and final summary in Section 11.

2 The language λ_{trace}

In this section we develop the syntax and semantics of our language model λ_{trace} .

2.1 Syntax

The syntax of the theory λ_{trace} is given in Figure 1. The base values include booleans and the unit value $()$. Expressions $\text{let } x = v \text{ in } e$ are included to implement let-polymorphism in the type system (Section 4). Functions, written $\lambda_z x.e$, possess a recursive binding mechanism where z is the self-variable. We assume the following syntactic sugarings:

$$\begin{aligned}
 e_1 \wedge e_2 &\triangleq \wedge e_1 e_2 & e_1 \vee e_2 &\triangleq \vee e_1 e_2 & \lambda x.e &\triangleq \lambda_z x.e \quad z \text{ not free in } e \\
 \lambda _ . e &\triangleq \lambda x.e \quad x \text{ not free in } e & e_1; e_2 &\triangleq (\lambda _ . e_2)(e_1)
 \end{aligned}$$

Events \mathcal{E} are named entities parameterized by constants c (we treat only the unary case in this presentation, but the extension to n -ary events is straightforward). These constants $c \in \mathcal{C}$ are abstract, with \mathcal{C} formally defined as a countably infinite set of arbitrary identifiers; this set could, for example, be strings or IP addresses. Ordered sequences of these events constitute traces η , which maintain the sequence of events experienced during program execution. We let $\hat{\eta}$ denote the string obtained from this sequence by removing delimiters $()$. We distinguish a subset of events \mathcal{E}_ϕ identified by assertions ϕ in a to-be-specified logical syntax (defined in Section 9). These are *check* events, and the semantics will require that run-time checks succeed in order for computation to progress. We presuppose existence of a meaning function Π such that $\Pi(\phi(c), \hat{\eta})$ holds iff $\phi(c)$ is valid for $\hat{\eta}$; we also leave the meaning function Π abstract until later (Section 9).

Parameterizing events and checks with constants c allow for a more expressive event language; for example, in Section 10, we show how the parameterized privileges of Java stack inspection can be encoded with the aid of these parameters.

$\eta, (\lambda_z.x.e)v$	$\rightsquigarrow \eta, e[v/x][\lambda_z.x.e/z]$	(β)
$\eta, \neg \text{true}$	$\rightsquigarrow \eta, \text{false}$	(<i>notT</i>)
$\eta, \neg \text{false}$	$\rightsquigarrow \eta, \text{true}$	(<i>notF</i>)
$\eta, \wedge \text{true}$	$\rightsquigarrow \eta, \lambda x.x$	(<i>andT</i>)
$\eta, \wedge \text{false}$	$\rightsquigarrow \eta, \lambda _.\text{false}$	(<i>andF</i>)
$\eta, \vee \text{true}$	$\rightsquigarrow \eta, \lambda _.\text{true}$	(<i>orT</i>)
$\eta, \vee \text{false}$	$\rightsquigarrow \eta, \lambda x.x$	(<i>orF</i>)
$\eta, \text{if true then } e_1 \text{ else } e_2$	$\rightsquigarrow \eta, e_1$	(<i>ifT</i>)
$\eta, \text{if false then } e_1 \text{ else } e_2$	$\rightsquigarrow \eta, e_2$	(<i>ifF</i>)
$\eta, \text{let } x = v \text{ in } e$	$\rightsquigarrow \eta, e[v/x]$	(<i>let</i>)
$\eta, \text{ev}(c)$	$\rightsquigarrow (\eta; \text{ev}(c)), ()$	(<i>event</i>)
$\eta, \text{ev}_\phi(c)$	$\rightsquigarrow (\eta; \text{ev}_\phi(c)), ()$	if $\Pi(\phi(c), \hat{\eta} \text{ev}_\phi(c))$ (<i>check</i>)
$\eta, E[e]$	$\rightarrow \eta', E[e']$	if $\eta, e \rightsquigarrow \eta', e'$ (<i>context</i>)

Fig. 2. λ_{trace} language semantics.

2.2 Semantics

The operational semantics of λ_{trace} is defined in Figure 2 via the call-by-value small step reduction relations \rightsquigarrow and \rightarrow on configurations η, e , where η is the run-time program event trace. We write \rightarrow^* to denote the reflexive, transitive closure of \rightarrow . Note that in the *event* reduction rule, an event $\text{ev}(c)$ encountered during execution is added to the end of the trace. The *check* rule specifies that when a configuration $\eta, \text{ev}_\phi(c)$ is encountered during execution, it is appended to the end of η like other events, and also $\phi(c)$ is required to be satisfied by the string $(\hat{\eta} \text{ev}_\phi(c))$, which is the concatenation of $\hat{\eta}$ and $\text{ev}_\phi(c)$ (see Definition 3.1), according to our meaning function Π . The reasons for treating checks as dynamic events are manifold; for one, some checks may ensure that other checks have occurred in the trace. Also, this scheme will simplify the definition of Π , as well as typing and verification. In case a check fails at run-time, execution is “stuck”; formally:

Definition 2.1

We say that a configuration η, e is *stuck* iff e is not a value and there does not exist η' and e' such that $\eta, e \rightarrow \eta', e'$. If $\epsilon, e \rightarrow^* \eta, e'$ and η, e' is stuck, then e is said to go *wrong*.

The following example demonstrates the basics of syntax and operational semantics.

Example 2.1 (File state)

Let the function w/file be defined as

$$w/\text{file} \triangleq \lambda \text{fn}.\lambda f.\text{open}(\text{fn}); f(\text{fn}); \text{close}(\text{fn})$$

And, let *readtwice* be defined as

$$\text{readtwice} \triangleq \lambda \text{fn}.\text{read}(\text{fn}); \text{read}(\text{fn}); ()$$

where $open(fn)$, $close(fn)$, and $read(fn)$ each induce events ev_{open} , ev_{close} , ev_{read} , respectively, parameterized by the filename fn , and $readtwice$ reads twice from file fn and returns the unit value $()$, i.e., it is evaluated only for effect. Then in the operational semantics, we have

$$\epsilon, w/file \text{ f } readtwice \rightarrow^* ev_{open}(f); ev_{read}(f); ev_{read}(f); ev_{close}(f), ()$$

3 Trace Effects

The goal of our static analysis is to conservatively approximate trace behavior of programs, specifically the trace that will be generated by a program during its evaluation, and to predict success or failure of program checks on the basis of this approximation. We will use a type system to reconstruct *trace effects*, which constitute the approximation.

In essence, trace effects H conservatively approximate traces η that may develop during execution, by representing a set of traces containing at least η . Trace effects are generated by the following grammar:

$$H ::= \epsilon \mid h \mid ev(c) \mid H;H \mid H|H \mid \mu h.H \quad \text{trace effects}$$

A trace effect may be an event $ev(c)$, a sequencing of trace effects $H_1;H_2$, a nondeterministic choice of trace effects $H_1|H_2$, or a recursively bound trace effect $\mu h.H$ that finitely represents the set of possibly infinite traces that may be generated by recursive functions. Trace effects may contain check events $ev_\phi(c)$, allowing verification of checks at the right points in trace approximations.

We now make precise the way in which a trace effect approximates the run-time trace of a program. We define a Labelled Transition System (LTS) interpretation of trace effects as sets of strings over the alphabet of events plus a \downarrow symbol to denote termination; abusing terminology, we also call these strings traces. Note that traces arising from program evaluation may be infinite, because programs may not terminate, but traces arising from the interpretation of trace effects will always be finite.

Definition 3.1

We write θ to denote possibly \downarrow terminated strings over the alphabet of events:

$$s ::= ev(c) \mid \epsilon \mid s s$$

$$\theta ::= s \mid s \downarrow$$

We endow strings with an equational theory to interpret ϵ as the empty string and string concatenation as usual— more precisely, for all s, s_1, s_2 , and s_3 the following equations hold:

$$s \epsilon = s \qquad \epsilon s = s \qquad (s_1 s_2) s_3 = s_1 (s_2 s_3)$$

The symbol Θ is defined to range over prefix-closed sets of traces θ .

Trace effects generate traces by viewing them as programs in a simple nondeterministic transition system.

Definition 3.2

The trace effect transition relation on closed trace effects is defined as follows:

$$\begin{aligned} ev(c) \xrightarrow{w(c)} \epsilon & \quad H_1|H_2 \xrightarrow{\epsilon} H_1 & \quad H_1|H_2 \xrightarrow{\epsilon} H_2 & \quad \mu h.H \xrightarrow{\epsilon} H[\mu h.H/h] \\ \epsilon; H \xrightarrow{\epsilon} H & \quad H_1; H_2 \xrightarrow{s} H'_1; H_2 & \text{ if } H_1 \xrightarrow{s} H'_1 \end{aligned}$$

We formally determine the sets of traces Θ associated with a closed trace effect in terms of the transition relation:

Definition 3.3

The interpretation of trace effects is defined as follows:

$$\llbracket H \rrbracket = \{s_1 \dots s_n \mid H \xrightarrow{s_1} \dots \xrightarrow{s_n} H'\} \cup \{s_1 \dots s_n \downarrow \mid H \xrightarrow{s_1} \dots \xrightarrow{s_n} \epsilon\}$$

Any trace effect interpretation is clearly prefix-closed. In this interpretation, an infinite trace is viewed as the set of its finite prefixes.

Note that prefix closure does not cause any loss of information, since the postpending of \downarrow to terminating traces allows them to be distinguished from their prefixes. In particular, this means that $(H_1; H_2) \neq H_1$ for arbitrary closed H_1 and $H_2 \neq \epsilon$.

Equivalence of trace effects is defined via their interpretation; i.e., $H_1 = H_2$ iff $\llbracket H_1 \rrbracket = \llbracket H_2 \rrbracket$. This relation is, in fact, undecidable: traces are equivalent to Basic Process Algebras (BPAs), as demonstrated in Section 9, and equivalence of BPAs is known to be undecidable (Burkart *et al.* 2001).

Some trace effects contain occurrences of checks; i.e., their interpretation contains traces of the form $\theta ev_\phi(c)$. We define *validity* of trace effects in terms of the satisfiability of these checks, given their context θ ; the check must hold for its immediate prefix:

Definition 3.4

A trace effect H is *valid* iff for all $\theta ev_\phi(c) \in \llbracket H \rrbracket$, $\Pi(\phi(c), \theta ev_\phi(c))$ holds.

3.1 Properties

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

Lemma 3.1

We note the following properties of trace effect equivalence for all closed H , H_1, H_2 , and H_3 :

1. $H|H = H$
2. $\epsilon; H = H = H; \epsilon$
3. $\mu h.H = H$
4. $H_1|H_2 = H_2|H_1$
5. $H_1; (H_2; H_3) = (H_1; H_2); H_3$
6. $H_1|(H_2; H_3) = (H_1|H_2); H_3$
7. $H_1; (H_2|H_3) = (H_1; H_2)|(H_1; H_3)$

$\alpha \in \mathcal{V}_{\text{Sing}}, t \in \mathcal{V}_{\text{Type}}, h \in \mathcal{V}_{\text{Eff}}$	type variables
$\beta \in \mathcal{V}_{\text{Sing}} \cup \mathcal{V}_{\text{Type}} \cup \mathcal{V}_{\text{Eff}}$	
$s ::= c \mid \alpha$	singleton types
$H ::= \epsilon \mid h \mid \text{ev}(s) \mid H; H \mid H \mid H \mid \mu h.H$	trace effects
$\tau ::= \beta \mid s \mid \{s\} \mid H \mid \tau \xrightarrow{H} \tau \mid \text{bool} \mid \text{unit}$	types
$\sigma ::= \forall \bar{\beta}. \tau$	type schemes
$\Gamma ::= \emptyset \mid \Gamma; x : \sigma$	type environments

Fig. 3. λ_{trace} type syntax.

8. $(H_1 \mid H_2); H_3 = (H_1; H_3) \mid (H_2; H_3)$
9. $H'[\mu h.H'/h] = \mu h.H'$ for all closed $\mu h.H'$.
10. Trace effect equivalence is homomorphic for all constructors

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

Lemma 3.2

Writing $H \subseteq H'$ iff $\llbracket H \rrbracket \subseteq \llbracket H' \rrbracket$, the following properties hold:

1. $H \subseteq H'$ is undecidable
2. $H \subseteq H \mid H'$
3. If $H \subseteq H'$ then $H \mid H'' \subseteq H' \mid H''$ and $H''; H \subseteq H''; H'$ and $H; H'' \subseteq H'; H''$ for all closed H'' .
4. If $H \subseteq H'$ then validity of H' implies validity of H .

Finally, noting that the syntax of traces η is the same as linear, variable-free trace effects, we abuse syntax and let η also range over linear, variable-free trace effects, interpreting traces η as the identical trace effect. This syntax is used, e.g., in the statement of subject reduction, Lemmas 4.2 and 6.10.

4 Hindley–Milner Style Typing for λ_{trace}

The syntax of types for λ_{trace} is given in Figure 3. This syntax is more liberal than what is actually allowed in type derivations; e.g., function domain and range types cannot be trace effects or “bare” singletons s , and the three different kinds of variables must occur only in the appropriate positions. Formally, we define as *well-kinded* only those type terms that have an interpretation in the ground tree model given in the sense of Definitions 5.2 and 5.3, below. Hereafter we restrict our presentation to well-kinded types. The syntax of types includes forms for booleans, unit, and function types of the form $\tau_1 \xrightarrow{H} \tau_2$, where *latent* effects H represent the traces that may result by use of the function. Events are side effects, and so these function types are a form of effect type (Talpin & Jouvelot 1992; Amtoft et al. 1999).

$\frac{\text{VAR}}{\Gamma(x) = \sigma} \quad \frac{\Gamma, \epsilon \vdash x : \sigma}{\Gamma, \epsilon \vdash x : \sigma}$	$\frac{\text{CONST}}{\Gamma, \epsilon \vdash \mathbf{c} : \Delta(\mathbf{c})}$	$\frac{\text{EVENT}}{\Gamma, H \vdash e : \{s\}} \quad \frac{\Gamma, H; \epsilon w(s) \vdash \epsilon w(e) : \text{unit}}$
$\frac{\text{\forall-INTRO}}{\Gamma, \epsilon \vdash v : \tau} \quad \frac{\bar{\beta} \# \text{fv}(\Gamma)}{\Gamma, \epsilon \vdash v : \forall \bar{\beta}. \tau}$	$\frac{\text{\forall-ELIM}}{\Gamma, \epsilon \vdash v : \forall \bar{\beta}. \tau} \quad \frac{\Gamma, \epsilon \vdash v : \tau[\bar{\tau}/\bar{\beta}]}{\Gamma, \epsilon \vdash v : \tau[\bar{\tau}/\bar{\beta}]}$	$\frac{\text{WEAKEN}}{\Gamma, H' \vdash e : \tau} \quad \frac{H' \sqsubseteq H}{\Gamma, H \vdash e : \tau}$
$\frac{\text{IF}}{\Gamma, H_1 \vdash e_1 : \text{bool} \quad \Gamma, H_2 \vdash e_2 : \tau \quad \Gamma, H_2 \vdash e_3 : \tau} \quad \frac{\Gamma, H_1; H_2 \vdash \text{if } e_1 \text{ then } e_2 \text{ else } e_3 : \tau}$		
$\frac{\text{APP}}{\Gamma, H_1 \vdash e_1 : \tau' \xrightarrow{H_3} \tau \quad \Gamma, H_2 \vdash e_2 : \tau'} \quad \frac{\Gamma, H_1; H_2; H_3 \vdash e_1 e_2 : \tau}$	$\frac{\text{FIX}}{\Gamma; x : \tau; z : \tau \xrightarrow{H} \tau', H \vdash e : \tau'} \quad \frac{\Gamma, \epsilon \vdash \lambda_2 x. e : \tau \xrightarrow{H} \tau'}$	
$\frac{\text{LET}}{\Gamma, \epsilon \vdash v : \sigma \quad \Gamma; x : \sigma, H \vdash e : \tau} \quad \frac{\Gamma, H \vdash \text{let } x = v \text{ in } e : \tau}$		

Fig. 4. λ_{trace} Hindley-Milner style typing rules.

In addition, since events and predicates are parameterized, we must be especially accurate with respect to our typing of singleton constants. Thus, we adopt a very simple form of singleton type $\{c\}$ (Stone 2000), where only atomic constants can have singleton type.

Types contain three kinds of variables: regular-type variables t , singleton-type variables α , and trace effect-type variables h . The metavariable β ranges over all kinds of variables. Universal type schemes $\forall \bar{\beta}. \tau$ bind any kind of type variable in τ , where $\bar{\beta}$ is a vector of type variables. Note that let-polymorphism over types, singletons, and trace effects are included in our system. We formally define some convenient notation for vectors of variables:

Definition 4.1

Variable vectors, denoted $\bar{\beta}$, range over sequences of distinct variables $\beta_1 \dots \beta_n$, with the empty vector denoted \emptyset . Vectors are equivalent up to reordering; hence, we may treat vectors $\beta_1 \dots \beta_n$ as analogous sets $\{\beta_1, \dots, \beta_n\}$, in particular adapting notation $\beta \in \bar{\beta}$ and $\bar{\beta}_1 \cup \bar{\beta}_2$ and $\bar{\beta}_1 \cap \bar{\beta}_2$ with the obvious meaning. We write $\bar{\beta}_1 \# \bar{\beta}_2$ iff $\bar{\beta}_1 \cap \bar{\beta}_2 = \emptyset$. For any type scheme σ , we write $\text{fv}(\sigma)$ to denote the free variables in σ , extending the notation to environments Γ in the obvious manner. We write τ as syntactic sugar for $\forall \emptyset. \tau$.

Without loss of generality, we equate type schemes up to α -renaming and extraneous bindings. That is, we assume the following axiom:

$$\forall \bar{\beta}_1 \bar{\beta}_2. \tau = \forall \bar{\beta}_1. \tau \text{ if } \bar{\beta}_2 \# \text{fv}(\tau)$$

Source code type derivation rules for judgements $\Gamma, H \vdash e : \tau$ are given in Figure 4, where Γ is an environment of variable-typing assumptions, H is the effect of the

expression e , and τ is the type of e . To assign types to constants, we posit a constant type binding environment Δ as follows.

Definition 4.2

Constants \mathbf{c} range over the set $\{(), \text{true}, \text{false}, \neg, \vee, \wedge\} \cup \mathcal{C}$. The binding environment Δ contains bindings for all constants, consisting of $c : \{c\}$ for all $c \in \mathcal{C}$, and the following:

$$\begin{aligned} () &: \text{unit} & \text{true} &: \text{bool} & \text{false} &: \text{bool} & \neg &: \text{bool} \xrightarrow{\epsilon} \text{bool} \\ \wedge &: \text{bool} \xrightarrow{\epsilon} \text{bool} \xrightarrow{\epsilon} \text{bool} & \vee &: \text{bool} \xrightarrow{\epsilon} \text{bool} \xrightarrow{\epsilon} \text{bool} \end{aligned}$$

Type safety will be defined in terms of *valid* typing judgments, which are derivable judgments where the top-level effect is valid. A subject reduction result will show that effects assigned by type analysis are conservative approximations of program trace behavior while validity of top-level effects will guarantee success of run-time checks.

Definition 4.3

A derivable type judgement $\Gamma, H \vdash e : \sigma$ is *valid* iff H is valid.

Example 4.1

Let w/file and readtwice be defined as in Example 2.1. Then the following judgments are derivable:

$$\begin{aligned} \emptyset, \epsilon \vdash w/\text{file} : \forall fn, h, \alpha. \{fn\} \xrightarrow{\epsilon} (\{fn\} \xrightarrow{h} \alpha) \xrightarrow{e_{\text{open}}; h; e_{\text{close}}} \alpha \\ \emptyset, \epsilon \vdash \text{readtwice} : \forall fn. \{fn\} \xrightarrow{e_{\text{read}}(fn); e_{\text{read}}(fn)} \text{unit} \\ \emptyset, e_{\text{open}}(\mathbf{f}); e_{\text{read}}(\mathbf{f}); e_{\text{read}}(\mathbf{f}); e_{\text{close}}(\mathbf{f}) \vdash w/\text{file} \ \mathbf{f} \ \text{readtwice} : \text{unit} \end{aligned}$$

4.1 Properties

One of the basic claims about our type system is that it is *conservative*, in that the addition of trace effects is a conservative extension to an underlying HM-style let-polymorphic type system: by using weakening before each if-then-else typing, any derivation in the underlying effect-free type system may be replayed here. This result is formalized as follows.

Definition 4.4

We define the “event erasure” function erase_η on expressions as follows:

$$\begin{aligned} \text{erase}_\eta(x) &= x \\ \text{erase}_\eta(\mathbf{c}) &= \mathbf{c} \\ \text{erase}_\eta(e(c)) &= () \\ \text{erase}_\eta(e_1 e_2) &= \text{erase}_\eta(e_1) \text{erase}_\eta(e_2) \\ \text{erase}_\eta(\lambda_z x. e) &= \lambda_z x. \text{erase}_\eta(e) \\ \text{erase}_\eta(\text{if } e_1 \text{ then } e_2 \text{ else } e_2) &= \text{if } \text{erase}_\eta(e_1) \text{ then } \text{erase}_\eta(e_2) \text{ else } \text{erase}_\eta(e_2) \\ \text{erase}_\eta(\text{let } x = e_1 \text{ in } e_2) &= \text{let } x = \text{erase}_\eta(e_1) \text{ in } \text{erase}_\eta(e_2) \end{aligned}$$

We also define the “trace effect erasure” function erase_H on types as follows:

$$\begin{aligned}\text{erase}_H(\tau) &= \tau \quad \tau \in \{t, \text{bool}, \text{unit}, \{s\}\} \\ \text{erase}_H(\tau_1 \xrightarrow{H} \tau_2) &= \text{erase}_H(\tau_1) \rightarrow \text{erase}_H(\tau_2)\end{aligned}$$

The function erase_H is extended to type environments Γ in the obvious manner.

Definition 4.5

We write $\Gamma \vdash e : \tau$ for judgments in a standard simple type system with let-polymorphism, e.g., that of Pierce (2002), extended with singleton constants and our CONST rule.

Lemma 4.1 (Conservativity)

$\Gamma, H \vdash e : \tau$ is derivable iff $\text{erase}_H(\Gamma) \vdash \text{erase}_\eta(e) : \text{erase}_H(\tau)$ is.

Proof

Straightforward by definition and induction on typing derivations. \square

Our type safety results are stated as follows. The proof of each is immediate by conservation of the HM-style system in the constraint subtyping system presented in Section 6, Lemma 6.3, and analogous results in that system, which are proven in detail in Section 6.

Theorem 4.1 (Type Safety)

If $\Gamma, H \vdash e : \tau$ is valid for closed e , then e does not go wrong.

Theorem 4.2 (Progress)

If $\Gamma, H \vdash e : \tau$ is derivable for closed e and η, e is irreducible with $\eta; H$ valid, then e is a value.

A subject reduction result for the HM-style system can also be proved via Lemma 6.3 and subject reduction in the constraint subtyping system.

Lemma 4.2 (Subject Reduction)

If $\Gamma, H \vdash e : \tau$ is derivable for closed e and $\eta, e \rightsquigarrow \eta', e'$, then $\Gamma, H' \vdash e' : \tau$ is derivable with $\eta'; H' \sqsubseteq \eta; H$.

A significant corollary of this result is one of our guiding intuitions: that trace effects conservatively approximate the set of possible run-time traces, formalised as follows. Prefix closure of effect interpretation ensures that traces produced at any point in the computation, not just at termination, will be approximated.

Corollary 4.1

If $\Gamma, H \vdash e : \tau$ is derivable for closed e and $\epsilon, e \rightarrow^* \eta, e'$ then $\hat{\eta} \in \llbracket H \rrbracket$.

5 Unification-based type inference

In this section we show that a type inference algorithm exists for the HM-style type system for λ_{trace} presented in the previous section. The type inference algorithm will be partly based on unification, in keeping with the HM subset of the system, but the presence of effect weakening requires a simple form of constraint solution. Thus,

$$\begin{array}{c}
\text{VAR} \\
\frac{\Gamma(x) = \forall \bar{\beta}. \tau / C}{\Gamma, \epsilon \vdash_{\bar{\beta}'} x : (\tau / C) [\bar{\beta}' / \bar{\beta}]} \\
\\
\text{CONST} \\
\Gamma, \epsilon \vdash_{\emptyset} \mathbf{c} : \Delta(\mathbf{c}) / \mathbf{true} \\
\\
\text{EVENT} \\
\frac{\Gamma, H \vdash_{\bar{\beta}} e : \tau / C}{\Gamma, H; \mathbf{ev}(\alpha) \vdash_{\bar{\beta} \cup \{\alpha\}} \mathbf{ev}(e) : \mathbf{unit} / C \wedge \tau \sqsubseteq \{\alpha\}} \\
\\
\text{IF} \\
\frac{\bar{\beta} = \bar{\beta}_1 \bar{\beta}_2 \bar{\beta}_3 \cup \{t\} \quad \Gamma, H_1 \vdash_{\bar{\beta}_1} e_1 : \tau_1 / C_1 \quad \Gamma, H_2 \vdash_{\bar{\beta}_2} e_2 : \tau_2 / C_2 \quad \Gamma, H_3 \vdash_{\bar{\beta}_3} e_3 : \tau_3 / C_3}{\Gamma, H_1; H_2 | H_3 \vdash_{\bar{\beta}} \mathbf{if} e_1 \mathbf{then} e_2 \mathbf{else} e_3 : t / C_1 \wedge C_2 \wedge C_3 \wedge \tau_1 \sqsubseteq \mathbf{bool} \wedge \tau_2 \sqsubseteq t \wedge \tau_3 \sqsubseteq t} \\
\\
\text{APP} \\
\frac{\bar{\beta}_1 \# \bar{\beta}_2 \quad \Gamma, H_1 \vdash_{\bar{\beta}_1} e_1 : \tau_1 / C_1 \quad \Gamma, H_2 \vdash_{\bar{\beta}_2} e_2 : \tau_2 / C_2}{\Gamma, H_1; H_2; h \vdash_{\bar{\beta}_1 \bar{\beta}_2 \cup \{t, h\}} e_1 e_2 : t / C_1 \wedge C_2 \wedge \tau_1 \sqsubseteq \tau_2 \xrightarrow{h} t} \\
\\
\text{FIX} \\
\frac{\Gamma; x : t; z : t \xrightarrow{h} t', H \vdash_{\bar{\beta}} e : \tau / C}{\Gamma, \epsilon \vdash_{\bar{\beta} \cup \{t, h, t'\}} \lambda z. x. e : t \xrightarrow{h} t' / C \wedge \tau \sqsubseteq t' \wedge H \sqsubseteq h} \\
\\
\text{LET} \\
\frac{\bar{\beta}_1 \# \bar{\beta}_2 \quad \Gamma, \epsilon \vdash_{\bar{\beta}_1} v : \tau' / D \quad \Gamma; x : \forall \bar{\beta}_1. \tau' / D, H \vdash_{\bar{\beta}_2} e : \tau / C}{\Gamma, H \vdash_{\bar{\beta}_1 \bar{\beta}_2} \mathbf{let} x = v \mathbf{in} e : \tau / C \wedge D}
\end{array}$$

Fig. 5. λ_{trace} type constraint inference rules.

our approach to inference will be constraint based, with type constraints interpreted as equality constraints, and effect constraints interpreted as containment constraints. In addition to being an effective approach in this context, we will see in Section 7 that it allows the same type inference algorithm to be used for a subtyping system, by imposing a subtyping interpretation of type constraints, and introducing a new constraint solution algorithm. That is, type inference is modular with respect to interpretation of constraints, in the spirit of systems such as HM(X) Sulzmann (2001).

5.1 Syntax and meaning of constraints

We begin by defining the syntax of constraints, where \sqsubseteq is the constraint relation between types.

Definition 5.1

Constraints C are defined as follows:

$$C ::= \mathbf{true} \mid \tau \sqsubseteq \tau \mid C \wedge C$$

Constraints are required to agree in kind, i.e. for any constraint $\tau \sqsubseteq \tau'$, we require that $\tau : k$ and $\tau' : k$ for some k .

$$\begin{array}{lll}
\rho(h, hs) & = & \rho(h) & h \notin hs \\
\rho(h, hs) & = & h & h \in hs \\
\rho(ev(s), hs) & = & ev(\rho(s)) \\
\rho(\epsilon, hs) & = & \epsilon \\
\rho(H_1; H_2, hs) & = & \rho(H_1, hs); \rho(H_2, hs) \\
\rho(H_1|H_2, hs) & = & \rho(H_1, hs)|\rho(H_2, hs) \\
\rho(\mu h.H, hs) & = & \mu h.\rho(H, hs \cup \{h\}) \\
\\
\rho(c) & = & c \\
\rho(\{s\}) & = & \{\rho(s)\} \\
\rho(unit) & = & unit \\
\rho(bool) & = & bool \\
\rho(\tau_1 \xrightarrow{H} \tau_2) & = & \rho(\tau_1) \xrightarrow{\rho(H)} \rho(\tau_2) \\
\rho(H) & = & \rho(H, \emptyset)
\end{array}$$

Fig. 6. Interpretations extended to types and effects.

The appropriate meaning of constraints is obtained by interpretation in a model \mathbb{T} , which for the HM system we define as a universe of monotypes, endowed with a partial order \preceq that is an equality relation on types, and a containment relation on effects. Later, to obtain a subtyping interpretation of constraints, we will redefine \mathbb{T} and \preceq appropriately.

Definition 5.2 (Ground Type Model)

Let \mathbb{T} be the set of *ground*, or variable free, types τ , denoted $\hat{\tau}$ and generated by the following grammar, where \hat{H} ranges over variable free trace effects:

$$\hat{\tau} ::= unit \mid bool \mid \{c\} \mid \hat{H} \mid \hat{\tau} \xrightarrow{\hat{H}} \hat{\tau}$$

Furthermore, \mathbb{T} is endowed with a partial order \preceq axiomatized as follows:

$$\frac{\llbracket H \rrbracket \subseteq \llbracket H' \rrbracket}{H \preceq H'}$$

Definition 5.3 (Interpretation of Constraints)

Interpretations ρ are total mappings from type variables β to \mathbb{T} . We impose the sanity condition that each kind of type variable is mapped to elements of \mathbb{T} of the appropriate form, i.e. for all α, t, h, ρ we require that $\rho(\alpha)$ is a singleton c , and $\rho(h)$ is a closed effect H , and $\rho(t)$ is one of the other element forms. Interpretations are extended to types and effects as in Figure 6, where in abuse of notation the interpretation of effects is parameterized by sets hs of effect variables, to prevent substitution of μ -bound variables. The relation $\rho \vdash C$, pronounced ρ *satisfies* or *solves* C , is axiomatized as follows:

$$\begin{array}{lll}
\rho \vdash \mathbf{true} & \frac{\rho(\tau_1) \preceq \rho(\tau_2)}{\rho \vdash \tau_1 \sqsubseteq \tau_2} & \frac{\rho \vdash C \quad \rho \vdash D}{\rho \vdash C \wedge D}
\end{array}$$

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

5.2 Type inference and constraint solution

Type inference judgments are of the form $\Gamma, H \vdash_{\bar{\beta}} e : \tau / C$, where H , τ , and C are the reconstructed top-level effect, type, and constraint, and $\bar{\beta}$ are the variables introduced during the derivation. Type schemes are now constrained, written $\forall \bar{\beta}. \tau / C$, and type environments Γ now bind variables to constrained type schemes. We write τ as shorthand for $\forall \emptyset. \tau / \mathbf{true}$. Type inference rules are defined in Figure 5; they are deterministic except for arbitrary choice of variables. We call *canonical* those derivations where fresh variables are chosen whenever possible, where a variable β is fresh with respect a judgement $\Gamma, H \vdash_{\bar{\beta}} e : \tau / C$ iff it occurs nowhere in the judgement. Note that disjointness conditions on fresh variables chosen in subderivation are imposed in derivations, by the IF, APP, and LET rules, ensuring “global” freshness of variable choice in canonical derivations.

To automatically obtain HM-style types for λ_{trace} expressions, it is also necessary to convert inferred types and constraints into unified, constraint-free form. To accomplish this, we define a unification algorithm, extended to also apply to effect constraints. Given a constraint generated by type inference, the resulting algorithm will generate a type substitution that solves the constraint.

Definition 5.4

Substitutions, written $[\tau_1/\beta_1, \dots, \tau_n/\beta_n]$, are mappings from type variables to types, extended to types, type schemes, type environments, and constraints in the usual manner. We write $\text{dom}([\tau_1/\beta_1, \dots, \tau_n/\beta_n])$ to mean β_1, \dots, β_n . We let ψ range over substitutions $[\tau_1/\beta_1, \dots, \tau_n/\beta_n]$, although the former is prefix notation while the latter is postfix. We write $\psi_1 \circ \psi_2$ to denote the substitution ψ such that $\psi(\tau) = \psi_1(\psi_2(\tau))$ for all τ . Extending the vector notation of Definition 4.1, we write $[\bar{\tau}/\bar{\beta}]$ to denote $[\tau_1/\beta_1, \dots, \tau_n/\beta_n]$, where $\bar{\tau} = \tau_1, \dots, \tau_n$ and $\bar{\beta} = \beta_1, \dots, \beta_n$.

Since any solution of a constraint C must unify the type constraints in C , and also satisfy the effect constraints in C , constraint solutions need to be defined generally, as follows:

Definition 5.5

A substitution ψ is a *solution* of a constraint C iff $\rho \vdash \psi(C)$ for all ρ .

To define how to construct solutions of constraint, here and for the subtyping version of our type analysis presented later, it is useful to interpret constraints as sets, rather than conjunctions, of atomic constraints. Thus, we introduce the following notation.

Definition 5.6

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

$$\hat{C} ::= \mathbf{true} \mid \tau \sqsubseteq \tau$$

$$\begin{aligned}
\text{soln}(C) &= \text{let } C_1, C_2 = \text{hsplit}(C) \text{ and } \psi_1 = \text{unify}(C_1) \text{ in } \text{soln}_{\mathbf{H}}(\psi_1(C_2)) \circ \psi_1 \\
\text{hsplit}(C) &= C_1, C_2 \text{ where } \text{set}(C_2) = \{H_1 \sqsubseteq H_2 \mid H_1 \sqsubseteq H_2 \in C\} \\
&\quad \text{and } C_1 = C - C_2 \\
\text{bounds}(h, C) &= H_1 | \cdots | H_n \text{ where } \{H_1, \dots, H_n\} = \{H \mid H \sqsubseteq h \in C\} \\
\text{soln}_{\mathbf{H}}(\emptyset) &= \emptyset \\
\text{soln}_{\mathbf{H}}(C) &= \text{let } \psi = [(\mu h. \text{bounds}(h, C))/h] \text{ in} \\
&\quad \text{soln}_{\mathbf{H}}(\psi(C - \{H \sqsubseteq h \mid H \sqsubseteq h \in C\})) \circ \psi
\end{aligned}$$

Fig. 7. Unification-based solution.

$$\begin{aligned}
\text{unify}(\mathbf{true}) &= \mathbf{true} \\
\text{unify}(C \wedge \tau \sqsubseteq \tau) &= \text{unify}(C) \\
\text{unify}(C \wedge \beta \sqsubseteq \tau) &= \mathbf{fail} \text{ if } \beta \in \text{fv}(\tau), \text{ else} \\
&\quad \text{unify}(C[\tau/\beta]) \circ [\tau/\beta] \\
\text{unify}(C \wedge \tau \sqsubseteq \beta) &= \mathbf{fail} \text{ if } \beta \in \text{fv}(\tau), \text{ else} \\
&\quad \text{unify}(C[\tau/\beta]) \circ [\tau/\beta] \\
\text{unify}(C \wedge \{s_1\} \sqsubseteq \{s_2\}) &= \text{unify}(C \wedge s_1 \sqsubseteq s_2) \\
\text{unify}(C \wedge \tau_1 \xrightarrow{h} \tau_2 \sqsubseteq \tau'_1 \xrightarrow{h'} \tau'_2) &= \text{unify}(C \wedge h \sqsubseteq h' \wedge \tau'_1 \sqsubseteq \tau_1 \wedge \tau_2 \sqsubseteq \tau'_2)
\end{aligned}$$

Fig. 8. Constraint set unification.

and given $C = \hat{C}_1 \wedge \cdots \wedge \hat{C}_n$, let $\text{set}(C) = \{\hat{C}_1, \dots, \hat{C}_n\}$. Then:

$$\begin{aligned}
C \sqsubseteq D &\iff \text{set}(C) \sqsubseteq \text{set}(D) \\
\hat{C} \in C &\iff \hat{C} \in \text{set}(C) \\
C_1 - C_2 = D &\iff \text{set}(C_1) - \text{set}(C_2) = \text{set}(D)
\end{aligned}$$

Now, since effect containment is known to be undecidable, as observed in Section 3, it seems there can be no general constraint solution algorithm. However, type inference generates effect constraints of a particular form, in particular it yields a system of lower bounds $H \sqsubseteq h$ on effect variables h , as is immediately demonstrated by observation of the type inference rules in Figure 5. For this restricted form of effect constraints, there exists a solution algorithm $\text{soln}_{\mathbf{H}}$, defined in Figure 7, that joins lower bounds in the solution of effect variables h . By composing this with the unification algorithm unify defined in Figure 8, a constraint solution algorithm soln is obtained for constraints generated by type inference. Note that inferred constraints need to be split into type and effect components by hsplit , before applying unify to the former and $\text{soln}_{\mathbf{H}}$ to the latter. Soundness of the type inference technique can be stated as follows; a proof of this result is given in Van Horn (2006a).

Theorem 5.1 (Soundness of Inference)

Given closed e ; then if $\emptyset, H \vdash_{\beta} e : \tau/C$ is derivable, and $\psi = \text{soln}(C)$, then $\emptyset, \psi(H) \vdash e : \psi(\tau)$ is derivable.

6 Constraint subtyping for λ_{trace}

In this section we propose a more expressive system—whereby a strictly larger class of valid programs is verifiable—based on subtyping constraints. The HM-style system developed in the previous section suffers from a certain degree of imprecision due to the interpretation of \sqsubseteq as an equivalence relation on types. This imprecision results in program effect characterizations that are sound but not as tight as can be achieved through subtyping. In general, tighter effect characterizations result in the verifiability of more programs. Consider the following example and its typing under the HM style system:

$$f \triangleq \lambda x.(\text{if true then } \lambda_{\dots}ev_1(c) \text{ else } x); x$$

Observe that the effect of x and $\lambda_{\dots}ev_1(c)$ must be weakened into equivalence as required by the typing rule for conditionals; hence, we derive

$$f : \forall ht.(t \xrightarrow{ev_1(c)|h} \text{unit}) \xrightarrow{\epsilon} (t \xrightarrow{ev_1(c)|h} \text{unit})$$

In other words, the effect of $\lambda_{\dots}ev_1(c)$ is forced to be subsumed by the effect of x , resulting in a lack of precision, further illustrated by application of f :

$$f(\lambda_{\dots}ev_2(c)) : t \xrightarrow{ev_1(c)|ev_2(c)} \text{unit}$$

Although this expression evaluates to $\lambda_{\dots}ev_2(c)$, which clearly has only the effect $ev_2(c)$ when applied, the latent effect of this function is given as $ev_1(c) \mid ev_2(c)$ by the type system. Supposing that the function is applied with a subsequent check for the occurrence of $ev_2(c)$, the program would fail to validate since $ev_2(c)$ *may*—but not *must*—occur according to this characterization of the program’s effect.

To remedy the situation and obtain a more expressive system, the subtyping relation will allow types, as well as effects, to be weakened as necessary. In treating the above example, the type of x will not be forced to be equivalent to $\lambda_{\dots}ev_1(c)$, but can be weakened via subsumption to allow typing of the conditional expression in the body of f , while retaining a purely abstract latent effect on x in a most general typing of f . The essence of the subtyping relation is an adaptation of effect weakening to latent effect on function types, for example:

$$t \xrightarrow{h} \text{unit} \sqsubseteq t \xrightarrow{ev_1(c)|h} \text{unit}$$

Thus we can derive a more precise typing for f :

$$f : \forall ht.(t \xrightarrow{h} \text{unit}) \xrightarrow{\epsilon} (t \xrightarrow{h} \text{unit})$$

and also for $f(\lambda_{\dots}ev_2(c))$:

$$f(\lambda_{\dots}ev_2(c)) : t \xrightarrow{ev_2(c)} \text{unit}$$

Again supposing the resulting function is applied with a subsequent check for $\mathcal{E}_2(c)$, the program would validate since this tighter characterization of the program's effect appropriately gives that $\mathcal{E}_2(c)$ *must* occur. Taken together with the conservativity result demonstrated in Lemma 6.3, which states that any HM-style derivation may be replayed in the subtyping system, this shows that a strictly larger class of valid programs is verifiable under the constraint subtyping system as compared to the HM-style unification-based system of the previous section.

Subtleties of the relation include contravariance (covariance) in function domain (range) types, respectively, and treatment of free variables in related types; these issues are resolved in Section 6.2 below.

The use of a constraint representation of types has several distinct benefits resulting from the greater precision of constraint types. For one, constraints may be recursive, allowing typing of self-referential expressions such as $\lambda x.xx$. Furthermore, constraints allow a representation of conjunctive and disjunctive types Eifrig *et al.* (1995), without requiring conjunctive and disjunctive type syntax and interpretation. While disjunction can be represented in effect terms via the nondeterministic choice operator (\mid), a term representation of effect conjunction would require the addition of a parallel and operation in effects, which is known to significantly increase the complexity of model-checking Burkart *et al.* (2001). In short, a constraint representation simplifies the subtyping analysis; while it yields a more cumbersome, less readable typing than a unified representation, model-checkable effects can still be extracted, as we show in Section 7. Since our goal is an automatic program analysis tool, this is agreeable.

6.1 Logical subtyping judgments and derivations

Logical subtyping judgments are of the form $\Gamma, C, H \vdash e : \tau$, with all syntactic forms as defined in previous sections. In any such judgement, we refer to C and H as the *top-level* constraint and effect. The constant type binding environment Δ is also as defined previously. Note that constraints C , previously occurring only in inference judgments, now occur in logical judgments; of course, a significant difference in the current system is the interpretation of \sqsubseteq as a subtyping relation, formalized below in Section 6.2. Note well that we adopt the interpretation of constraints given in Definition 5.3; subtyping will be realized purely through a redefinition of the model \mathbb{T} in Definition 6.3.

To allow more flexibility in the typing of singleton constants, we make a small extension to the type and effect language:

Definition 6.1

We extend the language of singleton types in Figure 3 with the form $s_1|s_2$. In addition, we define the semantics of an effect $\mathcal{E}(c_1|c_2)$ as equivalent to $\mathcal{E}(c_1)|\mathcal{E}(c_2)$.

Anticipating the interpretation, the meaning of a type $s_1|s_2$ subsumes the meaning of s_1 and s_2 —i.e., it is a disjunctive singleton type.

While the logical typing rules are based on previous subtyping constraint systems such as Eifrig *et al.* (1995) and Skalka and Pottier (2003), there is a difference, in

$\frac{\text{VAR} \quad \Gamma(x) = \sigma}{\Gamma, C, \epsilon \vdash x : \sigma}$	$\frac{\text{CONST}}{\Gamma, C, \epsilon \vdash \mathbf{c} : \Delta(\mathbf{c})}$	$\frac{\text{EVENT} \quad \Gamma, C, H \vdash e : \{s\}}{\Gamma, C, H; \omega(s) \vdash \omega(e) : \text{unit}}$
$\frac{\text{\forall-INTRO} \quad \Gamma, C, \epsilon \vdash v : \tau \quad \bar{\beta} \# \text{fv}(\Gamma)}{\Gamma, C, \epsilon \vdash v : \forall \bar{\beta}. \tau}$	$\frac{\text{\forall-ELIM} \quad \Gamma, C, \epsilon \vdash v : \forall \bar{\beta}. \tau \quad C \Vdash [\bar{\tau}/\bar{\beta}]}{\Gamma, C, \epsilon \vdash v : \tau[\bar{\tau}/\bar{\beta}]}$	
$\frac{\text{SUB} \quad \Gamma, C, H' \vdash e : \tau' \quad C \Vdash \tau' \sqsubseteq \tau \quad C \Vdash H' \sqsubseteq H}{\Gamma, C, H \vdash e : \tau}$		
$\frac{\text{IF} \quad \Gamma, C, H_1 \vdash e_1 : \text{bool} \quad \Gamma, C, H_2 \vdash e_2 : \tau \quad \Gamma, C, H_2 \vdash e_3 : \tau}{\Gamma, C, H_1; H_2 \vdash \text{if } e_1 \text{ then } e_2 \text{ else } e_3 : \tau}$		
$\frac{\text{APP} \quad \Gamma, C, H_1 \vdash e_1 : \tau' \xrightarrow{H} \tau \quad \Gamma, C, H_2 \vdash e_2 : \tau'}{\Gamma, C, H_1; H_2; H \vdash e_1 e_2 : \tau}$	$\frac{\text{FIX} \quad \Gamma; x : \tau'; z : \tau' \xrightarrow{H} \tau, C, H \vdash e : \tau}{\Gamma, C, \epsilon \vdash \lambda_z x. e : \tau' \xrightarrow{H} \tau}$	
$\frac{\text{LET} \quad \Gamma, C, \epsilon \vdash v : \sigma \quad \Gamma; x : \sigma, C, H \vdash e : \tau}{\Gamma, C, H \vdash \text{let } x = v \text{ in } e : \tau}$		

Fig. 9. λ_{trace} constraint typing, logical rules.

that in any judgement $\Gamma, C, H \vdash e : \tau$, the scope of quantified type variables $\bar{\beta}$ in type schemes in Γ effectively extend to C ; in particular, noting the form of the rule \forall -ELIM, we define:

Definition 6.2

Write $C \Vdash [\bar{\tau}/\bar{\beta}]$ iff $C \Vdash C[\bar{\tau}/\bar{\beta}]$.

Thus, given a top-level constraint $C \triangleq h_1 \sqsubseteq h_2$, the type scheme $\forall h_1 h_2. \text{unit} \xrightarrow{h_1} \text{unit}$ is more general than $\forall h_1. \text{unit} \xrightarrow{h_1} \text{unit}$, even though h_2 does not occur in the quantified type. Observe that given the former type scheme, the rule \forall -ELIM essentially imposes monomorphism on h_1 due to its relation with h_2 specified in C , whereas the latter type scheme allows any instantiation $[H_1/h_1]$ as long as there also exists a substitution $[H_2/h_2]$ such that $C \Vdash H_1 \sqsubseteq H_2$. This state of affairs is implicitly due to a simplification we have made in the type system, where all relevant typing constraints are subsumed in the top-level constraint C . In the type inference system presented in Section 7, where constraints must be reconstructed, rather than given *a priori*, we will return to constrained type scheme forms.

The logical subtyping derivation rules defined in Figure 9 are quite similar to the HM-style rules defined in Figure 4, other than the rule SUB, which incorporates type subsumption as well as weakening of top-level effects. As mentioned above, type subsumption allows weakening of latent effects on function types. The meaning

$\top^\emptyset, \perp^\emptyset : \text{Type}$	$\text{unit}^\emptyset : \text{Type}$	$\text{bool}^\emptyset : \text{Type}$	$\text{sc}^\emptyset : \text{Sing}$	$\{\text{sc}\}^\emptyset : \text{Type}$
$H^\emptyset : \text{Eff}$		$(\cdot \xrightarrow{\cdot})^{\text{Type}, \text{Eff}, \text{Type}} : \text{Type}$		

Fig. 10. Kinding rules for regular tree constructors.

of the subtyping relation and validity of type judgments are defined in terms of constraint solutions, formalized in the next section.

6.2 Interpretation of subtyping constraints

Following Trifonov and Smith (1996), subtyping is defined via interpretation in a regular tree model endowed with a *primitive subtyping* relation, a technique that allows us to accommodate recursive constraints. The relation is defined inductively via finite approximations. Intuitively, the subtyping relation is invariant on *bool* and *unit* types, is contravariant (resp. covariant) on the domain (resp. range) types of functions, and is covariant on latent function effects. As before, the relation \sqsubseteq on trace effects is defined via set containment in the LTS interpretation of related effects.

Definition 6.3 (Regular Tree Model)

Let the *tree constructor kinds* be defined as

$$k ::= \text{Type} \mid \text{Eff} \mid \text{Sing}$$

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

$$\begin{aligned} \text{sc} & ::= c \mid \text{sc}|\text{sc} \\ \text{tc} & ::= \top \mid \perp \mid \text{sc} \mid \text{unit} \mid \text{bool} \mid \{\text{sc}\} \mid H \mid \cdot \xrightarrow{\cdot} \end{aligned}$$

where each element of the alphabet is indexed by a signature, written tc^ς , and must be *well-kinded* according to the rules given in Figure 10.

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

A partial order over \mathbb{T} is then defined via an approximate relation over finite $\varphi \in \mathbb{T}$. First, define a *level- n cut* $\varphi \upharpoonright_n$ for $\varphi \in \mathbb{T}$ as the finite tree obtained by replacing all subtrees at level n of φ with \top . Then, \preceq_{fin} is the partial order over finite $\varphi \in \mathbb{T}$ axiomatized in Figure 11, and \preceq is the partial order over \mathbb{T} approximated by \preceq_{fin} axiomatized in Figure 11.

By retaining Definition 5.3 with \mathbb{T} and \preceq refigured in this manner, and extending interpretations as defined in Figure 6 to disjunctive singleton types as $\rho(s_1|s_2) =$




Fig. 11. Primitive subtyping for regular trees.

$\rho(s_1) | \rho(s_2)$, we obtain a subtyping interpretation of constraints. We immediately note that this imposes transitivity and reflexivity on \sqsubseteq :

Lemma 6.1

The relation \sqsubseteq is transitive and reflexive, by which we mean $C \Vdash \tau \sqsubseteq \tau$, and $C \Vdash \tau_1 \sqsubseteq \tau_2 \wedge \tau_2 \sqsubseteq \tau_3$ implies $C \Vdash \tau_1 \sqsubseteq \tau_3$.

It also is useful to characterize what sort of types cannot be placed in a \sqsubseteq relation, especially to establish a canonical forms result (Lemma 6.14).

Lemma 6.2

Let \asymp be the least symmetric relation on types (pronounced “clashes with”) axiomatised by the following relational schemas:

$$\begin{aligned} \text{unit} \asymp \text{bool} \quad \text{unit} \asymp \{s\} \quad \text{unit} \asymp \tau_1 \xrightarrow{H} \tau_2 \quad \text{bool} \asymp \{s\} \quad \text{bool} \asymp \tau_1 \xrightarrow{H} \tau_2 \\ \{s\} \asymp \tau_1 \xrightarrow{H} \tau_2 \end{aligned}$$

If C is solvable then for all τ_1, τ_2 such that $\tau_1 \asymp \tau_2$ it is the case that $C \not\Vdash \tau_1 \sqsubseteq \tau_2$.

Given our interpretation of constraints, we can now define validity of type judgments. Note that top-level effects may contain free variables that must be interpreted via top-level constraints in order to ascertain validity of expression effects.

Definition 6.4

A judgement $\Gamma, C, H \vdash e : \tau$ is *satisfiable* iff it is derivable and C has a solution. The judgement is *valid* iff it is satisfiable and there exists a solution ρ of C such that $\rho(H)$ is valid.

6.3 Properties

In this section we demonstrate our type safety and progress results, which along with correctness of type reconstruction will constitute the main results of this paper. The result proceeds via a subject reduction argument, modified to account for trace effects. Type safety is explicitly demonstrated for the constraint subtyping system; to allow application to the unified system of Section 4, we demonstrate the following conservativity result. Note that this also allows conservativity of HM typings à la Lemma 4.1 to be extended to constraint subtyping.

Lemma 6.3

If $\Gamma, H \vdash e : \tau$ is derivable, then so is $\Gamma, \mathbf{true}, H \vdash e : \tau$.

Proof

Straightforward by induction and case analysis on the last rule instance in the derivation of $\Gamma, H \vdash e : \tau$. Each derivation step in $\Gamma, H \vdash e : \tau$ can be mimicked by a same-named rule instance in the constraint subtyping system, other than instances of WEAKEN, which can be simulated by instances of SUB, by reflexivity of \sqsubseteq and since it is easy to show that $H \sqsubseteq H'$ implies $C \Vdash H \sqsubseteq H'$ for all C . \square

In later proofs the following observations about effect subtyping will be important; since interpretation of effect subtyping is defined in terms of containment in the LTS semantics defined in Section 3, the stated properties follow by Lemma 3.2.

Lemma 6.4

The following properties hold:

1. $C \Vdash H \sqsubseteq H|H'$
2. If $C \Vdash H \sqsubseteq H'$ then $C \Vdash H|H'' \sqsubseteq H'|H''$ and $C \Vdash H'';H \sqsubseteq H'';H'$ and $C \Vdash H;H'' \sqsubseteq H';H''$.

On the way to proving type safety and progress, we demonstrate a suite of standard results for the current type system. This includes weakening and instantiation; both properties follow in a similar manner to corresponding results in Skalka and Pottier (2003):

Lemma 6.5 (Weakening)

If $\Gamma, C, H \vdash e : \tau$ is derivable and $C' \Vdash C$, then so is $\Gamma, C', H \vdash e : \tau$ by a derivation of the same structure.

Lemma 6.6 (Instantiation)

If $\Gamma, C, H \vdash e : \tau$ is derivable, then so is the judgement $\psi(\Gamma), \psi(C), \psi(H) \vdash e : \psi(\tau)$ by a derivation of the same structure.

Although the following Lemma seems obscure, it brings to light some subtle issues related to polymorphism and generalization, including the requirement that top-level effects be ϵ for generalization and instantiation, and that generalisation cannot be performed on free environment variables. It also highlights the soundness of generalization over variables in the top-level constraint, provided the requirements on instantiation. In essence, the result shows that the \forall introduction and elimination forms do not allow over generalization.

Lemma 6.7

Consecutive instances of \forall -INTRO and \forall -ELIM may be suppressed.

Proof

Assume the following derivation structure:

$$\frac{\frac{\Gamma, C, \epsilon \vdash e : \tau \quad \bar{\beta} \# \text{fv}(\Gamma)}{\Gamma, C, \epsilon \vdash e : \forall \bar{\beta}. \tau} \quad C \Vdash [\bar{\tau}/\bar{\beta}]}{\Gamma, C, \epsilon \vdash e : \tau[\bar{\tau}/\bar{\beta}]}$$

where $\Gamma, C, \epsilon \vdash e : \tau$ is derivable by assumption. Thus $\Gamma[\bar{\tau}/\bar{\beta}], C[\bar{\tau}/\bar{\beta}], \epsilon[\bar{\tau}/\bar{\beta}] \vdash e : \tau[\bar{\tau}/\bar{\beta}]$ is derivable by a derivation of the same structure, by Lemma 6.6. But $\epsilon[\bar{\tau}/\bar{\beta}] = \epsilon$, and $\Gamma[\bar{\tau}/\bar{\beta}] = \Gamma$ since $\bar{\beta} \# \text{fv}(\Gamma)$ by assumption; therefore, the result follows by Lemma 6.5 since $C \Vdash C[\bar{\tau}/\bar{\beta}]$ by assumption and Definition 6.2. \square

The following result allows normalization of type derivations. When inducting on the structure of type derivations, normal form derivations allow type judgments to be *inverted*; their subderivations can be deconstructed from root judgments, given the form of the root expression.

Lemma 6.8 (Normalization)

If $\Gamma, C, H \vdash e : \tau$ is derivable, then it must follow by an instance of SUB from a judgement \mathcal{J} such that

- i. if $e = \lambda_z x. e'$ then \mathcal{J} follows by FIX;
- ii. if $e = e_1 e_2$ then \mathcal{J} follows by APP;
- iii. if $e = \text{if } e_1 \text{ then } e_2 \text{ else } e_3$ then \mathcal{J} follows by IF;
- iv. if $e = w(e')$ then \mathcal{J} follows by EVENT;
- v. if e is one of `true`, `false`, `()`, `∨`, `∧`, `¬`, `c`, then \mathcal{J} follows by CONST;
- vi. if $e = x$ then \mathcal{J} follows by VAR and \forall -ELIM;
- vii. if $e = (\text{let } x = v \text{ in } e')$ the \mathcal{J} follows by LET.

Proof

By Lemma 6.7, any consecutive instances of \forall -INTRO and \forall -ELIM can be suppressed. In effect, this allows us to restrict all instances of \forall -ELIM to immediately follow instances of VAR, and all instances of \forall -INTRO to immediately precede LET. Furthermore, by Lemma 6.1 it is easy to show that consecutive instances of SUB can be collapsed into a single instance, and due to reflexivity of SUB, trivial instances of SUB can be inserted anywhere in a derivation other than following \forall -INTRO. Since any derivation of a judgement $\Gamma, C, H \vdash e : \tau$ must contain a syntax-directed rule instance corresponding to the form of e , with a judgement \mathcal{J} —as the consequence, the result follows by suppressing, collapsing, or inserting non-syntax-directed rule instances as appropriate between \mathcal{J} and $\Gamma, C, H \vdash e : \tau$. \square

Corollary 6.1

If $\Gamma, C, H \vdash v : \tau$ is derivable, then so is $\Gamma, C, \epsilon \vdash v : \tau'$ with $C \Vdash \tau' \sqsubseteq \tau$ and $C \Vdash \epsilon \sqsubseteq H$.

Proof

Straightforward by the previous Lemma, and by observing that if a judgement $\Gamma, C, H' \vdash v : \tau'$ follows by a syntax-directed rule applicable to a value v , then $H' = \epsilon$. \square

To demonstrate the *let* and β reduction cases of subject reduction, we prove a values substitution Lemma as follows:

Lemma 6.9 (Substitution)

If both $\Gamma; x : \sigma', C, H \vdash e : \sigma$ and $\Gamma, C, \epsilon \vdash v : \sigma'$ are derivable, then so is $\Gamma, C, H \vdash e[v/x] : \sigma$.

Proof

By induction on the derivation of $\Gamma; x : \sigma', C, H \vdash e : \sigma$ and case analysis on the last step in the derivation. We restrict our consideration to the most interesting cases.

Case \forall -INTRO. In this case $\sigma = \forall \bar{\beta}. \tau$, where by definition of \forall -INTRO we can reconstruct:

$$\frac{\Gamma; x : \sigma', C, H \vdash e : \tau \quad \bar{\beta} \# \text{fv}(\Gamma; x : \sigma')}{\Gamma; x : \sigma', C, H \vdash e : \forall \bar{\beta}. \tau}$$

But by the induction hypothesis the judgement $\Gamma, C, H \vdash e[v/x] : \tau$ is derivable, and since clearly $\text{fv}(\Gamma) \subseteq \text{fv}(\Gamma; x : \sigma')$ therefore $\bar{\beta} \# \text{fv}(\Gamma)$, so this case holds via an instance of \forall -INTRO.

Case VAR. In this case $e = y$, $H = \epsilon$, and by definition of VAR we have that $(\Gamma; x : \sigma')(y) = \sigma$. Now, suppose on the one hand that $x \neq y$. Then $e[v/x] = y$, and $\Gamma(y) = \sigma$, so this case holds via an instance of VAR. Suppose on the other hand that $x = y$; then $e[v/x] = v$ and $\sigma = \sigma'$, so the result follows by an assumption of the Lemma.

Case FIX. In this case $e = \lambda_z y. e'$ and $\sigma = \tau_1 \xrightarrow{H'} \tau_2$ and $H = \epsilon$, and by definition of FIX we can reconstruct:

$$\frac{\Gamma; x : \sigma'; y : \tau_1; z : \tau_1 \xrightarrow{H'} \tau_2, C, H' \vdash e' : \tau_2}{\Gamma; x : \sigma', C, \epsilon \vdash \lambda_z y. e' : \tau_1 \xrightarrow{H'} \tau_2}$$

Supposing that $x \neq y$ and $x \neq z$, it is the case that:

$$\Gamma; x : \sigma'; y : \tau_1; z : \tau_1 \xrightarrow{H'} \tau_2 = \Gamma; y : \tau_1; z : \tau_1 \xrightarrow{H'} \tau_2; x : \sigma'$$

Hence the following is derivable by the induction hypothesis:

$$\Gamma; y : \tau_1; z : \tau_1 \xrightarrow{H'} \tau_2, C, H' \vdash e'[v/x] : \tau_2$$

and by FIX:

$$\frac{\Gamma; y : \tau_1; z : \tau_1 \xrightarrow{H'} \tau_2, C, H' \vdash e'[v/x] : \tau_2}{\Gamma; x : \sigma', C, \epsilon \vdash \lambda_z y. (e'[v/x]) : \tau_1 \xrightarrow{H'} \tau_2}$$

so the result follows, since $\lambda_z y. (e'[v/x]) = e[v/x]$ by definition of substitution. Supposing that $x = y$, it is the case that

$$\Gamma; x : \sigma'; y : \tau_1; z : \tau_1 \xrightarrow{H'} \tau_2 = \Gamma; y : \tau_1; z : \tau_1 \xrightarrow{H'} \tau_2$$

and $e[v/x] = v$ by definition of substitution, so we can reconstruct by FIX:

$$\frac{\Gamma; y : \tau_1; z : \tau_1 \xrightarrow{H'} \tau_2, C, H' \vdash e' : \tau_2}{\Gamma, C, \epsilon \vdash (\lambda_z y. e')[v/x] : \tau_1 \xrightarrow{H'} \tau_2}$$

The result follows similarly supposing that $x = z$. \square

Type safety ultimately relies on the preservation of type validity by \rightarrow^* . Since the relation \rightarrow^* is predicated on both \rightsquigarrow and \rightarrow forms of reduction, we, therefore, need to prove subject reduction results for each form. First, for \rightsquigarrow :

Lemma 6.10 (Subject Reduction \rightsquigarrow)

If $\Gamma, C, H \vdash e : \tau$ is derivable and $\eta, e \rightsquigarrow \eta', e'$, then $\Gamma, C, H' \vdash e' : \tau$ is derivable with $C \Vdash \eta'; H' \sqsubseteq \eta; H$.

Proof

By induction on the derivation of $\Gamma, C, H \vdash e : \tau$ and case analysis on the rule form of $\eta, e \rightsquigarrow \eta', e'$. We restrict our consideration to the most interesting representative cases.

Case β . The following assertions hold by assumption in this subcase:

$$e = (\lambda_z x. e'')v \quad e' = e''[v/x][\lambda_z x. e''/z] \quad \eta' = \eta$$

By Lemma 6.8, reconstruct the following, moving from the root upwards; by SUB:

$$\frac{\Gamma, C, H_1; H_2; H_3 \vdash (\lambda_z x. e'')v : \tau_2 \quad C \Vdash \tau_2 \sqsubseteq \tau \quad C \Vdash H_1; H_2; H_3 \sqsubseteq H}{\Gamma, C, H \vdash (\lambda_z x. e'')v : \tau}$$

and by APP:

$$\frac{\Gamma, C, H_1 \vdash \lambda_z x. e'' : \tau_1 \xrightarrow{H_3} \tau_2 \quad \Gamma, C, H_2 \vdash v : \tau_1}{\Gamma, C, H_1; H_2; H_3 \vdash (\lambda_z x. e'')v : \tau_2}$$

and by SUB:

$$\frac{\Gamma, C, \epsilon \vdash \lambda_z x. e'' : \tau'_1 \xrightarrow{H'_3} \tau'_2 \quad C \Vdash \tau'_1 \xrightarrow{H'_3} \tau'_2 \sqsubseteq \tau_1 \xrightarrow{H_3} \tau_2 \quad C \Vdash \epsilon \sqsubseteq H_1}{\Gamma, C, H_1 \vdash \lambda_z x. e'' : \tau_1 \xrightarrow{H_3} \tau_2}$$

and by FIX:

$$\frac{\Gamma; x : \tau'_1; z : \tau'_1 \xrightarrow{H'_3} \tau'_2, C, H'_3 \vdash e'' : \tau'_2}{\Gamma, C, \epsilon \vdash \lambda_z x. e'' : \tau'_1 \xrightarrow{H'_3} \tau'_2}$$

and by SUB and Corollary 6.1:

$$\frac{\Gamma, C, \epsilon \vdash v : \tau_3 \quad C \Vdash \epsilon \sqsubseteq H_2 \quad C \Vdash \tau_3 \sqsubseteq \tau_1}{\Gamma, C, H_2 \vdash v : \tau_1}$$

Now, observe that by contravariance of \sqsubseteq in function domain types and transitivity of \sqsubseteq we have $C \Vdash \tau_3 \sqsubseteq \tau'_1$, hence by SUB and the above we can derive $\Gamma, C, H_2 \vdash v : \tau'_1$, so by successive applications of Lemma 6.9 we have that $\Gamma, C, H'_3 \vdash e' : \tau'_2$ is derivable. But $C \Vdash H'_3 \sqsubseteq \epsilon; \epsilon; H'_3$ by Lemma 3.1, and since by the above and definition of \sqsubseteq :

$$C \Vdash H'_3 \sqsubseteq H_3 \quad C \Vdash \epsilon \sqsubseteq H_1 \quad C \Vdash \epsilon \sqsubseteq H_2$$

therefore by Lemma 6.4:

$$C \Vdash H'_3 \sqsubseteq H_1; H_2; H_3$$

and since $C \Vdash H_1; H_2; H_3 \sqsubseteq H$ and $C \Vdash \tau'_2 \sqsubseteq \tau$ by the above and transitivity of \sqsubseteq , by SUB:

$$\frac{\Gamma, C, H'_3 \vdash e' : \tau'_2 \quad C \Vdash \tau'_2 \sqsubseteq \tau \quad C \Vdash H'_3 \sqsubseteq H}{\Gamma, C, H \vdash e' : \tau}$$

so this case holds.

Case let. In this case $e = \text{let } x = v \text{ in } e_0$ and $e' = e_0[v/x]$ and $\eta = \eta'$ by definition. By Lemma 6.8, reconstruct the following; by SUB:

$$\frac{\Gamma, C, H_0 \vdash \text{let } x = v \text{ in } e_0 : \tau_0 \quad C \Vdash H_0 \sqsubseteq H \quad C \Vdash \tau_0 \sqsubseteq \tau}{\Gamma, C, H \vdash \text{let } x = v \text{ in } e_0 : \tau}$$

and by LET:

$$\frac{\Gamma, C, \epsilon \vdash v : \sigma \quad \Gamma; x : \sigma, C, H_0 \vdash e_0 : \tau_0}{\Gamma, C, H_0 \vdash \text{let } x = v \text{ in } e_0 : \tau_0}$$

But then $\Gamma, C, H_0 \vdash e' : \tau_0$ is derivable by Lemma 6.9, hence by the above and SUB:

$$\frac{\Gamma, C, H_0 \vdash e' : \tau_0 \quad C \Vdash H_0 \sqsubseteq H \quad C \Vdash \tau_0 \sqsubseteq \tau}{\Gamma, C, H \vdash e' : \tau}$$

so this case holds.

Case event. In this case $e = \text{ev}(c)$ and $e' = ()$ and $\eta' = \eta; \text{ev}(c)$. By Lemma 6.8 and definition of the SUB, EVENT, and CONST, we can reconstruct:

$$\frac{\Gamma, C, H'; \text{ev}(s) \vdash \text{ev}(c) : \text{unit} \quad C \Vdash \text{unit} \sqsubseteq \tau \quad C \Vdash H'; \text{ev}(s) \sqsubseteq H}{\Gamma, C, H \vdash \text{ev}(c) : \tau}$$

and

$$\frac{\frac{\Gamma, C, \epsilon \vdash c : \{c\} \quad C \Vdash \{c\} \sqsubseteq \{s\} \quad C \Vdash \epsilon \sqsubseteq H'}{\Gamma, C, H' \vdash c : \{s\}}}{\Gamma, C, H'; \text{ev}(s) \vdash \text{ev}(c) : \text{unit}}$$

But then by applications of CONST and SUB:

$$\frac{\Gamma, C, \epsilon \vdash () : \text{unit} \quad C \Vdash \text{unit} \sqsubseteq \tau \quad C \Vdash \epsilon \sqsubseteq H'}{\Gamma, C, H' \vdash () : \tau}$$

and since $C \Vdash \{c\} \sqsubseteq \{s\}$ by the above, therefore $C \Vdash c \sqsubseteq s$ and hence $C \Vdash \text{ev}(c) \sqsubseteq \text{ev}(s)$ by Definition 5.3, thus by Lemma 6.4 we have $C \Vdash \eta; \text{ev}(c); H' \sqsubseteq \eta; \text{ev}(s); H'$, so this case holds. \square

Before proving subject reduction for \rightarrow , we need to establish some typing results related to evaluation contexts. First, if e is a redex, then the effect of $E[e]$ should reflect that the trace predicted for e is the “first effect that will happen” in the evaluation of $E[e]$. Formally, we characterize this as follows:

Lemma 6.11 (Context Inversion)

If $\Gamma, C, H \vdash E[e] : \tau$ is derivable for closed $E[e]$, then there exists $H_1; H_2$ such that $\Gamma, C, H_1; H_2 \vdash E[e] : \tau$ is derivable with a subderivation concluding in $\Gamma, H_1 \vdash e : \tau'$ for e in the hole, where $C \Vdash H_1; H_2 \sqsubseteq H$.

Proof

By induction on E . In the basis $E = []$ and $E[e] = e$ so the result follows by assumption and Lemma 3.1, taking $H_2 = \epsilon$. Otherwise, the remaining forms of $E - E'e', vE'$, if E' then e_1 else e_2 , and $\text{ev}(E')$ —ensure that H will at least reflect that the

effect of e “happens first.” Taking $E = E'e'$, for example, so that $E[e] = E'[e]e'$, by Lemma 6.8 and SUB and APP we can reconstruct

$$\frac{\Gamma, C, H'_1; H'_2; H'_3 \vdash E'[e]e' : \tau' \quad C \Vdash \tau' \sqsubseteq \tau \quad C \Vdash H'_1; H'_2; H'_3 \sqsubseteq H}{\Gamma, C, H \vdash E'[e]e' : \tau}$$

and

$$\frac{\Gamma, C, H'_1 \vdash E'[e] : \tau_1 \xrightarrow{H'_3} \tau_2 \quad \Gamma, C, H'_2 \vdash e' : \tau_1}{\Gamma, C, H'_1; H'_2; H'_3 \vdash E'[e]e' : \tau'}$$

But then by the induction hypothesis, there exists H''_2 such that $\Gamma, C, H_1; H''_2 \vdash E'[e] : \tau_1 \xrightarrow{H'_3} \tau_2$ is derivable with $C \Vdash H_1; H''_2 \sqsubseteq H'_1$, so the result follows by an instance of APP, Lemmas 3.1 and 6.4, taking $H_2 = H''_2; H'_2; H'_3$. The other cases follow in a similar manner, case $E = vE'$ with a little help from Corollary 6.1. \square

Corollary 6.2

If $\Gamma, C, H \vdash E[ew(c)] : \tau$ is derivable, there exists H' such that $C \Vdash ew(c); H' \sqsubseteq H$.

We also show that typing is preserved by contextual substitution:

Lemma 6.12 (Context Substitution)

If $\Gamma, C, H_1; H_2 \vdash E[e] : \tau$ is derivable with a subderivation concluding in $\Gamma, C, H_1 \vdash e : \tau'$ for e in the hole, and $\Gamma, C, H'_1 \vdash e' : \tau'$ is derivable, then so is $\Gamma, C, H'_1; H_2 \vdash E[e'] : \tau$.

Proof

Straightforward by induction on E and Lemma 6.11. \square

Now, the main result for \rightarrow reduction:

Lemma 6.13 (Subject Reduction \rightarrow)

If $\Gamma, C, H \vdash e : \tau$ is derivable for closed e and $\eta, e \rightarrow \eta', e'$, then $\Gamma, C, H' \vdash e' : \tau$ is derivable with $C \Vdash \eta'; H' \sqsubseteq \eta; H$.

Proof

The following hold by assumption and definition of \rightarrow :

$$e = E[e_1] \text{ with } e_1 \text{ a redex} \quad e' = E[e_2] \quad \eta, e_1 \rightsquigarrow \eta', e_2$$

Also, by Lemma 6.11 there exists a derivation of $\Gamma, C, H_1; H_2 \vdash e : \tau$ with $C \Vdash H_1; H_2 \sqsubseteq H$, and with a subderivation concluding in $\Gamma, C, H_1 \vdash e_1 : \tau'$ for e_1 in the hole. Furthermore, by Lemma 6.10 we have that $\Gamma, H'_1 \vdash e_2 : \tau'$ is derivable with:

$$C \Vdash \eta'; H'_1 \sqsubseteq \eta; H_1$$

Now, the following judgement is derivable by Lemma 6.12:

$$\Gamma, C, H'_1; H_2 \vdash E[e_2] : \tau$$

and by Lemma 3.1, Lemma 6.4, preceding facts, and Lemma 6.1:

$$\begin{aligned} C \Vdash \eta; H_1; H_2 &\sqsubseteq \eta; H \\ C \Vdash \eta'; H'_1; H_2 &\sqsubseteq \eta; H_1; H_2 \\ C \Vdash \eta'; H'_1; H_2 &\sqsubseteq \eta; H \end{aligned}$$

The result follows taking $H' = H'_1; H_2$. \square

Note that as a corollary of this result, we may formalise the intuition that “trace effects predict run-time program traces”:

Corollary 6.3 (Trace Approximation)

If $\Gamma, C, H \vdash e : \tau$ is derivable for closed e and $\epsilon, e \rightarrow^* \eta, e'$ and $C \Vdash H \sqsubseteq H'$ for closed H' , then $\hat{\eta} \in \llbracket H' \rrbracket$.

Proof

By Lemma 6.13 and induction on the length of reduction there exists H'' such that $\Gamma, C, H'' \vdash e' : \tau$ is derivable with $C \Vdash \eta; H'' \sqsubseteq H$. Now, suppose that ρ is a solution of C , meaning that $\llbracket \eta; \rho(H'') \rrbracket \subseteq \llbracket \rho(H) \rrbracket$, therefore clearly $\hat{\eta} \in \llbracket \rho(H) \rrbracket$. Since also $\llbracket \rho(H) \rrbracket \subseteq H'$, the result follows. \square

Using well-known strategy, we demonstrate progress on the basis of a canonical forms lemma, establishing the form of values denoted by a particular type:

Lemma 6.14 (Canonical Forms)

If $\Gamma, C, H \vdash v : \tau$ is satisfiable for closed v , then the following conditions hold:

- i. If $\tau = \text{bool}$ then v is either true or false.
- ii. If $\tau = \text{unit}$ then $v = ()$.
- iii. If there exists s such that $\tau = \{s\}$ then there exists c such that $v = c$.
- iv. If there exists τ_1, H, τ_2 such that $\tau = \tau_1 \xrightarrow{H} \tau_2$ then either v is some function $\lambda_2 x.e$, or v is in $\{\wedge, \vee, \neg\}$.

Proof

Suppose that $\tau = \text{bool}$, and suppose on the contrary that $v \notin \{\text{true}, \text{false}\}$. Then by Lemma 6.8 and remaining possible forms of v it is the case that $\Gamma, C, \epsilon \vdash v : \tau'$ such that $C \Vdash \tau' \sqsubseteq \text{bool}$ where τ' is of the form unit , $\{c\}$, or $\tau_1 \xrightarrow{H} \tau_2$. Lemma 6.2 implies the contradiction. The other cases follow in a similar manner. \square

Finally, on the basis of the proceeding, we can prove our main type safety results.

Theorem 6.1 (Progress)

Suppose $\Gamma, C, H \vdash e : \tau$ is satisfiable, there exists a solution ρ of C such that $\rho(\eta; H)$ is valid, and η, e is irreducible. Then e is a value.

Proof

Suppose on the contrary that e is not a value. Then by definition of \rightarrow , e is of the form $E[f]$, with $\Gamma, C, H' \vdash f : \tau'$ derivable as a subderivation of $\Gamma, C, H \vdash e : \tau$, and where one of the following cases holds by definition of \rightsquigarrow :

Case (f is one of $\{\text{if } v \text{ then } e_1 \text{ else } e_2, \wedge v, \vee v, \neg v\}$, where v is not a boolean value). In this case it is easy to show by inversion of either IF or APP that $\Gamma, C, H'' \vdash v : \text{bool}$. The contradiction follows by Lemma 6.14.

Case (f is $ev(v)$ and v is not a singleton constant c). Similar to the previous case.

Case (f is $\phi(c)$ where $\Pi(\phi(c), \hat{\eta} ev_\phi(c))$ does not hold). By Corollary 6.2, there exists H_2 such that $C \Vdash ev_\phi(c); H_2 \sqsubseteq H$. By assumption there exists a solution ρ of C such that $\rho(\eta; H)$ is valid, and $\llbracket \rho(\eta; ev_\phi(c); H_2) \rrbracket \subseteq \llbracket \rho(\eta; H) \rrbracket$ by Definition 5.3, hence $\rho(\eta; ev_\phi(c); H_2)$ is also valid by Lemma 3.2. But interpretations of effects are prefix-closed, so clearly $\hat{\eta} ev_\phi(c) \in \llbracket \rho(\eta; ev_\phi(c); H_2) \rrbracket$, so $\Pi(\phi(c), \hat{\eta} ev_\phi(c))$ holds by Definition 3.4, which is a contradiction. \square

$$\begin{array}{c}
\text{C-FN} \\
(\tau_1 \xrightarrow{H} \tau_2 \sqsubseteq \tau'_1 \xrightarrow{H'} \tau'_2) \rightsquigarrow_{\text{close}} (\tau'_1 \sqsubseteq \tau_1 \wedge \tau_2 \sqsubseteq \tau'_2 \wedge H \sqsubseteq H') \\
\\
\text{C-TRANS} \qquad \qquad \qquad \text{C-SINGLETON} \\
(\tau_1 \sqsubseteq \tau_2 \wedge \tau_2 \sqsubseteq \tau_3) \rightsquigarrow_{\text{close}} \tau_1 \sqsubseteq \tau_3 \qquad \qquad \qquad \{\tau_1\} \sqsubseteq \{\tau_2\} \rightsquigarrow_{\text{close}} \tau_1 \sqsubseteq \tau_2 \\
\\
\text{C-CONTEXT} \\
\frac{C' \sqsubseteq C \quad C' \rightsquigarrow_{\text{close}} D \quad D \not\sqsubseteq C}{C \rightarrow_{\text{close}} C \wedge D}
\end{array}$$

Fig. 12. Constraint closure rules.

Theorem 6.2 (Type Safety)

If $\Gamma, C, H \vdash e : \tau$ is valid for closed e , then e does not go wrong.

Proof

Suppose on the contrary that $e, e \rightarrow^n \eta, e'$ with η, e' irreducible and e' not a value. Then $\Gamma, C, H' \vdash e' : \tau$ is derivable with $C \Vdash \eta; H' \sqsubseteq H$ by Lemma 6.13, induction on n , and Lemma 6.1. Now, since $\Gamma, C, H \vdash e : \tau$ is valid by assumption; therefore, there exists a solution ρ of C such that $\rho(H)$ is valid by Definition 6.4. But since $\llbracket \rho(\eta; H') \rrbracket \sqsubseteq \llbracket \rho(H) \rrbracket$ by Definition 5.3; therefore, $\rho(\eta; H')$ is valid by Lemma 3.2. The contradiction follows by Theorem 6.1. \square

7 Subtyping constraint inference

Like the HM-style type system of Figure 4, typings specified by the subtyping system of Figure 9 are inferable. We now give a type inference algorithm for constraint subtyping, and prove that it is sound and complete with respect to logical judgments. Furthermore, our completeness result yields a principal types property for λ_{trace} constraint subtyping. The inference algorithm and form of judgments are the same as for the HM style system, defined in Figure 5. The difference is in the constraint solution technique. Rather than a unification-based solver as in the HM style system, we use a *closure* technique. The technique is standard, with extensions to accommodate trace effects and singleton types. To exploit model-checking techniques to verify inferred effects, we also develop an *effect extraction* algorithm to represent top-level effects in term form. Our technical development culminates in Corollary 7.1 and Lemma 7.22, formally establishing correctness of the composition of inference, closure, and effect extraction, but before the proofs we provide some initial definitions and intuitions.

Closure and consistency. In contrast to unification, which literally generates a solution to the constraint system, closure only verifies that a constraint is solvable. The closure algorithm is the iteration of a single-step rewrite relation defined in Figure 12. Each rewrite rule makes explicit constraints that are implicit in existing ones for example, observe that closure of a constraint $\{s_1\} \sqsubseteq \{s_2\}$ adds $s_1 \sqsubseteq s_2$ to the closure. When no new constraints can be added, closure terminates. Formally:

$\vdash \mathbf{true} : ok$	$\frac{\vdash C : ok \quad \vdash D : ok}{\vdash C \wedge D : ok}$	$\vdash H \sqsubseteq H' : ok$	$\vdash \tau \sqsubseteq \tau' : ok$
$\frac{\tau, \beta : k}{\vdash \tau \sqsubseteq \beta : ok}$	$\frac{\tau, \beta : k}{\vdash \beta \sqsubseteq \tau : ok}$	$\vdash \tau_1 \xrightarrow{H} \tau_2 \sqsubseteq \tau'_1 \xrightarrow{H'} \tau'_2 : ok$	$\vdash \{\tau_1\} \sqsubseteq \{\tau_2\} : ok$

Fig. 13. Constraint consistency rules.

$hextract_C H = hextract_C(H, \emptyset)$
$hextract_C(\epsilon, hs) = \epsilon$
$hextract_C(ev(\alpha), hs) = ev(bounds_C \alpha)$
$hextract_C(h, hs) = h \quad h \in hs$
$hextract_C(h, hs) = \mu h. hextract_C((bounds_C h), hs \cup \{h\}) \quad h \notin hs$
$hextract_C(H_1; H_2, hs) = (hextract_C(H_1, hs)); (hextract_C(H_2, hs))$
$hextract_C(H_1 H_2, hs) = (hextract_C(H_1, hs))(hextract_C(H_2, hs))$
$bounds_C \beta = \tau_1 \dots \tau_n \quad \text{where } \{\tau_1, \dots, \tau_n\} = \{\tau \mid \tau \sqsubseteq \beta \in C \text{ and } \tau \notin \mathcal{V}_k\}$

Fig. 14. $hextract$, $hextract_C$, and $bounds$ functions.**Definition 7.1 (Constraint Closure)**

The rewrite relations \rightsquigarrow_{close} and \rightarrow_{close} are defined in Figure 12. C is *closed* iff there does not exist D such that $C \rightarrow_{close} D$. The relation \rightarrow_{close}^* is the reflexive, transitive closure of \rightarrow_{close} . We define $close(C)$ as a closed constraint such that $C \rightarrow_{close}^* close(C)$.

When closure terminates, all “atomic elements” of the constraint system are made evident. A simple structural consistency check, defined in Figure 13, can then be used to ensure that all constraints entailed by the closure are solvable. For example, observe that clashing types in the sense of Lemma 6.2 cannot populate a consistent constraint. The definition of consistency exploits kinding rules extended to type terms to ensure consistency of constraints involving variables. Formally:

Definition 7.2

A type term τ has kind k , written $\tau : k$, iff the root constructor of any interpretation of τ has kind k as specified in Figure 11.

Note that sanity conditions on interpretations imply that if $\beta \in \mathcal{V}_k$, then $\beta : k$.

Extracting effects from constraints. While demonstrating solvability is sufficient for constraints in general, model checking of inferred trace effects requires that they be presented in term form, not constraint form. Thus, we define another algorithm in Figure 14, called $hextract$, that extracts a term representation of the inferred top-level trace effect from a given constraint. The algorithm $hextract$ is defined in a curried

style that integrates nicely with a combinator fixpoint construction proof technique, used for proving soundness of *hextract* in Section 7.3. The definition of *hextract*_C specifies a family of functions, each element of which has is defined in terms of a fixed parameter *C*. Note that *hextract*_C is not defined on the μ -bound form of trace effects, because this form does not occur in inferred types and constraints. However, *hextract*_C does introduce μ -bindings, in order to solve recursively constrained effect variables.

Like the solution technique for the HM style system, *hextract* exploits the property that inferred constraints define a system of lower bounds on effect variables. However, in the subtyping constraint system, there is the added complication of singleton union types, as presented in Definition 6.1, although it turns out that singleton constraints also define a system of lower bounds on variables. The function *bounds*_C obtains the set of lower bounds of a given variable in *C*.

7.1 Soundness of inference

Before delving into our main results, here are some auxiliary lemmas concerning variable freshness and constraint entailment that will be useful in later proofs.

Lemma 7.1

If $\Gamma, H \vdash_{\bar{\beta}} e : \tau / C$ is canonically derivable, then $\bar{\beta} \# \text{fv}(\Gamma)$ and also $\text{fv}(\tau, C, H) \subseteq \text{fv}(\Gamma) \cup \bar{\beta}$.

Lemma 7.2

All of the following properties hold:

1. If $C \Vdash \tau' \sqsubseteq \tau$, then $\text{fv}(\tau) \subseteq \text{fv}(\tau', C)$.
2. $C \wedge D \Vdash D$.
3. If $C \Vdash \tau \sqsubseteq \tau'$ then $\psi(C) \Vdash \psi(\tau) \sqsubseteq \psi(\tau')$.

We now demonstrate soundness of type inference with respect to constraint subtyping; that is, we show that inferred judgments are derivable in the logical constraint subtyping system of Figure 9. For this purpose it is necessary to convert constrained type schemes in inference environments to unconstrained type schemes in logical environments.

Definition 7.3

We define the following environment transformation function:

$$\begin{aligned} [\emptyset] &= \emptyset \\ [\Gamma; x : \forall \bar{\beta}. \tau / C] &= [\Gamma]; x : \forall \bar{\beta}. \tau \end{aligned}$$

Now, soundness can be proved by a straightforward induction on inference derivations. The proof technique essentially rewrites inference derivations into logical derivations by pushing reconstructed constraints from the root to the leaves.

Lemma 7.3 (Soundness of Subtype Inference)

If $\Gamma, H \vdash_{\bar{\beta}} e : \tau / C$ is canonically derivable, then so is $[\Gamma], C \wedge D, H \vdash e : \tau$ for any *D*.

Proof

By induction on the derivation of $\Gamma, H \vdash_{\bar{\beta}} e : \tau/C$ and case analysis on the last rule instance.

Case VAR. In this case $e = x$ and $H = \epsilon$ and $\tau/C = \tau_0[\bar{\beta}/\bar{\beta}_0]/C_0[\bar{\beta}/\bar{\beta}_0]$ by definition of inference derivations, where

$$\frac{\Gamma(x) = \forall \bar{\beta}_0. \tau_0 / C_0}{\Gamma, \epsilon \vdash_{\bar{\beta}} x : \tau_0[\bar{\beta}/\bar{\beta}_0] / C_0[\bar{\beta}/\bar{\beta}_0]}$$

But since clearly $C_0[\bar{\beta}/\bar{\beta}_0] \wedge D \Vdash [\bar{\beta}/\bar{\beta}_0]$ for arbitrary D , therefore by an instance of VAR and an instance of \forall -ELIM we can derive

$$\frac{\frac{[\Gamma](x) = \forall \bar{\beta}_0. \tau_0}{[\Gamma], C_0[\bar{\beta}/\bar{\beta}_0], \epsilon \vdash x : \forall \bar{\beta}_0. \tau_0} \quad C_0[\bar{\beta}/\bar{\beta}_0] \Vdash [\bar{\beta}/\bar{\beta}_0]}{[\Gamma], C_0[\bar{\beta}/\bar{\beta}_0], \epsilon \vdash x : \tau_0[\bar{\beta}/\bar{\beta}_0]}}$$

Case LET. In this case $e = \text{let } x = v \text{ in } e_0$ and $C = C_1 \wedge C_2$ and $\bar{\beta} = \bar{\beta}_1 \bar{\beta}_2$ by definition of inference derivations, and we have

$$\frac{\bar{\beta}_1 \# \bar{\beta}_2 \quad \Gamma, \epsilon \vdash_{\bar{\beta}_1} v : \tau_1 / C_1 \quad \Gamma; \forall \bar{\beta}_1. \tau_1 / C_1, \epsilon \vdash_{\bar{\beta}_2} e_0 : \tau_2 / C_2}{\Gamma, H \vdash_{\bar{\beta}_1 \bar{\beta}_2} \text{let } x = v \text{ in } e_0 : \tau / C_1 \wedge C_2}$$

But by the induction hypothesis both of

$$[\Gamma], C_1 \wedge C_2, \epsilon \vdash v : \tau_1 \quad [\Gamma]; x : \forall \bar{\beta}_1. \tau_1, C_1 \wedge C_2, H \vdash e_0 : \tau_2$$

are derivable, and by Lemma 7.1 we have that $\bar{\beta}_1 \# \text{fv}(\Gamma)$, hence $\bar{\beta}_1 \# \text{fv}([\Gamma])$, so by applications of \forall -INTRO and LET:

$$\frac{\frac{[\Gamma], C_1 \wedge C_2, \epsilon \vdash v : \tau_1 \quad \bar{\beta}_1 \# \text{fv}([\Gamma])}{[\Gamma], C_1 \wedge C_2, \epsilon \vdash v : \forall \bar{\beta}_1. \tau_1} \quad [\Gamma]; x : \forall \bar{\beta}_1. \tau_1, C_1 \wedge C_2, H \vdash e_0 : \tau_2}{[\Gamma], C_1 \wedge C_2, H \vdash \text{let } x = v \text{ in } e_0 : \tau_2}}$$

Case FIX. In this case:

$$e = \lambda_z x. e' \quad \bar{\beta} = \bar{\beta}' \cup \{t, h, t'\} \quad H = h \quad \tau = t \xrightarrow{h} t' \quad C = C' \wedge t' \sqsubseteq t' \wedge H \sqsubseteq h$$

where we have

$$\frac{\Gamma; x : t; z : t \xrightarrow{h} t', H' \vdash_{\bar{\beta}'} e' : \tau' / C}{\Gamma, \epsilon \vdash_{\bar{\beta}' \cup \{t, h, t'\}} \lambda_z x. e' : t \xrightarrow{h} t' / C' \wedge t' \sqsubseteq t' \wedge H' \sqsubseteq h}$$

by definition of inference derivations. By the induction hypothesis, the judgement

$$[\Gamma]; x : t; z : t \xrightarrow{h} t', C, H' \vdash e' : \tau'$$

is derivable, so by applications of SUB and FIX:

$$\frac{\frac{[\Gamma]; x : t; z : t \xrightarrow{h} t', C, H' \vdash e' : \tau' \quad C' \Vdash t' \sqsubseteq t' \quad C' \Vdash H' \sqsubseteq h}{[\Gamma]; x : t; z : t \xrightarrow{h} t', C, h \vdash e' : \tau'}}{[\Gamma], C, \epsilon \vdash \lambda_z x. e' : t \xrightarrow{h} t'}}$$

Case EVENT. In this case we have

$$e = ew(e') \quad H = H'; w(\alpha) \quad \tau = \text{unit} \quad C = C' \wedge \tau' \sqsubseteq \{\alpha\} \quad \bar{\beta} = \bar{\beta}' \cup \{\alpha\}$$

by definition of inference derivations, and as the last derivation step:

$$\frac{\Gamma, H' \vdash_{\bar{\beta}'} e' : \tau' / C'}{\Gamma, H'; w(\alpha) \vdash_{\bar{\beta}' \cup \{\alpha\}} ew(e') : \text{unit} / C' \wedge \tau' \sqsubseteq \{\alpha\}}$$

By the induction hypothesis, the judgement $[\Gamma], C, H' \vdash e' : \tau$ is derivable, and since $C \Vdash \tau' \sqsubseteq \{\alpha\}$, by instances of SUB and LET:

$$\frac{[\Gamma], C, H' \vdash e' : \tau \quad C \Vdash \tau' \sqsubseteq \{\alpha\}}{[\Gamma], C, H' \vdash e' : \{\alpha\}} \quad \frac{[\Gamma], C, H'; \{\alpha\} \vdash ew(e') : \text{unit}}{[\Gamma], C, H' \vdash e' : \tau}$$

Case APP. In this case by definition of inference derivations we have

$$e = e_1e_2 \quad C = C_1 \wedge C_2 \wedge \tau_1 \sqsubseteq \tau_2 \xrightarrow{h} t \quad H = H_1; H_2; h \quad \tau = t \quad \bar{\beta} = \bar{\beta}_1\bar{\beta}_2 \cup \{t, h\}$$

and as the last derivation step:

$$\frac{\bar{\beta}_1\#\bar{\beta}_2 \quad \Gamma, H_1 \vdash_{\bar{\beta}_1} e_1 : \tau_1 / C_1 \quad \Gamma, H_2 \vdash_{\bar{\beta}_2} e_2 : \tau_2 / C_2}{\Gamma, H_1; H_2; h \vdash_{\bar{\beta}_1\bar{\beta}_2 \cup \{t, h\}} e_1e_2 : t / C_1 \wedge C_2 \wedge \tau_1 \sqsubseteq \tau_2 \xrightarrow{h} t}$$

But by the induction hypothesis, both of $\Gamma, H_1, C \vdash e_1 : \tau_1$ and $\Gamma, H_2, C \vdash e_2 : \tau_2$ are derivable, and since $C \Vdash \tau_1 \sqsubseteq \tau_2 \xrightarrow{h} t$ by Lemma 7.2, we can derive by applications of SUB and APP:

$$\frac{\frac{\Gamma, H_1, C \vdash e_1 : \tau_1 \quad C \Vdash \tau_1 \sqsubseteq \tau_2 \xrightarrow{h} t}{\Gamma, H_1, C \vdash e_1 : \tau_2 \xrightarrow{h} t} \quad \Gamma, H_2, C \vdash e_2 : \tau_2}{\Gamma, H_1; H_2; h, C \vdash e_1e_2 : t}$$

so this case holds.

Case IF. In this case we have

$$e = \text{if } e_1 \text{ then } e_2 \text{ else } e_3 \quad \tau = t \quad H = H_1; H_2 | H_3$$

$$C = C_1 \wedge C_2 \wedge C_3 \wedge \tau_1 \sqsubseteq \text{bool} \wedge \tau_2 \sqsubseteq t \wedge \tau_3 \sqsubseteq t \quad \bar{\beta} = \bar{\beta}_1\bar{\beta}_2\bar{\beta}_3 \cup \{t\}$$

by definition of inference derivations, where as the last derivation step

$$\frac{\bar{\beta}_1\#\bar{\beta}_2\#\bar{\beta}_3 \quad \Gamma, H_1 \vdash_{\bar{\beta}_1} e_1 : \tau_1 / C_1 \quad \Gamma, H_2 \vdash_{\bar{\beta}_2} e_2 : \tau_2 / C_2 \quad \Gamma, H_3 \vdash_{\bar{\beta}_3} e_3 : \tau_3 / C_3}{\Gamma, H_1; H_2 | H_3 \vdash_{\bar{\beta}} \text{if } e_1 \text{ then } e_2 \text{ else } e_3 : t / C}$$

By the induction hypothesis, each of:

$$[\Gamma], C, H_1 \vdash e_1 : \tau_1 \quad [\Gamma], C, H_2 \vdash e_2 : \tau_2 \quad [\Gamma], C, H_3 \vdash e_3 : \tau_3$$

is derivable, and by Lemma 7.2 and Lemma 6.4:

$$C \Vdash H_2 \sqsubseteq H_2 | H_3 \quad C \Vdash H_3 \sqsubseteq H_3 | H_3 \quad C \Vdash \tau_1 \sqsubseteq \text{bool} \quad C \Vdash \tau_2 \sqsubseteq t \quad C \Vdash \tau_3 \sqsubseteq t$$

so by an application of SUB:

$$\frac{[\Gamma], C, H_1 \vdash e_1 : \tau_1 \quad C \Vdash \tau_1 \sqsubseteq \text{bool}}{[\Gamma], C, H_1 \vdash e_1 : \text{bool}}$$

and also by an application of SUB:

$$\frac{[\Gamma], C, H_2 \vdash e_2 : \tau_2 \quad C \Vdash H_2 \sqsubseteq H_2|H_3 \quad C \Vdash \tau_2 \sqsubseteq t}{[\Gamma], C, H_2|H_3 \vdash e_2 : t}$$

and by another application of SUB:

$$\frac{[\Gamma], C, H_3 \vdash e_3 : \tau_3 \quad C \Vdash H_3 \sqsubseteq H_2|H_3 \quad C \Vdash \tau_3 \sqsubseteq t}{[\Gamma], C, H_2|H_3 \vdash e_3 : t}$$

so by an application of IF:

$$\frac{[\Gamma], C, H_1 \vdash e_1 : \text{bool} \quad [\Gamma], C, H_2|H_3 \vdash e_2 : t \quad [\Gamma], C, H_2|H_3 \vdash e_3 : t}{[\Gamma], C, H_1; H_2|H_3 \vdash \text{if } e_1 \text{ then } e_2 \text{ else } e_3 : t}$$

So this case holds. Case CONST is immediate by nearly exact correspondence of the CONST rule forms in the logical and inference systems, so the result follows. \square

7.2 Completeness of inference

As is usually the case, completeness is a more difficult manner than soundness, and the development requires more formal overhead. To wit, we extend the \sqsubseteq relation to constrained types, to relate constrained type schemes (taken from logical judgments) with unconstrained type schemes (taken from inference judgments), and to relate constrained type environments with unconstrained type environments. This is all with a view to relating the generality of type inference judgments with logical type judgments. The extended \sqsubseteq relation between type schemes also takes into account generalization of type variables.

Definition 7.4

The relation \sqsubseteq is extended to constrained types, type schemes, and type environments as follows:

$$\frac{C \Vdash C' \quad C \Vdash \tau' \sqsubseteq \tau}{\tau'/C' \sqsubseteq \tau/C}$$

$$\frac{\forall [\bar{\tau}/\bar{\beta}] . C \Vdash [\bar{\tau}/\bar{\beta}] \Rightarrow \exists [\bar{\tau}'/\bar{\beta}'] . (\tau'/C')[\bar{\tau}'/\bar{\beta}'] \sqsubseteq (\tau[\bar{\tau}/\bar{\beta}]/C)}{C \Vdash \forall \bar{\beta}' . \tau'/C' \sqsubseteq \forall \bar{\beta} . \tau}$$

$$\frac{\forall x \in \text{dom}(\Gamma) . C \Vdash \Gamma(x) \sqsubseteq \Gamma'(x)}{C \Vdash \Gamma \sqsubseteq \Gamma'}$$

Lemma 7.4

Definition 7.4 preserves reflexivity and transitivity of \sqsubseteq .

To formalize “maximum” generalization of type schemes, we define the *gen* function, and to specify the generalizable variables in a judgement, we define the *genvars* function.

Definition 7.5

Define $gen(\Gamma, C, \forall \bar{\beta}. \tau) = \forall \bar{\beta}'. \tau$ where $\bar{\beta}' = \text{fv}(\tau, C) - \text{fv}(\Gamma)$. Define $genvars(\Gamma, C, H, \forall \bar{\beta}. \tau) = \text{fv}(C, H, \tau) - \text{fv}(\Gamma)$.

Lemma 7.5

Given $\forall \bar{\beta}. \tau$, and Γ such that $\bar{\beta} \# \text{fv}(\Gamma)$, then $C \Vdash gen(\Gamma, C, \forall \bar{\beta}. \tau) \sqsubseteq \forall \bar{\beta}. \tau$ for any C . Furthermore, for all σ and x such that $\Gamma(x) = \sigma$, we have $gen(\Gamma, C, \sigma) = \sigma$.

Note that in the second part of the above definition, while extraneous bindings in σ may be eliminated in its generalization, the equivalence still holds since we equate type schemes up to elimination of extraneous variables in Section 4.

Next, we formalize the principality relation on type judgments, relating inference judgments with logical judgments. We will show that most general types are inferred, according to this relation. The relation is parameterized by a substitution ψ , which may instantiate free variables in inference judgments for relation to logical judgments; intuitively, ψ imposes the unnecessary assumptions made in type environments in logical derivations.

Definition 7.6

Given $genvars(\Gamma, C, H, \forall \bar{\beta}_0. \tau) = \bar{\beta}$. The relation:

$$\Gamma', H' \vdash_{\bar{\beta}'} e : \tau' / C' \sqsubseteq_{\psi} \Gamma, C, H \vdash e : \forall \bar{\beta}_0. \tau$$

holds iff $C \Vdash \psi(\Gamma') \sqsubseteq \Gamma$, and for all $\psi'' = [\bar{\tau}' / \bar{\beta}']$ where $C \Vdash [\bar{\tau}' / \bar{\beta}']$ there exists $\psi' = [\bar{\tau}' / \bar{\beta}']$ such that

$$C \Vdash \psi' \circ \psi(C') \quad C \Vdash \psi' \circ \psi(H') \sqsubseteq \psi''(H) \quad C \Vdash \psi' \circ \psi(\tau') \sqsubseteq \psi''(\tau)$$

The relation \sqsubseteq_{ψ} takes into account generalizability of top-level effects, but for typing of values, for which effects are ϵ , it is simpler to speak in terms of maximal generalization of types alone. The next Lemmas establish relevant properties.

Lemma 7.6

Supposing that $\text{dom}(\psi) \subseteq \text{fv}(\Gamma')$ and

$$\Gamma', H' \vdash_{\bar{\beta}'} e : \tau' / C' \sqsubseteq_{\psi} \Gamma, C, H \vdash e : \sigma$$

Then $C' \Vdash \forall \bar{\beta}'. \psi(\tau' / C') \sqsubseteq gen(\Gamma, C, \sigma)$, and $H = \epsilon$ implies $H' = \epsilon$.

Lemma 7.7

Given ψ and judgments $\Gamma, C, \epsilon \vdash e : \sigma$ and $\Gamma', \epsilon \vdash_{\bar{\beta}'} e : \tau' / C'$ where $\text{dom}(\psi) \subseteq \text{fv}(\Gamma')$ and $C \Vdash \psi(\Gamma') \sqsubseteq \Gamma$. Then $C \Vdash \forall \bar{\beta}'. \psi(\tau' / C') \sqsubseteq gen(\Gamma, C, \sigma)$ implies

$$\Gamma', \epsilon \vdash_{\bar{\beta}'} e : \tau' / C' \sqsubseteq_{\psi} \Gamma, C, \epsilon \vdash e : \sigma$$

Now, a few more auxiliary properties and definitions, including a definition of substitution restriction, before the main completeness result, which follows by induction on logical type derivations.

Lemma 7.8

If $C \Vdash \tau_1 \sqsubseteq \tau'_1$ and for all ψ_1 with $C \Vdash \psi_1$ there exists ψ_2 such that $C \Vdash \psi_2(\tau_2) \sqsubseteq \psi_1(\tau_1)$, then $C \Vdash \psi_2(\tau_2) \sqsubseteq \psi_1(\tau'_1)$.

Proof

By Lemma 7.2 it is the case that $\psi_1(C) \Vdash \psi_1(\tau_1) \sqsubseteq \psi_1(\tau'_1)$, so $C \Vdash \psi_1(\tau_1) \sqsubseteq \psi_1(\tau'_1)$ by Lemma 6.4, thus $C \Vdash \psi_2(\tau_2) \sqsubseteq \psi_1(\tau'_1)$ by Lemma 6.4. \square

Definition 7.7

Given $\bar{\beta} \subseteq \text{dom}(\psi)$; write $\psi|_{\bar{\beta}}$ to denote ψ' such that $\text{dom}(\psi') = \bar{\beta}$ and $\psi'(\bar{\beta}) = \psi(\bar{\beta})$.

Lemma 7.9

For all $\bar{\beta} \subseteq \text{dom}(\psi)$, if $C \Vdash \psi$ then $C \Vdash \psi|_{\bar{\beta}}$.

Lemma 7.10 (Completeness of Subtype Inference)

Given ψ , Γ' , and derivable judgement $\Gamma, C, H \vdash e : \sigma$ such that $C \Vdash \psi(\Gamma') \sqsubseteq \Gamma$ and $\text{dom}(\psi) \subseteq \text{fv}(\Gamma')$. Then a judgement $\Gamma', H' \vdash_{\bar{\beta}'} e : \tau'/C'$ is canonically derivable, where

$$\Gamma', H' \vdash_{\bar{\beta}'} e : \tau'/C' \sqsubseteq_{\psi} \Gamma, C, H \vdash e : \sigma$$

Proof

By induction on the derivation of $\Gamma, C, H \vdash e : \sigma$, and case analysis on the last rule instance in the derivation.

Case VAR. In this case by definition of subtyping derivations, $H = \epsilon$ and $e = x$ and we can reconstruct

$$\frac{\Gamma(x) = \sigma}{\Gamma, C, \epsilon \vdash x : \sigma}$$

But by VAR in the inference system we can construct

$$\frac{\Gamma'(x) = \forall \bar{\beta}' . \tau'/C'}{\Gamma', \epsilon \vdash_{\bar{\beta}'} x : \tau'[\bar{\beta}''/\bar{\beta}']/C'[\bar{\beta}''/\bar{\beta}']}$$

where $\bar{\beta}''$ are fresh by canonical form derivations. Since $\text{dom}(\psi) = \text{fv}(\Gamma')$, observe that

$$\psi(\forall \bar{\beta}' . \tau'/C') = \forall \bar{\beta}' . \psi(\tau'/C')$$

Now, letting $\sigma = \forall \bar{\beta}' . \tau$, take some arbitrary $[\bar{\tau}/\bar{\beta}]$ such that $C \Vdash [\bar{\tau}/\bar{\beta}]$; then we have by assumption and Definition 7.4

$$\exists [\bar{\tau}'/\bar{\beta}'] . (\psi(\tau'/C'))[\bar{\tau}'/\bar{\beta}'] \sqsubseteq (\tau[\bar{\tau}/\bar{\beta}]/C)$$

Let

$$\tau'' = \tau'[\bar{\beta}''/\bar{\beta}'] \quad C'' = C'[\bar{\beta}''/\bar{\beta}']$$

Since $\bar{\beta}''$ are fresh; therefore, $\bar{\beta}' \# \bar{\beta}'' \# \text{dom}(\psi)$; hence,

$$(\psi(\tau''/C''))[\bar{\tau}'/\bar{\beta}'] = (\psi(\tau'/C'))[\bar{\tau}'/\bar{\beta}']$$

so

$$(\psi(\tau''/C''))[\bar{\tau}'/\bar{\beta}'] \sqsubseteq (\tau[\bar{\tau}/\bar{\beta}]/C)$$

Since $\forall \bar{\beta}' . \tau = \text{gen}(\Gamma, C, \sigma)$ by Lemma 7.5, the result follows by Definition 7.4 and Lemma 7.7.

Case LET. In this case $e = \text{let } x = v \text{ in } e'$ and $\sigma = \tau$ and by definition of subtyping derivations we can reconstruct

$$\frac{\Gamma, C, \epsilon \vdash v : \sigma' \quad \Gamma; x : \sigma', C, H \vdash e' : \tau}{\Gamma, C, H \vdash \text{let } x = v \text{ in } e' : \tau}$$

By the induction hypothesis, a judgement $\Gamma', \epsilon \vdash_{\bar{\beta}_1} v : \tau_1 / C_1$ is canonically derivable, where there exists $\psi_1 = [\bar{\tau}_1 / \bar{\beta}_1]$ such that

$$C \Vdash \psi_1 \circ \psi(C_1)$$

and by Lemma 7.6:

$$C \Vdash \forall \bar{\beta}_1. \psi(\tau_1 / C_1) \sqsubseteq \text{gen}(\Gamma, C, \sigma')$$

But $\text{dom}(\psi) \# \bar{\beta}_1$ by Lemma 7.1, hence

$$\forall \bar{\beta}_1. \psi(\tau_1 / C_1) = \psi(\forall \bar{\beta}_1. \tau_1 / C_1)$$

so by Lemmas 7.4 and 7.5:

$$C \Vdash \psi(\forall \bar{\beta}_1. \tau_1 / C_1) \sqsubseteq \sigma'$$

Thus, $C \Vdash \psi(\Gamma'; x : \forall \bar{\beta}_1. \tau_1 / C_1) \sqsubseteq \Gamma; x : \sigma$ by Definition 7.4. By Lemma 7.1 and the induction hypothesis, letting $\bar{\beta} = \text{genvars}(\Gamma, C, H, \tau)$ we have that

$$\Gamma'; x : \forall \bar{\beta}_1. \tau_1 / C_1, H_2 \vdash_{\bar{\beta}_2} e' : \tau_2 / C_2$$

is canonically derivable, where there exists ψ_2 with $\text{dom}(\psi_2) = [\bar{\tau}_2 / \bar{\beta}_2]$ such that for arbitrary $\psi'' = [\bar{\tau} / \bar{\beta}]$ with $C \Vdash [\bar{\tau} / \bar{\beta}]$:

$$C \Vdash \psi_2 \circ \psi(C_2) \quad C \Vdash \psi_2 \circ \psi(\tau_2) \sqsubseteq \psi''(\tau) \quad C \Vdash \psi_2 \circ \psi(H_2) \sqsubseteq \psi''(H)$$

Assume $\bar{\beta}_2 \# \bar{\beta}_1$ wlog, so that as an instance of LET:

$$\frac{\Gamma', \epsilon \vdash_{\bar{\beta}_1} v : \tau_1 / C_1 \quad \Gamma'; x : \forall \bar{\beta}_1. \tau_1 / C_1, H_2 \vdash_{\bar{\beta}_2} e' : \tau_2 / C_2}{\Gamma', H_2 \vdash_{\bar{\beta}_1 \bar{\beta}_2} \text{let } x = v \text{ in } e' : \tau_2 / C_1 \wedge C_2}$$

Now, we have by assumption and Lemma 7.1 that $\bar{\beta}_2 \# \bar{\beta}_1 \# \text{fv}(\Gamma)$, and also by Lemma 7.1:

$$\text{fv}(\tau_1, C_1) \subseteq \text{fv}(\Gamma) \cup \bar{\beta}_1 \quad \text{fv}(\tau_2, C_2, H_2) \subseteq \text{fv}(\Gamma) \cup \bar{\beta}_2$$

so clearly:

$$\psi_1 \circ \psi(C_1) = \psi_1 \circ \psi_2 \circ \psi(C_1) \quad \psi_2 \circ \psi(C_2) = \psi_1 \circ \psi_2 \circ \psi(C_2)$$

$$\psi_2 \circ \psi(\tau_2) = \psi_1 \circ \psi_2 \circ \psi(\tau_2) \quad \psi_2 \circ \psi(H_2) = \psi_1 \circ \psi_2 \circ \psi(H_2)$$

therefore by the above:

$$C \Vdash \psi_1 \circ \psi_2 \circ \psi(C_1 \wedge C_2) \quad C \Vdash \psi_1 \circ \psi_2 \circ \psi(\tau_2) \sqsubseteq \psi''(\tau)$$

$$C \Vdash \psi_1 \circ \psi_2 \circ \psi(H_2) \sqsubseteq \psi''(H)$$

and $\text{dom}(\psi_1 \circ \psi_2) = \bar{\beta}_1 \bar{\beta}_2$, so this case holds by Definition 7.6.

Case \forall -INTRO. In this case by definition of subtyping derivations we have that $e = v$ and $\sigma = \forall \bar{\beta}. \tau$, and we can reconstruct

$$\frac{\Gamma, C, \epsilon \vdash v : \tau \quad \bar{\beta} \# \text{fv}(\Gamma)}{\Gamma, C, \epsilon \vdash v : \forall \bar{\beta}. \tau}$$

By the induction hypothesis and Lemma 7.6 there exists derivable $\Gamma', \epsilon \vdash_{\bar{\beta}'} v : \tau' / C'$, with $C \Vdash \forall \bar{\beta}'. \psi(\tau' / C') \sqsubseteq \text{gen}(\Gamma, C, \tau)$. But $C \Vdash \text{gen}(\Gamma, C, \tau) \sqsubseteq \forall \bar{\beta}. \tau$ by Lemma 7.5, therefore by Lemma 7.4 we have $C \Vdash \forall \bar{\beta}'. \psi(\tau' / C') \sqsubseteq \forall \bar{\beta}. \tau$, so this case holds by Lemma 7.7.

Case \forall -ELIM. In this case by definition of subtyping derivations we have that $e = v$ and $\sigma = \tau[\bar{\tau} / \bar{\beta}]$, and we can reconstruct

$$\frac{\Gamma, C, \epsilon \vdash v : \forall \bar{\beta}. \tau \quad C \Vdash [\bar{\tau} / \bar{\beta}]}{\Gamma, C, \epsilon \vdash v : \tau[\bar{\tau} / \bar{\beta}]}$$

By the induction hypothesis and Lemma 7.6, there exists derivable $\Gamma', \epsilon \vdash_{\bar{\beta}'} v : \tau' / C'$, with $C \Vdash \forall \bar{\beta}'. \psi(\tau' / C') \sqsubseteq \text{gen}(\Gamma, C, \forall \bar{\beta}. \tau)$. Let $\forall \bar{\beta}_0. \tau = \text{gen}(\Gamma, C, \forall \bar{\beta}. \tau)$ and $\bar{\beta}_1 = \{\beta \mid \beta \in \bar{\beta}_0 - \bar{\beta}\}$. Then $\tau[\bar{\tau} / \bar{\beta}] = \tau[\bar{\tau} / \bar{\beta}][\bar{\beta}_1 / \bar{\beta}_1]$. Let $\psi_0 = [\bar{\tau} / \bar{\beta}][\bar{\beta}_1 / \bar{\beta}_1]$ and let $\text{gen}(\Gamma, C, \tau[\bar{\tau} / \bar{\beta}]) = \forall \bar{\beta}_2. \tau[\bar{\tau} / \bar{\beta}]$. It is easy to show that $\bar{\beta}_2 \sqsubseteq \bar{\beta}_1 \cup \text{fv}(\bar{\tau})$; hence for arbitrary $[\bar{\tau}_2 / \bar{\beta}_2]$ such that $C \Vdash [\bar{\tau}_2 / \bar{\beta}_2]$, we have $\text{dom}([\bar{\tau}_2 / \bar{\beta}_2] \circ \psi_0) \subseteq \bar{\beta}_0$, hence there exists $\bar{\tau}_3$ such that $[\bar{\tau}_3 / \bar{\beta}_0] = [\bar{\tau}_2 / \bar{\beta}_2] \circ \psi_0$. Thus by Definition 7.4 and above facts there exists $[\bar{\tau}' / \bar{\beta}']$ such that the following entailments hold:

$$C \Vdash (\psi(C'))[\bar{\tau}' / \bar{\beta}'] \quad C \Vdash (\psi(\tau'))[\bar{\tau}' / \bar{\beta}'] \sqsubseteq \tau[\bar{\tau}_3 / \bar{\beta}_0]$$

the latter of which is to say by the above equivalences:

$$C \Vdash (\psi(\tau'))[\bar{\tau}' / \bar{\beta}'] \sqsubseteq (\tau[\bar{\tau} / \bar{\beta}])([\bar{\tau}_2 / \bar{\beta}_2])$$

so this case holds by Definition 7.4 and Lemma 7.7.

Case EVENT. In this case $e = \text{ev}(e_0)$ and $\sigma = \text{unit}$ and $H = H_0; \text{ev}(s)$ by definition of subtyping derivations, where we can reconstruct the following rule instance:

$$\frac{\Gamma, C, H_0 \vdash e_0 : \{s\}}{\Gamma, C, H_0; \text{ev}(s) \vdash \text{ev}(e_0) : \text{unit}}$$

Let $\text{genvars}(\Gamma, C, H_0, \{s\}) = \bar{\beta}$. Then by the induction hypothesis there exists derivable $\Gamma', H' \vdash_{\bar{\beta}'} \text{ev}(e_0) : \tau' / C'$, where for all $\psi_0 = [\bar{\tau} / \bar{\beta}]$ such that $C \Vdash [\bar{\tau} / \bar{\beta}]$ there exists $\psi' = [\bar{\tau}' / \bar{\beta}']$ such that

$$C \Vdash \psi' \circ \psi(C') \quad C \Vdash \psi' \circ \psi(\tau') \sqsubseteq \psi_0(\{s\}) \quad C \Vdash \psi' \circ \psi(H') \sqsubseteq \psi_0(H_0)$$

and by an instance of EVENT we can canonically infer

$$\frac{\Gamma', H' \vdash_{\bar{\beta}'} \text{ev}(e_0) : \tau' / C'}{\Gamma', H'; \text{ev}(\alpha) \vdash_{\bar{\beta}' \cup \{\alpha\}} \text{ev}(e') : \text{unit} / C' \wedge \tau' \sqsubseteq \{s\}}$$

Now, since α is canonically chosen fresh, by Lemma 7.1 $\alpha \# \bar{\beta}' \# \text{dom}(\psi)$ and

$$\text{fv}(\tau', C', H') \subseteq \text{dom}(\psi) \cup \bar{\beta}'$$

Therefore, letting $\psi'' = \psi' \circ [\psi_0(s)/\alpha]$:

$$\begin{aligned} \psi'' \circ \psi(\alpha) &= \psi_0(s) & \psi'' \circ \psi(C') &= \psi' \circ \psi(C') & \psi'' \circ \psi(\tau') &= \psi' \circ \psi(\tau') \\ \psi'' \circ \psi(H') &= \psi' \circ \psi(H') \end{aligned}$$

so by the above facts and Lemma 6.4:

$$C \Vdash \psi'' \circ \psi(C' \wedge \tau' \sqsubseteq \{\alpha\}) \quad C \Vdash \psi'' \circ \psi(H'; ev(\alpha)) \sqsubseteq \psi_0(H_0; ev(s))$$

and $\text{dom}(\psi'') = \bar{\beta}' \cup \{\alpha\}$, so this case holds by Definition 7.6.

Case Fix. In this case we have $e = \lambda_z x.e'$ and $H = \epsilon$ and $\tau = \tau_1 \xrightarrow{H'} \tau_2$ by definition of subtyping derivations, with the following as the last derivation step:

$$\frac{\Gamma; x : \tau_1; z : \tau_1 \xrightarrow{H_1} \tau_2, C, H_1 \vdash e' : \tau_2}{\Gamma, C, \epsilon \vdash \lambda_z x.e' : \tau_1 \xrightarrow{H_1} \tau_2}$$

Let $\psi_0 = [\tau_1/t_1][\tau_2/t_2][H_1/h] \circ \psi$. Since t_1, t_2, h are chosen fresh in canonical derivations, therefore $C \Vdash \psi_0(t_1 \xrightarrow{h} t_2) \sqsubseteq \tau_1 \xrightarrow{H_1} \tau_2$ and by Definition 7.4:

$$C \Vdash \psi_0(\Gamma'; x : t_1; z : t_1 \xrightarrow{h} t_2) \sqsubseteq (\Gamma; x : \tau_1; z : \tau_1 \xrightarrow{H_1} \tau_2)$$

Therefore, since $\text{genvars}(\Gamma; x : \tau_1; z : \tau_1 \xrightarrow{H_1} \tau_2, C, H_1, \tau_2) = \emptyset$, by the induction hypothesis a judgement of the following form is derivable:

$$\Gamma'; x : t_1; z : t_1 \xrightarrow{h} t_2, H' \vdash_{\bar{\beta}''} e' : \tau'/C'$$

where there exists $\psi' = [\bar{\tau}/\bar{\beta}'']$ such that

$$C \Vdash \psi' \circ \psi_0(C') \quad C \Vdash \psi' \circ \psi_0(H') \sqsubseteq H_1 \quad C \Vdash \psi' \circ \psi_0(\tau') \sqsubseteq \tau_2$$

But letting $\psi'' = \psi' \circ [\tau_1/t_1][\tau_2/t_2][H_1/h]$, equivalently:

$$C \Vdash \psi'' \circ \psi(C') \quad C \Vdash \psi'' \circ \psi(H') \sqsubseteq H_1 \quad C \Vdash \psi'' \circ \psi(\tau') \sqsubseteq \tau_2$$

And, freshness of $\{t_1, t_2, h\}$ means $\{t_1, t_2, h\} \# \bar{\beta}'' \# \text{dom}(\psi)$, therefore $\psi'' \circ \psi(t_1) = \tau_1$ and $\psi'' \circ \psi(t_2) = \tau_2$ and $\psi'' \circ \psi(h_1) = H_1$, thus

$$C \Vdash \psi'' \circ \psi(C' \wedge H' \sqsubseteq h \wedge \tau' \sqsubseteq t_2) \quad C \Vdash \psi'' \circ \psi(t_1 \xrightarrow{h} t_2) \sqsubseteq \tau_1 \xrightarrow{H_1} \tau_2$$

and $\text{dom}(\psi') = \bar{\beta}'' \cup \{t_1, t_2, h\}$. By an application of Fix:

$$\frac{\Gamma'; x : t_1; z : t_1 \xrightarrow{h} t_2, H' \vdash_{\bar{\beta}''} e' : \tau'/C'}{\Gamma', \epsilon \vdash_{\bar{\beta}'' \cup \{t_1, t_2, h\}} \lambda_z x.e' : \tau_1 \xrightarrow{h} t_2/C' \wedge H' \sqsubseteq h \wedge \tau' \sqsubseteq t_2}$$

so this case holds by Definition 7.6.

Case App. In this case $e = e_1 e_2$ and $H = H_1; H_2; H_3$ by definition of subtyping derivations, and as the last step in the derivation:

$$\frac{\Gamma, C, H_1 \vdash e_1 : \tau' \xrightarrow{H_3} \tau \quad \Gamma, C, H_2 \vdash e_2 : \tau'}{\Gamma, C, H_1; H_2; H_3 \vdash e_1 e_2 : \tau}$$

Let

$$\begin{aligned} \text{genvars}(\Gamma, C, H_1, \tau' \xrightarrow{H_3} \tau) &= \bar{\beta} & \text{genvars}(\Gamma, C, H_2, \tau') &= \bar{\beta}' \\ \text{genvars}(\Gamma, C, H_1; H_2; H_3, \tau) &= \bar{\beta}'' \end{aligned}$$

Let ψ' be an arbitrary substitution such that $\text{dom}(\psi') = \bar{\beta}\bar{\beta}'$ and $C \Vdash \psi'$, and let:

$$\psi_{\bar{\beta}} = \psi' \upharpoonright_{\bar{\beta}} \quad \psi_{\bar{\beta}'} = \psi' \upharpoonright_{\bar{\beta}'} \quad \psi_{\bar{\beta}''} = \psi' \upharpoonright_{\bar{\beta}''}$$

Then by Lemma 7.9 and the induction hypothesis there exists a canonically derivable judgement $\Gamma', H'_1 \vdash_{\bar{\beta}_1} e_1 : \tau_1/C_1$ and substitution $\psi_1 = [\bar{\tau}_1/\bar{\beta}_1]$ such that

$$C \Vdash \psi_1 \circ \psi(C_1) \quad C \Vdash \psi_1 \circ \psi(\tau_1) \sqsubseteq \psi_{\bar{\beta}}(\tau' \xrightarrow{H_3} \tau) \quad C \Vdash \psi_1 \circ \psi(H'_1) \sqsubseteq \psi_{\bar{\beta}}(H_1)$$

Also by Lemma 7.9 and the induction hypothesis there exists canonically derivable $\Gamma', H'_2 \vdash_{\bar{\beta}_2} e_2 : \tau_2/C_2$ and substitution $\psi_2 = [\bar{\tau}_2/\bar{\beta}_2]$ such that

$$C \Vdash \psi_2 \circ \psi(C_2) \quad C \Vdash \psi_2 \circ \psi(\tau_2) \sqsubseteq \psi_{\bar{\beta}'}(\tau') \quad C \Vdash \psi_2 \circ \psi(H'_2) \sqsubseteq \psi_{\bar{\beta}'}(H_2)$$

Note that

$$\begin{aligned} \psi_{\bar{\beta}}(H_1) &= \psi_{\bar{\beta}''}(H_1) & \psi_{\bar{\beta}'}(H_2) &= \psi_{\bar{\beta}''}(H_2) & \psi_{\bar{\beta}}(H_3) &= \psi_{\bar{\beta}''}(H_3) & \psi_{\bar{\beta}}(\tau) &= \psi_{\bar{\beta}''}(\tau) \\ \psi_{\bar{\beta}}(\tau') &= \psi_{\bar{\beta}'}(\tau') \end{aligned}$$

Now, let $\psi_3 = \psi_2 \circ [\psi_{\bar{\beta}''}(\tau)/t][\psi_{\bar{\beta}''}(H_3)/h]$. Taking $\bar{\beta}_1 \# \bar{\beta}_2$ wlog, by Lemma 7.1 it is the case that $\bar{\beta}_1 \# \bar{\beta}_2 \# \text{dom}(\psi)$ and

$$\text{fv}(\tau_1, H'_1, C_1) = \text{fv}(\Gamma) \cup \bar{\beta}_1 \quad \text{fv}(\tau_2, H'_2, C_2) = \text{fv}(\Gamma) \cup \bar{\beta}_2$$

therefore

$$\begin{aligned} \psi_3 \circ \psi(\tau_1) &= \psi_1 \circ \psi(\tau_1) & \psi_3 \circ \psi(H_1) &= \psi_1 \circ \psi(H_1) & \psi_3 \circ \psi(C_1) &= \psi_1 \circ \psi(C_1) \\ \psi_3 \circ \psi(\tau_2) &= \psi_2 \circ \psi(\tau_2) & \psi_3 \circ \psi(H_2) &= \psi_2 \circ \psi(H_2) & \psi_3 \circ \psi(t) &= \psi_{\bar{\beta}''}(\tau) \\ \psi_3 \circ \psi(h) &= \psi_{\bar{\beta}''}(H_3) \end{aligned}$$

Combining these equivalences with preceding facts and Lemma 6.4, we have

$$C \Vdash \psi_3 \circ \psi(H'_1; H'_2; h) \sqsubseteq \psi_{\bar{\beta}''}(H_1; H_2; H_3) \quad C \Vdash \psi_3 \circ \psi(t) \sqsubseteq \psi_{\bar{\beta}''}(\tau)$$

We can also assert

$$\begin{aligned} C \Vdash \psi_3 \circ \psi(C_1 \wedge C_2) & \quad C \Vdash \psi_{\bar{\beta}}(\tau' \xrightarrow{H_3} \tau) \sqsubseteq \psi_3 \circ \psi(\tau_2 \xrightarrow{h} t) \\ C \Vdash \psi_3 \circ \psi(\tau_1) & \sqsubseteq \psi_{\bar{\beta}}(\tau' \xrightarrow{H_3} \tau) \end{aligned}$$

so by transitivity of \sqsubseteq , etc., we have

$$C \Vdash \psi_3 \circ \psi(C_1 \wedge C_2 \wedge \tau_1) \sqsubseteq \tau_2 \xrightarrow{h} t$$

Finally, by APP we can derive

$$\frac{\Gamma', H'_1 \vdash_{\bar{\beta}_1} e_1 : \tau_1/C_1 \quad \Gamma', H'_2 \vdash_{\bar{\beta}_2} e_2 : \tau_2/C_2}{\Gamma', H'_1; H'_2; h \vdash_{\bar{\beta}_1 \bar{\beta}_2 \{t, h\}} e_1 e_2 : t/C_1 \wedge C_2 \wedge \tau_1 \sqsubseteq \tau_2 \xrightarrow{h} t}$$

and since $\text{dom}(\psi_3) = \bar{\beta}_1 \bar{\beta}_2 \{t, h\}$, this case holds.

Case IF. In this case $e = \text{if } e_1 \text{ then } e_2 \text{ else } e_3$, and $H = H_1; H_2$, and as the last step in the derivation:

$$\frac{\Gamma, C, H_1 \vdash e_1 : \text{bool} \quad \Gamma, C, H_2 \vdash e_2 : \tau \quad \Gamma, C, H_2 \vdash e_3 : \tau}{\Gamma, C, H_1; H_2 \vdash \text{if } e_1 \text{ then } e_2 \text{ else } e_3 : \tau}$$

Let

$$\text{genvars}(\Gamma, C, H_1, \text{bool}) = \bar{\beta} \quad \text{genvars}(\Gamma, C, H_2, \tau) = \bar{\beta}'$$

and taking arbitrary $\text{dom}(\psi') = \bar{\beta}\bar{\beta}'$ such that $C \Vdash \psi'$, define $\psi_{\bar{\beta}} = \psi' \upharpoonright_{\bar{\beta}}$ and $\psi_{\bar{\beta}'} = \psi' \upharpoonright_{\bar{\beta}'}$. Then by Lemma 7.9 and the induction hypothesis, there exist derivable judgments:

$$\Gamma', H'_1 \vdash_{\bar{\beta}_1} e_1 : \tau_1/C_1 \quad \Gamma', H'_2 \vdash_{\bar{\beta}_2} e_2 : \tau_2/C_2 \quad \Gamma', H'_3 \vdash_{\bar{\beta}_3} e_3 : \tau_3/C_3$$

and substitutions ψ_1, ψ_2, ψ_3 with $\text{dom}(\psi_1) = \bar{\beta}_1, \text{dom}(\psi_2) = \bar{\beta}_2, \text{dom}(\psi_3) = \bar{\beta}_3$ such that:

$$C \Vdash \psi_1 \circ \psi(C_1) \quad C \Vdash \psi_1 \circ \psi(H'_1) \sqsubseteq \psi_{\bar{\beta}}(H_1) \quad C \Vdash \psi_1 \circ \psi(\tau_1) \sqsubseteq \psi_{\bar{\beta}}(\text{bool})$$

$$C \Vdash \psi_2 \circ \psi(C_2) \quad C \Vdash \psi_2 \circ \psi(H'_2) \sqsubseteq \psi_{\bar{\beta}'}(H_2) \quad C \Vdash \psi_2 \circ \psi(\tau_2) \sqsubseteq \psi_{\bar{\beta}'}(\tau)$$

$$C \Vdash \psi_3 \circ \psi(C_3) \quad C \Vdash \psi_3 \circ \psi(H'_3) \sqsubseteq \psi_{\bar{\beta}'}(H_2) \quad C \Vdash \psi_3 \circ \psi(\tau_3) \sqsubseteq \psi_{\bar{\beta}'}(\tau)$$

Assume $\bar{\beta}_1 \# \bar{\beta}_2 \# \bar{\beta}_3$ wlog. Note the following equivalences:

$$\psi_{\bar{\beta}}(\text{bool}) = \text{bool} \quad \psi_{\bar{\beta}'}(\tau) = \psi'(\tau) \quad \psi_{\bar{\beta}}(H_1) = \psi'(H_1) \quad \psi_{\bar{\beta}'}(H_2) = \psi'(H_2)$$

and let $\psi_4 = \psi_1 \circ \psi_2 \circ \psi_3 \circ [\tau/t]$. Then since $\bar{\beta}_1 \# \bar{\beta}_2 \# \bar{\beta}_3 \# t \# \text{dom}(\psi)$ by assumption, canonical freshness of t , and Lemma 7.1, we have $\psi_4(t) = \tau$, and for all $i \in [1..3]$:

$$\psi_i \circ \psi(C_i) = \psi_4 \circ \psi(C_i) \quad \psi_i \circ \psi(H'_i) = \psi_4 \circ \psi(H'_i) \quad \psi_i \circ \psi(\tau_i) = \psi_4 \circ \psi(\tau_i)$$

By the above facts and equivalences and Lemma 6.4 we can then assert

$$C \Vdash \psi_4 \circ \psi(t) \sqsubseteq \psi'(\tau) \quad C \Vdash \psi_4 \circ \psi(H'_1; H'_2 | H'_3) \sqsubseteq \psi'(H_1; H_2)$$

$$C \Vdash \psi_4 \circ \psi(\tau_1) \sqsubseteq \text{bool} \quad C \Vdash \psi_4 \circ \psi(\tau_2) \sqsubseteq t \quad C \Vdash \psi_4 \circ \psi(\tau_3) \sqsubseteq t$$

$$C \Vdash \psi_4 \circ \psi(C_1 \wedge C_2 \wedge C_3)$$

and letting $C' = C_1 \wedge C_2 \wedge C_3 \wedge \tau_1 \sqsubseteq \text{bool} \wedge \tau_2 \sqsubseteq t \wedge \tau_3 \sqsubseteq t$, observe $C \Vdash \psi_4 \circ \psi(C')$, and by IF we can derive

$$\frac{\bar{\beta}_1 \# \bar{\beta}_2 \# \bar{\beta}_3 \quad \bar{\beta}_4 = \bar{\beta}_1 \bar{\beta}_2 \bar{\beta}_3 \cup \{t\} \quad \Gamma', H'_1 \vdash_{\bar{\beta}_1} e_1 : \tau_1/C_1 \quad \Gamma', H'_2 \vdash_{\bar{\beta}_2} e_2 : \tau_2/C_2 \quad \Gamma', H'_3 \vdash_{\bar{\beta}_3} e_3 : \tau_3/C_3}{\Gamma', H'_1; H'_2 | H'_3 \vdash_{\bar{\beta}_4} \text{if } e_1 \text{ then } e_2 \text{ else } e_3 : t/C'}$$

Therefore since $\text{dom}(\psi_4) = \bar{\beta}_1 \bar{\beta}_2 \bar{\beta}_3 \cup \{t\}$, this case holds.

Case SUB. In this case by definition of subtyping derivations, we have as the last step in the derivation:

$$\frac{\Gamma, C, H' \vdash e : \tau' \quad C \Vdash \tau' \sqsubseteq \tau \quad C \Vdash H' \sqsubseteq H}{\Gamma, C, H \vdash e : \tau}$$

Letting $genvars(\Gamma, C, H', \tau') = \bar{\beta}$ and taking arbitrary ψ' such that $dom(\psi') = \bar{\beta}$ and $C \Vdash \psi'$, by the induction hypothesis there exists derivable $\Gamma', C_0 \vdash_{\bar{\beta}_0} e : \tau_0/C_0$ and substitution ψ_0 such that

$$C \Vdash \psi_0 \circ \psi(C_0) \quad C \Vdash \psi_0 \circ \psi(H_0) \sqsubseteq \psi'(H') \quad C \Vdash \psi_0 \circ \psi(\tau_0) \sqsubseteq \psi'(\tau')$$

But by Lemma 7.8 we have

$$C \Vdash \psi_0 \circ \psi(H_0) \sqsubseteq \psi'(H) \quad C \Vdash \psi_0 \circ \psi(\tau_0) \sqsubseteq \psi'(\tau)$$

and since $genvars(\Gamma, C, H, \tau) = genvars(\Gamma, C, H', \tau')$ by Lemma 7.2, this case holds. Case CONST follows trivially, so the result follows. \square

7.3 Closure, consistency, and hextraction

Now we prove correctness of the closure, consistency, and effect extraction techniques. Since our closure and consistency algorithms are standard with respect to type constraints, we mainly appeal to existing results, especially those in Palsberg & O'Keefe (1995), to establish correctness in that regard. The focus of our proof development is on satisfaction of singleton and effect constraints, and extraction of effect terms from constraints via *hextract*. In fact, correctness of *hextract* provides a constructive proof of the satisfiability of closed, consistent effect constraints.

We begin by observing that closure does not change the meaning of constraints, i.e. a constraint is logically equivalent to its closure:

Lemma 7.11

$$C = close(C).$$

One benefit of closure is that it provides a way to split constraints into distinct type, effect, and singleton components, the solutions for which can be composed to solve the entire constraint. We formalize this with the following definition and Lemma, where $f|_S$ denotes the restriction of the domain of a function f to S as usual.

Definition 7.8

For each kind k , the k component of a constraint C is a constraint D such that

$$set(D) = \{\tau_1 \sqsubseteq \tau_2 \mid \tau_1 \sqsubseteq \tau_2 \in C \text{ and } \tau_1, \tau_2 : k\}$$

Given that D is the k component of C , the constraint C has a k solution iff there exists ρ such that $\rho \Vdash D$.

Lemma 7.12

Suppose C is closed, and ρ_1 , ρ_2 , and ρ_3 are *Type*, *Sing*, and *Eff* solutions of C , respectively. Then

$$(\rho_1|_{\mathcal{V}_{Type}}) \circ (\rho_2|_{\mathcal{V}_{Sing}}) \circ (\rho_3|_{\mathcal{V}_{Eff}}) \Vdash C$$

The following Lemma establishes the crucial property that effect and singleton constraints form a system of lower bounds on variables. This property implies that any consistent closure has a singleton and effect solution. It also suggests a technique

for deriving term representations of singletons and effects, by joining lower bounds of variables. This is not just convenient, but demonstrates that while effect constraint solution is not decidable in general, it is so for the restricted form generated by inference.

Lemma 7.13

Suppose that $\emptyset, H \vdash_{\bar{\beta}} e : \tau/C$ is derivable and $\tau_1 \sqsubseteq \tau_2 \in \text{close}(C)$. If $\tau_2 : \text{Sing}$ then τ_2 is a singleton variable α , and if $\tau_2 : \text{Eff}$ then τ_2 is an effect variable h .

Now, we show that the singleton component of a closed, consistent constraint has a solution that can be obtained via the bounds_C function.

Definition 7.9

Suppose that f is a partial mapping from type variables to types. Write $f \preceq \rho$ iff there exists ρ' such that $\rho = \rho' \circ f$.

Lemma 7.14

Suppose $\emptyset, H \vdash_{\bar{\beta}} e : \tau/C$ is derivable and $\text{close}(C)$ is consistent. Then there exists a Sing solution ρ of $\text{close}(C)$ such that $\text{bounds}_{\text{close}(C)} \upharpoonright_{\mathcal{V}_{\text{Sing}}} \preceq \rho$.

Next, we show that the effect component of a closed, consistent constraint has a solution, that can be obtained via the hextract function. We show that this is the case constructively, by demonstrating that hextract is defined on top-level effects, and that hextract is a solution of given constraints. The first step is to show that top-level effects have concrete lower bounds so that their term form is closed.

Definition 7.10

Let C be closed and consistent. Then τ_1 is a *component of τ_2 in C* iff τ_1 is a subterm of τ_2 , or there exists a subterm τ'_1 of τ_1 such that $\tau'_2 \sqsubseteq \tau'_1 \in C$ and τ_2 is a component of τ'_2 in C .

Definition 7.11

Let C be closed and consistent. Then the *abstract components of τ in C* , denoted $\text{abstract}(\tau, C)$ are the variable components of τ that have no lower bound in C . More precisely:

$$\text{abstract}(\tau, C) = \{\beta \mid \beta \text{ is a component of } \tau \text{ in } C \wedge \neg \exists \tau'. \tau' \sqsubseteq \beta \in C\}$$

On the basis of these definitions, we can now show that top-level effects have concrete lower bounds in inferred constraints, meaning that hextract is defined and closed on top-level effects. The result follows by a tedious and unilluminating induction on derivations. Essentially, the property holds because any effect variable occurring in a top-level effect is the upper bound of a function effect by definition of inference, and since top-level expressions are closed, the function must be a closed lambda abstraction, which have defined, concrete effects.

Lemma 7.15 (Top-level effects are concrete)

Suppose $\emptyset, H \vdash_{\bar{\beta}} e : \tau/C$ is derivable and $\text{close}(C)$ is consistent. Then H has no abstract components in $\text{close}(C)$.

$$\begin{array}{ll}
HX \ f \ C \ \epsilon & = \ \epsilon \\
HX \ f \ C \ \varpi(\alpha) & = \ \varpi(\text{bounds}_C \ \alpha) \\
HX \ f \ C \ h & = \ f \ C \ (\text{bounds}_C \ h) & \text{if } \text{bounds}_C \ h \text{ defined} \\
HX \ f \ C \ h & = \ f \ C \ h & \text{if } \text{bounds}_C \ h \text{ not defined} \\
HX \ f \ C \ H_1; H_2 & = \ (f \ C \ H_1); (f \ C \ H_2) \\
HX \ f \ C \ H_1|H_2 & = \ (f \ C \ H_1)|(f \ C \ H_2)
\end{array}$$

Fig. 15. The HX combinator.*Lemma 7.16*

Suppose $\emptyset, H \vdash_{\bar{\beta}} e : \tau/C$ is derivable and $\text{close}(C)$ is consistent. Then $(\text{hextract } \text{close}(C) \ H)$ is defined and closed.

Having established that hextract is defined on top-level effects, we now show that hextract is a solution of given constraints. The recursive definition of hextract makes a direct proof of this property difficult. Instead, we demonstrate the property via a combinator fixpoint construction. We define a nonrecursive combinator HX in Figure 15, that characterizes the solution of closed, consistent effect constraints, and then show that hextract is a fixpoint of this combinator. The first step is to show that HX correctly characterizes solutions of closed, consistent effect constraints, as follows.

Lemma 7.17

Suppose $\emptyset, H \vdash_{\bar{\beta}} e : \tau/C$ is derivable and $\text{close}(C)$ is consistent. If f is a fixpoint of HX , then there exists an *Eff* solution ρ of $\text{close}(C)$ such that $(f \ \text{close}(C)) \preceq \rho$.

Proof

Let D be the *Eff* component of $\text{close}(C)$, and suppose that f is a fixpoint of HX . Let ρ' be a *Sing* solution of $\text{close}(C)$, which must exist by Lemma 7.14, and let

$$g = \rho' \circ (f \ \text{close}(C))$$

To proceed, we number the clauses of HX from top to bottom, starting with 1. Since we assume that f is a fixpoint of HX , clauses 1–2 and 5–6 establish the following equalities:

$$\begin{array}{lll}
g(\epsilon) = \epsilon & g(\varpi(\alpha)) = \rho'(\varpi(\text{bounds}_C \ \alpha)) & g(H_1; H_2) = g(H_1); g(H_2) \\
g(H_1|H_2) = g(H_1)|g(H_2) & &
\end{array}$$

Thus, g is an interpretation, as defined in Figure 6. Furthermore, clause 3 establishes that for all $H \sqsubseteq h \in D$, it is the case that $g(H) \preceq g(h)$, meaning that g satisfies every constraint in D by Lemma 7.13, and is therefore a solution of D . Letting $\rho = g$, the result follows by Definition 7.9. \square

It would now be possible to show that any closed, consistent effect constraint has a solution, by showing that HX is monotonic over a partial ordering of solutions, since classic results in Tarski (1955) would then imply that HX has a fixpoint. Instead, we prove that hextract is a fixpoint of HX in Lemma 7.17, establishing

both correctness of *hextract* and existence of a solution of closed consistent effect constraints. We first prove the essential auxiliary Lemma for proving the tricky μ case of Lemma 7.17, and then Lemma 7.17 itself.

Lemma 7.18

Let $h \notin hs$ and $H' = \text{hextract}_C(h, hs)$. Then

$$\text{hextract}_C(H, hs \cup \{h\})[H'/h] = \text{hextract}_C(H, hs)$$

Proof

The proof proceeds by induction on the call tree of hextract_C and case analysis on H . In case $H = \epsilon$ or $H = v(\alpha)$ the result is immediate. In case $H = H_1|H_2$ or $H_1;H_2$ the result follows in a straightforward manner via the induction hypothesis. The more interesting case is that in which H is an effect variable h' . Here we consider subcases $h' \neq h$ and $h' = h$, where the former breaks down into subcases with $h' \in hs$ on the one hand and $h' \notin hs$ on the other. If $h' \neq h$ and $h' \in hs$, then by definition of hextract_C :

$$\text{hextract}_C(h', hs \cup \{h\})[H'/h] = \text{hextract}_C(h', hs) = h'$$

so the result follows.

On the other hand, if $h' \neq h$ and $h' \notin hs$, then by definition of hextract_C

$$\text{hextract}_C(h', hs \cup \{h\})[H'/h] = (\mu h'. \text{hextract}_C(\text{bounds}_C h', hs \cup \{h, h'\})) [H'/h]$$

and

$$\text{hextract}_C(h', hs) = \mu h'. \text{hextract}_C(\text{bounds}_C h', hs \cup \{h'\})$$

Let $H'' = \text{hextract}_C(\text{bounds}_C h', hs \cup \{h, h'\})$. Since $h' \neq h$ by assumption:

$$(\mu h'. H'') [H'/h] = \mu h'. (H'' [H'/h])$$

and by the induction hypothesis:

$$\text{hextract}_C(\text{bounds}_C h', hs \cup \{h'\}) = H'' [H'/h]$$

Hence

$$\text{hextract}_C(h', hs \cup \{h\}) [H'/h] = \text{hextract}_C(h', hs) = \mu h'. (H'' [H'/h])$$

so the result follows.

Now, consider the subcase $h' = h$. Then by the definition of hextract_C :

$$\text{hextract}_C(h, hs \cup \{h\}) = h$$

and clearly

$$\text{hextract}_C(h, hs \cup \{h\}) [H'/h] = H'$$

and $\text{hextract}_C(H, hs) = H'$ in this case, so the result follows. \square

Lemma 7.19

hextract is a fixpoint of HX .

Proof

We begin by noting the obvious point $\text{hextract}_C H = \text{hextract}_C(H, \emptyset)$. To prove the result, it suffices to show that for all C and for all H in the domain of HX $f C$

we have

$$HX \text{ hextract } C H = \text{hextract } C H$$

The proof proceeds by case analysis on H .

Case $H = \epsilon$. By definition of HX and hextract :

$$HX \text{ hextract } C \epsilon = \text{hextract } C \epsilon = \epsilon$$

so this case holds.

Case $H = \nu(\alpha)$. By definition of HX and hextract :

$$HX \text{ hextract } C \nu(\alpha) = \text{hextract } C \nu(\alpha) = \nu(\text{bounds}_C \alpha)$$

so this case holds.

Case $H = (H_1|H_2)$. By definition of HX :

$$HX \text{ hextract } C (H_1|H_2) = (\text{hextract } C H_1)|(\text{hextract } C H_2)$$

and by definition of hextract :

$$\text{hextract } C (H_1|H_2, C) = \text{hextract}(H_1, C)|\text{hextract}(H_2, C)$$

so this case holds. Case $H = H_1;H_2$ follows similarly.

Case $H = h$. By definition of HX :

$$HX \text{ hextract } C h = \text{hextract } C (\text{bounds}_C h)$$

and by definition of hextract :

$$\text{hextract } C h = \text{hextract}_C(h, \emptyset) = \mu h.\text{hextract}_C((\text{bounds}_C h), \{h\})$$

Now, let $H' = \mu h.\text{hextract}_C(\text{bounds}_C h, \{h\})$. By Lemma 7.18 we have

$$\text{hextract}_C((\text{bounds}_C h), \emptyset) = (\text{hextract}_C(\text{bounds}_C h, \{h\})) [H'/h]$$

and by properties of trace effect equivalence noted in Lemma 3.1:

$$\rho((\text{hextract}_C((\text{bounds}_C h), \{h\})) [H'/h]) = \rho(H')$$

so the result follows in this case, which was the last to be proven. \square

Having shown that hextract is a fixpoint of HX , we can immediately demonstrate that hextract implements a solution of given constraints when applied to top-level effects.

Lemma 7.20 (Correctness of hextract)

Suppose $\emptyset, H \vdash_{\bar{\beta}} e : \tau/C$ is derivable and $\text{close}(C)$ is consistent. Then C has an Eff solution ρ such that $(\text{hextract } \text{close}(C)) \preceq \rho$.

Proof

By Lemma 7.19, hextract is a fixpoint of HX . The result follows immediately by Lemma 7.17. \square

Composition of the above Lemmas establishing solvability of singleton and effect constraints, with standard results for closure and consistency of type constraints, allows us to obtain that consistency of closure is equivalent to satisfiability.

Lemma 7.21 (Correctness of Closure)

If $\emptyset, H \vdash_{\bar{\beta}} e : \tau / C$ is derivable, then C has a solution iff $\text{close}(C)$ is consistent.

Proof

That solvability of C entails consistency of C follows via *reductio ad absurdum*, since if we assume on the contrary that $\text{close}(C)$ is inconsistent, this implies that $\text{close}(C)$ contains a clashing constraint and thus the contradiction by Lemma 6.2. The other direction proceeds as follows. By Lemma 7.14, there exists a *Sing* solution of $\text{close}(C)$. By Lemma 7.20, there exists an *Eff* solution of $\text{close}(C)$. Given this, we can show that $\text{close}(C)$ has a *Type* solution by appeal to standard results, particularly those in Palsberg and O’Keefe (1995), since our closure and consistency techniques for *Type* constraints are standard. The result follows by Lemma 7.12 and Lemma 7.11. \square

Finally, we can tie all the results together to state high-level properties of the system. Since we have demonstrated completeness of inference in Lemma 7.10 and shown that consistent closure is equivalent to satisfiability in Lemma 7.21, we immediately obtain a completeness result for inference.

Corollary 7.1 (Implementation Completeness)

If a closed expression e is typable in the system of Figure 9, then a judgement $\emptyset, H \vdash_{\bar{\beta}} e : \tau / C$ is derivable and $\text{close}(C)$ is consistent.

However, since we have only shown that *hextract* is a fixpoint of HX , not the *least* fixpoint, the best we can demonstrate for the algorithm is soundness. We informally conjecture that *hextract* implements a least solution of the given constraint, since it adds no extraneous constraints in the construction, but in any case soundness is sufficient to obtain a type safety result for the implementation, as follows.

Lemma 7.22 (Implementation Soundness)

Suppose $\emptyset, H \vdash_{\bar{\beta}} e : \tau / C$ is derivable and $\text{close}(C)$ is consistent. Then $\emptyset, H, C \vdash e : \tau$ is satisfiable, and there exists a solution ρ of C such that $(\text{hextract } \text{close}(C) H) = \rho(H)$.

Proof

If $\emptyset, H \vdash_{\bar{\beta}} e : \tau / C$ is derivable then so is $\emptyset, H, C \vdash e : \tau$ by Lemma 7.3. If $\text{close}(C)$ is consistent then C has a solution ρ' by Lemma 7.21, and by Lemma 7.20 C has an *Eff* solution ρ'' such that $(\text{hextract } \text{close}(C)) \leq \rho''$. Let $\rho = \rho' \circ \rho'' \upharpoonright_{\gamma_{\text{Eff}}}$. Since $(\text{hextract } \text{close}(C) H)$ is defined and closed by Lemma 7.16, therefore,

$$\rho''(H) = (\text{hextract } \text{close}(C) H) = \rho(H)$$

and ρ is a solution of C by Lemma 7.12. \square

8 The stack-based variation

In this section we define a stack-based variation on the framework of the previous sections, allowing properties of the run-time stack at a program point to be verified at compile-time. Instead of keeping track of *all* events, only events for functions on the current call stack are maintained, in keeping with a general stack-based

$S ::= \text{nil} \mid S :: \eta$ <i>history stacks</i>		
$S, (\lambda_{z,x}.e)v$	\rightsquigarrow	$S :: \epsilon, \cdot e[v/x][\lambda_{z,x}.e/z] \cdot$ (β)
$S :: \eta, ev(c)$	\rightsquigarrow	$S :: \eta; ev(c), ()$ $(event)$
$S :: \eta, ev_{\phi}(c)$	\rightsquigarrow	$S :: \eta; ev_{\phi}(c), ()$ $(check)$
		if $\Pi(\phi(c), (\widehat{S :: \eta})ev_{\phi}(c))$
$S :: \eta, \cdot v \cdot$	\rightsquigarrow	S, v (pop)

Fig. 16. Semantics of λ_{trace}^S (selected rules).

security model (as in, e.g., Jensen *et al.* 1999). Assertions ϕ in this framework are run-time assertions about the *active* event sequence, not all events. While the stack-based model is somewhat distinct from the trace-based model, we show that this variation requires only a minor “postprocessing” of inferred trace effects for a sound analysis. There are results showing how it is possible to directly model-check stack properties of a Push-Down Automata (PDA) computation Esparza *et al.* (2001); our approach based on postprocessing trace effects represents an alternative method, which may also prove useful for modelling features such as exceptions: the raising of an exception implies any subsequent effect is discarded.

We note that our system captures a more fine-grained stack-based model than has been previously proposed; in particular, the use of stacks of traces allows the ordering of events within individual stack frames to be taken into account, along with the ordering of frames themselves.

8.1 Syntax and semantics

The values, expressions, and evaluation contexts of λ_{trace}^S are exactly those of λ_{trace} , extended with an expression form $\cdot e \cdot$ and evaluation context form $\cdot E \cdot$ for delimiting the scope of function activations. We impose the requirement that in any function $\lambda_{z,x}.e$, there exist no subexpressions $\cdot e' \cdot$ of e . The operational semantics of λ_{trace}^S is a relation on *configurations* S, e , where S ranges over stacks of traces, defined in Figure 16. The active security context for run-time checks is obtained from the trace stack in configurations, by appending traces in the order they appear in the stack. To formalize this, we define the notation $[S]$ inductively as follows:

$$[\text{nil}] = \epsilon \qquad [S :: \eta] = [S]; \eta$$

Abusing notation, we write \hat{S} as syntactic sugar for $[S]$.

Selected rules for the reduction relations \rightsquigarrow and \rightarrow on configurations are then specified in Figure 16 (those not specified are obtained by replacing metavariables η with S in reductions for other expression forms in Figure 2). The trace interpretation function Π is defined as for λ_{trace} .

$$\begin{aligned}
\text{stackify}(\epsilon) &= \epsilon \\
\text{stackify}(\epsilon; H) &= \text{stackify}(H) \\
\text{stackify}(ev(c); H) &= ev(c); \text{stackify}(H) \\
\text{stackify}(h; H) &= h | \text{stackify}(H) \\
\text{stackify}((\mu h.H_1); H_2) &= (\mu h. \text{stackify}(H_1)) | \text{stackify}(H_2) \\
\text{stackify}((H_1 | H_2); H) &= \text{stackify}(H_1; H) | \text{stackify}(H_2; H) \\
\text{stackify}((H_1; H_2); H_3) &= \text{stackify}(H_1; (H_2; H_3)) \\
\text{stackify}(H) &= \text{stackify}(H; \epsilon)
\end{aligned}$$

Fig. 17. The *stackify* algorithm.

8.2 Stackified trace effects

Although λ_{trace}^S uses stack rather than trace contexts at run-time, we are able to use the type and verification framework developed previously, assigning types to λ_{trace}^S expressions in the same manner as λ_{trace} expressions. The only additional requirement will be to process trace effects to *stackify* them, yielding an approximation of the stack contexts that will evolve at run-time. The trick is to use μ -delimited scope in trace types, since this corresponds to function scope in inferred types as discussed in Section 7, and function activations and deactivations induce pushes and pops at run-time. The *stackify* algorithm is defined inductively in Figure 17. This algorithm works over effects that are sequences; for effects that are not sequences, the last clause massages into sequence form.

The last three clauses use trace effect equalities characterized in Section 3 to “massage” trace effects into appropriate form. Observe that the range of *stackify* consists of trace effects that are all tail-recursive; stacks are therefore finite-state transition systems and more efficient model-checking algorithms are possible Esparza *et al.* (2001).

Example 8.1

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

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

In the third example, since $\mu h.a$ precedes b , but $\mu h.a$ denotes the effect of a function call, stackification specifies that no events precede b since a will be popped before encountering b . In the last example, since b is preceded by h , which represents recursive μ -scope and hence a recursive call, any events preceding b will be popped before b is encountered, hence stackification specifies that no events precede b . The ϵ s in the last two examples are artifacts of the transformation, that could be cleaned up with some minor alterations to *stackify*.

8.3 Properties

In this section we prove a trace approximation result for $\lambda_{\text{trace}}^{\text{S}}$ (Theorem 8.1), showing how stackified effects generated by λ_{trace} type inference approximate program trace behavior in the $\lambda_{\text{trace}}^{\text{S}}$ model. The result is obtained by simulation of $\lambda_{\text{trace}}^{\text{S}}$ in λ_{trace} (Lemma 8.8), and reuse of trace approximation (Corollary 6.3) for λ_{trace} . Because we are concerned only with showing that stackification is correct, by proving that it yields an approximation of program event traces, we omit consideration of checks, since simulation of checks would be a needless complication.

As the basis of the simulation, we define an embedding $\langle e \rangle$ of $\lambda_{\text{trace}}^{\text{S}}$ expressions in λ_{trace} expressions; expressions are essentially unchanged, except that distinguished events push and pop are inserted in programs to mark function invocation and return. Later, we will define a trace transformation that deduces the stack context from the push and pop markers.

Definition 8.1

Let c_{dummy} be a distinguished singleton constant, and let

$$\text{push} \triangleq \text{ev}_{\text{push}}(c_{\text{dummy}}) \qquad \text{pop} \triangleq \text{ev}_{\text{pop}}(c_{\text{dummy}})$$

Then the $\lambda_{\text{trace}}^{\text{S}}$ to λ_{trace} encoding of expressions, denoted $\langle e \rangle$, is defined inductively as follows:

$$\begin{aligned} \langle \lambda_x.e \rangle &= \lambda_x.\text{push}; (\lambda x.\text{pop}; x) \langle e \rangle \\ \langle \cdot e \cdot \rangle &= (\lambda x.\text{pop}; x) \langle e \rangle \\ \langle \mathbf{c} \rangle &= \mathbf{c} \\ \langle \text{ev}(e) \rangle &= \text{ev}(\langle e \rangle) \\ \langle x \rangle &= x \\ \langle e_1 e_2 \rangle &= \langle e_1 \rangle \langle e_2 \rangle \\ \langle \text{if } e_1 \text{ then } e_2 \text{ else } e_3 \rangle &= \text{if } \langle e_1 \rangle \text{ then } \langle e_2 \rangle \text{ else } \langle e_3 \rangle \\ \langle \text{let } x = v \text{ in } e \rangle &= \text{let } \langle x \rangle = \langle v \rangle \text{ in } \langle e \rangle \end{aligned}$$

Given this encoding, it is clear that traces of images in the encoding will be of a particular form ς —every pop event has a matching push, with possible nestings, corresponding to the activation control flow structure. We specify this form grammatically:

$$\begin{aligned} s &::= \text{push}; \text{pop} \mid \text{push}; s; \text{pop} \mid \epsilon \mid \text{ev}(c) \mid s; s & \text{ev}(c) \notin \{\text{push}, \text{pop}\} \\ \varsigma &::= s \mid \text{push} \mid \varsigma; \varsigma \end{aligned}$$

The point of the encoding is that a stack context can faithfully be obtained from traces ς . The proof of simulation will also require that simulation of stack contexts in traces ς include a reflection of stack structure—every stack frame connector $::$ should be simulated by an unmatched push. To this end, we define the following.

Definition 8.2

The stack context simulation function sctx is defined as

$$\begin{aligned} \text{sctx}(\epsilon) &= \epsilon \\ \text{sctx}(\text{push}) &= \epsilon \\ \text{sctx}(ev(c)) &= ev(c) \\ \text{sctx}(\text{push}; \text{pop}) &= \epsilon \\ \text{sctx}(\text{push}; \zeta; \text{pop}) &= \epsilon \\ \text{sctx}(\zeta_1; \zeta_2) &= \text{sctx}(\zeta_1) \text{sctx}(\zeta_2) \end{aligned}$$

A trace ζ simulates a stack S , written $S \sim \zeta$, iff the relation can be derived given the following rules:

$$\frac{\text{sctx}(\zeta) = \epsilon}{\text{nil} \sim \zeta} \qquad \frac{S \sim \zeta \quad \text{sctx}(s) = \hat{\eta}}{S :: \eta \sim \zeta; \text{push}; s}$$

Since stack traces ζ are special forms of trace effects, interpretations and equivalence of them are defined as for trace effects in general. Thus, we can characterize sctx with the following Lemmas.

Lemma 8.1

If $S \sim \zeta$ then $\hat{S} = \text{sctx}(\zeta)$.

Lemma 8.2

If $\zeta = \zeta'$ then $\text{sctx}(\zeta) = \text{sctx}(\zeta')$.

The following auxiliary lemma will allow us to “pop” the simulated stack by postpending a pop event, which is essential to simulating λ_{trace}^S using λ_{trace} machinery.

Lemma 8.3

If $S :: \eta \sim \zeta$, then $S \sim \zeta; \text{pop}$.

Proof

By definition of \sim , we have that $\zeta = (\zeta'; \text{push}; s)$ where $\text{sctx}(s) = \hat{\eta}$ and $S \sim \zeta'$. The proof then proceeds by case analysis on S .

Case $S = \text{nil}$. Since $S \sim \zeta'$, therefore $\text{sctx}(\zeta') = \epsilon$ by definition of \sim . But also $\text{sctx}(\text{push}; s; \text{pop}) = \epsilon$ by definition of sctx , and also by definition of sctx we have $\text{sctx}(\zeta; \text{pop}) = \text{sctx}(\zeta') \text{sctx}(\text{push}; s; \text{pop})$, i.e. $\text{sctx}(\zeta; \text{pop}) = \epsilon$, so $S \sim \zeta; \text{pop}$ in this case by definition of \sim .

Case $S = S' :: \eta$. Since $S \sim \zeta'$, therefore $\zeta' = \zeta''; \text{push}; s'$ where $S' \sim \zeta''$ and $\text{sctx}(s') = \hat{\eta}$ by definition of \sim . But since $\text{sctx}(\text{push}; s; \text{pop}) = \epsilon$ by definition of sctx , and also $\text{sctx}(s'; \text{push}; s; \text{pop}) = \text{sctx}(s') \text{sctx}(\text{push}; s; \text{pop})$, therefore $\text{sctx}(s'; \text{push}; s; \text{pop}) = \text{sctx}(s') = \hat{\eta}$, so $S \sim \zeta; \text{pop}$ in this case by definition of \sim . \square

After the following auxiliary lemma, we show that \rightsquigarrow reduction is simulated in the encoding (\cdot) . Note that the proof allows one step of λ_{trace}^S reduction to be simulated by multiple steps in λ_{trace} reduction. The essential feature of the result is the simulation of stacks by traces ζ .

Lemma 8.4

$$\langle e \rangle[\langle v \rangle/x] = \langle e[v/x] \rangle.$$

Lemma 8.5

Given $\varsigma_\epsilon = \varsigma$ such that $S \sim \varsigma_\epsilon$. If $S, e \rightsquigarrow S', e'$ then $\varsigma, \langle e \rangle \rightarrow^* \varsigma', \langle e' \rangle$ where there exists $\varsigma'_\epsilon = \varsigma'$ such that $S' \sim \varsigma'_\epsilon$.

Proof

By case analysis on the rule form of the reduction $S, e \rightsquigarrow S', e'$.

Case β . In this case $e = (\lambda_z x. e_0)v$, $e' = \cdot e_0[v/x][e/z]$, and $S' = S :: \epsilon$. By Definition 8.1 we have

$$\langle e \rangle = \langle (\lambda_z x. e_0) \rangle \langle v \rangle \quad \langle \lambda_z x. e_0 \rangle = \lambda_z x. \text{push}; (\lambda x. \text{pop}; x) \langle e_0 \rangle$$

and since $\langle v \rangle$ is a value by Definition 8.1, therefore by β :

$$\begin{aligned} \varsigma, (\lambda_z x. \text{push}; (\lambda x. \text{pop}; x) \langle e_0 \rangle) \langle v \rangle \\ \rightsquigarrow \\ \varsigma, (\text{push}; (\lambda x. \text{pop}; x) \langle e_0 \rangle) [\langle v \rangle/x] [\langle \lambda_z x. e_0 \rangle/z] \end{aligned}$$

But by definition of substitution and Lemma 8.4:

$$\begin{aligned} (\text{push}; (\lambda x. \text{pop}; x) \langle e_0 \rangle) [\langle v \rangle/x] [\langle \lambda_z x. e_0 \rangle/z] \\ = \\ \text{push}; (\lambda x. \text{pop}; x) \langle e_0[v/x][\lambda_z x. e_0/z] \rangle \end{aligned}$$

and by *context*, *event*, and *seq*:

$$\begin{aligned} \varsigma, \text{push}; (\lambda x. \text{pop}; x) \langle e_0[v/x][\lambda_z x. e_0/z] \rangle \\ \rightarrow^* \\ \varsigma, \text{push}, (\lambda x. \text{pop}; x) \langle e_0[v/x][\lambda_z x. e_0/z] \rangle \end{aligned}$$

But since $S \sim \varsigma_\epsilon$ by assumption, therefore $(S :: \epsilon) \sim (\varsigma_\epsilon; \text{push}; \epsilon)$ by definition of \sim , and by Definition 8.1:

$$\langle \cdot e_0[v/x][\lambda_z x. e_0/z] \cdot \rangle = (\lambda x. \text{pop}; x) \langle e_0[v/x][\lambda_z x. e_0/z] \rangle$$

so this case holds, since $\varsigma_\epsilon; \text{push}; \epsilon = \varsigma; \text{push}$.

Case *pop*. In this case $e = \cdot v$, $S = S' :: \eta$, and $e' = v$. By Definition 8.1:

$$\langle \cdot v \cdot \rangle = (\lambda x. \text{pop}; x) \langle v \rangle$$

and since $\langle v \rangle$ is a value by Definition 8.1, by β :

$$\varsigma, (\lambda x. \text{pop}; x) \langle v \rangle \rightsquigarrow \varsigma, \text{pop}; \langle v \rangle$$

and by *event*, *seq*, and *context*:

$$\varsigma, \text{pop}; \langle v \rangle \rightarrow^* \varsigma; \text{pop}, \langle v \rangle$$

Furthermore, by assumption $S' :: \eta \sim \varsigma_\epsilon$, therefore $S' \sim \varsigma_\epsilon; \text{pop}$ by Lemma 8.3, so this case holds since $\varsigma_\epsilon; \text{pop} = \varsigma; \text{pop}$.

The rest of the proof follows in a straightforward manner by homomorphism of $\langle e \rangle$ in the remaining cases. \square

To extend the simulation result to \rightarrow reduction, the encoding (\cdot) needs to be extended to evaluation contexts.

Definition 8.3

The encoding (\cdot) is extended to evaluation contexts as follows, where a homomorphic extension to other context forms is elided:

$$\begin{aligned} ([\]) &= [\] \\ (\cdot E \cdot) &= (\lambda x. \text{pop}; x)(E) \\ (Ee) &= (E)(e) \\ (vE) &= (v)(E) \\ (\text{if } E \text{ then } e_1 \text{ else } e_2) &= \text{if } (E) \text{ then } (e_1) \text{ else } (e_2) \\ &\vdots \end{aligned}$$

After the following auxiliary lemma, we prove a simulation result for \rightarrow reduction, which follows easily on the basis of Lemma 8.5. This is immediately followed by the main simulation result, for \rightarrow^* reduction.

Lemma 8.6

$$(E[e]) = (E)[(e)].$$

Lemma 8.7

Given $\varsigma_\epsilon = \varsigma$ such that $S \sim \varsigma_\epsilon$. If $S, E[e] \rightarrow S', E[e']$ then $\varsigma, (E[e']) \rightarrow^* \varsigma', (E[e'])$ where there exists $\varsigma'_\epsilon = \varsigma'$ such that $S' \sim \varsigma'_\epsilon$.

Proof

By *context* it is the case that $S, e \rightsquigarrow S', e'$. Therefore $\varsigma, (e) \rightarrow^* \varsigma', (e')$ where there exists $\varsigma'_\epsilon = \varsigma'$ such that $S' \sim \varsigma'_\epsilon$, by Lemma 8.5. This means $\varsigma, (E)[(e)] \rightarrow^* \varsigma', (E)[(e')]$ by *context*, and since $(E)[(e)] = (E[e])$ and $(E)[(e')] = (E[e'])$ by Lemma 8.6, the result follows. \square

Lemma 8.8 (λ_{trace}^S Simulation)

If $\text{nil}, e \rightarrow^* S, e'$, then $\epsilon, (e) \rightarrow^* \varsigma, (e')$ with $\hat{S} = \text{sctx}(\varsigma)$.

Proof

Since $\text{nil} \sim \epsilon$ by definition of \sim , therefore by Lemma 8.7 and induction on the length of the reduction $\text{nil}, e \rightarrow^* S, e'$, we can assert $\epsilon, (e) \rightarrow^* \varsigma, (e')$ where there exists $\varsigma_\epsilon = \varsigma$ such that $S \sim \varsigma_\epsilon$. By Lemma 8.1 it is the case that $\text{sctx}(\varsigma_\epsilon) = \hat{S}$, so by Lemma 8.2 the result follows. \square

Having established the dynamic simulation, we now turn to stackification of trace effects. Our strategy is to show that the sctx transformation of traces ς is contained in the interpretation of stackified effect approximations of ς . The simulation result just proved will then allow us to make the bridge to λ_{trace}^S stack contexts. First, it will ease matters to define an alternate transformation of traces ς , one that “lines up” better with stackification. We do so as follows; note the structural similarity with stackification:

Definition 8.4

Let ζ trace sctxs transformation be defined inductively as follows:

$$\begin{aligned}
 \text{sctxs}(\epsilon) &= \epsilon \\
 \text{sctxs}(\epsilon; \zeta) &= \text{sctxs}(\zeta) \\
 \text{sctxs}(\text{push}; \zeta) &= \text{sctxs}(\zeta) \\
 \text{sctxs}(ev(c); \zeta) &= ev(c); \text{sctxs}(\zeta) \\
 \text{sctxs}((\text{push}; \zeta; \text{pop}); \zeta') &= \text{sctxs}(\zeta) | \text{sctxs}(\zeta') \\
 \text{sctxs}((\text{push}; \text{pop}); \zeta) &= \text{sctxs}(\zeta) \\
 \text{sctxs}((\zeta_1; \zeta_2); \zeta_3) &= \text{sctxs}(\zeta_1; (\zeta_2; \zeta_3)) \\
 \text{sctxs}(\zeta) &= \text{sctxs}(\zeta; \epsilon)
 \end{aligned}$$

We immediately observe that sctx transformations are contained in trace effect interpretations of sctxs transformations:

Lemma 8.9

$$\text{sctx}(\zeta) \in \llbracket \text{sctxs}(\zeta) \rrbracket.$$

Now, the crux of our technique will be to use push and pop events to line up traces ζ with effects H , and inductively show that the sctxs transformation is approximated by stackification. However, we want stackification to be more “aware” of push and pop events, and in particular to remove them in stackification, in keeping with the sctxs transformation, so we define an alternate function *stackify'* with these modifications. The connection with *stackify* is easily made later, since the modifications are minor; this is because μ -scope in effect always corresponds to function scope, which is always delimited by push and pop events in images of $\llbracket \cdot \rrbracket$ (as we prove below).

Definition 8.5

Let *stackify'* be defined equivalently to *stackify*, except for the μ case, defined as follows:

$$\text{stackify}'((\mu h.\text{push}; H_1; \text{pop}); H_2) = (\mu h.\text{stackify}'(H_1)) | \text{stackify}'(H_2)$$

The following definition formalizes this special form of effects, inferred for images of $\llbracket \cdot \rrbracket$, as the *stackifiable* property of effects.

Definition 8.6

H is *stackifiable* iff one of the following conditions holds inductively:

1. $H \equiv h$
2. $H \equiv ev(c)$ and $ev(c) \neq \text{push}, \text{pop}$
3. $H \equiv H_1 | H_2$ and both H_1 and H_2 are stackifiable
4. $H \equiv H_1; H_2$ and both H_1 and H_2 are stackifiable
5. $H \equiv \mu h.\text{push}; H; \text{pop}$ and H is stackifiable

Now, following an important auxiliary lemma necessary to deal with μ -bound effects, we show that if H approximates a trace ζ , then *stackify'*(H) approximates sctxs(ζ).

Lemma 8.10

Given stackifiable H . Then

$$\text{stackify}'(H[\mu h.\text{push}; H; \text{pop}/h]) = \text{stackify}'(H)[\mu h.\text{stackify}'(H)/h]$$

Proof

Straightforward by induction on h . \square

Lemma 8.11

For all H , define

$$\llbracket H \rrbracket_{pre} = \{\theta \mid \theta \in \llbracket H \rrbracket \text{ and } \theta \neq (\theta' \downarrow)\}$$

Then, stackifiable H and $\hat{\zeta} \in \llbracket H \rrbracket$ implies $\llbracket \text{sctxs}(\zeta) \rrbracket_{pre} \subseteq \llbracket \text{stackify}'(H) \rrbracket$.

Proof

(Sketch) By case analysis on the match forms of ζ in the definition clauses of sctxs and induction on the recursive call tree of $\text{sctxs}(\zeta)$. By the definitions of sctxs and $\text{stackify}'$, and correspondence of the match forms in the defining clauses, it is easy (and tedious) to “line up” the different cases of ζ and H , and demonstrate the result by induction. The most interesting case is $\zeta = \text{push}; \zeta'; \text{pop}; \zeta''$. It follows in this case that H must be of the form $H = (\mu h.\text{push}; H'; \text{pop}); H''$, where $\text{push}; \zeta'; \text{pop} \in \llbracket \mu h.\text{push}; H'; \text{pop} \rrbracket$ and $\hat{\zeta}'' \in \llbracket H'' \rrbracket$, since the push and pop events occur nowhere but at the limits of μ -scope in stackifiable effects, hence $\hat{\zeta}' \in \llbracket H'[\mu h.\text{push}; H'; \text{pop}/h] \rrbracket$. Therefore, by the induction hypothesis:

$$\begin{aligned} \llbracket \text{sctxs}(\zeta') \rrbracket_{pre} &\subseteq \text{stackify}'(H'[\mu h.\text{push}; H'; \text{pop}/h]) \\ \llbracket \text{sctxs}(\zeta'') \rrbracket_{pre} &\subseteq \text{stackify}'(H'') \end{aligned}$$

Since

$$\begin{aligned} \text{sctxs}(\text{push}; \zeta'; \text{pop}; \zeta'') &= \text{sctxs}(\zeta') | \text{sctxs}(\zeta'') \\ \text{stackify}'((\mu h.\text{push}; H'; \text{pop}); H'') &= \mu h.\text{stackify}'(H') | \text{stackify}'(H'') \end{aligned}$$

It remains to be shown that $\llbracket \text{sctxs}(\zeta') \rrbracket_{pre} \subseteq \llbracket \mu h.\text{stackify}'(H') \rrbracket$. Observing the equality $\mu h.\text{stackify}'(H') = \text{stackify}'(H')[\mu h.\text{stackify}'(H')/h]$ implied by Lemma 3.1, we can assert $\mu h.\text{stackify}'(H') = \text{stackify}'(H'[\mu h.\text{push}; H'; \text{pop}/h])$ by Lemma 8.10. The result follows. \square

Now, we show that type inference generates stackifiable effects. The typing rules actually do not “cleanly” place push and pop events at the limits of μ -scope, but rather there will be intervening cruft such as ϵ s and dummy μ -bindings, due to the nature of the encoding of functions. Therefore, a simple transformation f is necessary to obtain an equivalent stackifiable form of effects.

Lemma 8.12

Given satisfiable $\emptyset, H \vdash_{\beta} e : \tau/C$ for closed e . Then $\emptyset, H' \vdash_{\beta'} (\!| e \!|) : \tau/C'$ is satisfiable, where there exists a function f such that each of the following conditions hold:

1. $f(\text{hextract } C \ H')$ is stackifiable
2. $f(\text{hextract } C \ H') = (\text{hextract } C \ H')$
3. $\text{stackify}'(\text{hextract } C \ H) = \text{stackify}'(f(\text{hextract } C' \ H'))$

Proof

(Sketch) Observe that any effect is stackifiable if the only occurrence of push and pop is at the limits of μ -bound subeffects, and that all μ -bound subeffects have such delimitation. Also, the only place where μ -bound effects occur in $(\text{hextract } C \ H)$ will be in *hextraction* of effect variable components of H in C (in the sense of Definition 7.10). Now, if e is closed, and H has a variable component h , then it is demonstrable by definition of type inference and Lemma 7.15 that h must have the effect H_λ of one or more concrete functions $\lambda_z x.e$ as lower bounds. Furthermore, since $\llbracket \lambda_z x.e \rrbracket = \lambda_z x.\text{push}; (\lambda x.\text{pop}; x)\llbracket e \rrbracket$, the lower bounds induced by the transformation $\llbracket \lambda_z x.e \rrbracket$ will be equivalent to $\text{push}; H_\lambda; \text{pop}$ (this is approximate; the described effect transformation is also carried recursively through H_λ , as the expression transformation is carried recursively through the function body). In short, if $(\text{hextract } C \ H)$ contains a subeffect $\mu h.H_\lambda$, then $f(\text{hextract}(H', C'))$ will have the same shape as $(\text{hextract } C \ H)$, and will contain $\mu h.\text{push}; H_\lambda; \text{pop}$ in the position corresponding to $\mu h.H_\lambda$ in $(\text{hextract } C \ H)$. Otherwise, observe that $f(\text{hextract}(H', C'))$ is homomorphic to $(\text{hextract } C \ H)$, establishing (2). Since $(\text{hextract } C \ H)$ contains no occurrences of push or pop, therefore $f(\text{hextract}(H', C'))$ contains none aside from those at the limits of μ -scope, and is therefore stackifiable, establishing (1), while (3) follows by definition of *stackify* and *stackify'*. \square

Finally, we can demonstrate correctness of *stackify*, by demonstrating a trace approximation result for λ_{trace}^S via on the strength of the preceding results and trace approximation in λ_{trace} .

Theorem 8.1 (Trace Approximation stackify)

If $\emptyset, H \vdash_{\beta} e : \tau/C$ is satisfiable and also $\text{nil}, e \rightarrow^* S, e'$, then $\hat{S} \in \llbracket \text{stackify}(\text{hextract } C \ H) \rrbracket$.

Proof

First note that e must be top-level; otherwise, it would not be typable as in the premise. Thus by Lemma 8.8, we have $\epsilon, \llbracket e \rrbracket \rightarrow^* \zeta, \llbracket e' \rrbracket$ with $\text{sctx}(\zeta) = \hat{S}$. By Lemma 8.12 there exists a satisfiable judgement $\emptyset, H' \vdash_{\beta'} \llbracket e \rrbracket : \tau/C'$, with $\hat{\zeta} \in \llbracket \text{hextract}(H', C') \rrbracket$ by Corollary 6.3. Therefore, by Lemma 8.9 and Lemma 8.11 and the equality $\text{sctx}(\zeta) = \hat{S}$, we may assert $\hat{S} \in \llbracket \text{stackify}'(\text{hextract}(H', C')) \rrbracket$, so the result follows by Lemma 8.12. \square

9 Verification of trace effects

We have described the syntax and semantics of λ_{trace} , which include a trace component in configurations. We have also described two type and effect systems that conservatively approximate run-time traces, and shown that type and effects can be automatically inferred. However, we have been abstract so far with respect to the definition of checks and effect verification, basing safety of computation and validity of typing on a yet-to-be-defined notion of check and effect validity. To fill in these details, we define in this Section:

- A logic for run-time checks, including a syntax for expressing checks as predicates in the logic, along with a notion of validity for these checks that can be automatically verified in the run-time system.

- A means of verifying validity of trace effects, as defined in Definition 3.4, where check events that are predicted by the trace effect analysis are automatically shown to either succeed or fail in the relevant context.

The latter point is necessary to address because, even though validity of trace effects has been defined, the notion is logical but not algorithmic; in particular, $\llbracket H \rrbracket$ may be an infinite set. We accomplish automated verification using a temporal logic and model-checking techniques, allowing us to reuse existing algorithms and results for trace effect verification. We use a linear-time logic because the run-time validity of a predicate is based on one linear trace η only.

9.1 Trace effects are BPAs

We interpret trace effects as labelled transition systems (Definition 3.3), with a form that is appropriate for our application. This form is close to that of Basic Process Algebra (BPA) terms; in fact, the latter can faithfully represent trace types. Such a representation allows us to apply known model-checking algorithms to our analysis, since several exist for verifying properties of BPAs Steffen & Burkart (1992).

Definition 9.1

BPA expressions are defined as follows:

$$p ::= a \mid \epsilon \mid p \cdot p \mid p + p \mid X \quad \text{BPA expressions}$$

Intuitively, each a denotes a transition action, ϵ is the empty process, \cdot denotes sequencing, and $+$ denotes nondeterministic choice. Each X in a BPA process is defined via a set of declarations Δ of the form $X \triangleq p$. We let P range over sets of BPA expressions. The definition of BPA processes is completed via the following operational semantics, where $\epsilon \cdot p$ is considered equivalent to p :

$$\frac{p \xrightarrow{a} p'}{p + q \xrightarrow{a} p'} \quad \frac{q \xrightarrow{a} q'}{p + q \xrightarrow{a} q'} \quad \frac{p \xrightarrow{a} p'}{p \cdot q \xrightarrow{a} p' \cdot q} \quad \frac{p \xrightarrow{a} p' \quad X \triangleq p \in \Delta}{X \xrightarrow{a} p'} \quad a \xrightarrow{a} \epsilon$$

Now, we define a straightforward transformation from trace effects to BPA processes. We define actions a in the BPA encoding to be events $w(c)$. Note that since BPA declarations may include only guarded expressions, the translation is meaningful just in case the transformation of a μ -bound trace effect yields a guarded declaration; we therefore assume a trivial pre-processing of traces to ensure this property, that, e.g., could insert ϵ before every trace effect variable h :

Definition 9.2

The translation from closed, ϵ -free trace effects H to BPA expression, declaration pairs p, Δ is inductively defined as follows. The translation is parameterized by an environment Ψ that maps trace effect variables h to their BPA variable representations

X in the encoding:

$$\begin{aligned}
\text{BPA}(\epsilon, \Psi) &= \epsilon, \emptyset \\
\text{BPA}(ev(c), \Psi) &= ev(c), \emptyset \\
\text{BPA}(H_1; H_2, \Psi) &= \text{let } p_1, \Delta_1 = \text{BPA}(H_1, \Psi) \text{ in} \\
&\quad \text{let } p_2, \Delta_2 = \text{BPA}(H_2, \Psi) \text{ in} \\
&\quad p_1 \cdot p_2, \Delta_1 \cup \Delta_2 \\
\text{BPA}(H_1 | H_2, \Psi) &= \text{let } p_1, \Delta_1 = \text{BPA}(H_1, \Psi) \text{ in} \\
&\quad \text{let } p_2, \Delta_2 = \text{BPA}(H_2, \Psi) \text{ in} \\
&\quad p_1 + p_2, \Delta_1 \cup \Delta_2 \\
\text{BPA}(\mu h.H, \Psi) &= \text{let } p, \Delta = \text{BPA}(H, (\Psi; h : X)) \text{ in} \\
&\quad X, \Delta \cup \{X \triangleq p\} \quad X \text{ fresh} \\
\text{BPA}(h, \Psi) &= \Psi(h), \emptyset
\end{aligned}$$

For brevity, we write $\text{BPA}(H)$ for $\text{BPA}(H, \emptyset)$

The correctness of the translation falls out as a simple corollary:

Corollary 9.1

$$\llbracket H \rrbracket = \llbracket \text{BPA}(H) \rrbracket.$$

9.2 Verified checks in the linear μ -calculus

While a variety of model-checking logics are available, we use the μ -calculus Kozen (1983) because it is powerful and is syntactically close to trace effects H . Furthermore, efficient techniques for the automated verification of μ -calculus formulas on BPA processes have been developed (Esparza 1994; Burkart *et al.* 2001). We use the *linear* variant of the μ -calculus Esparza (1994) because at run-time only one linear trace of events is known, and so trace effects H and run-time traces η can be compared only in a linear-time logic.

Definition 9.3

The syntax of the linear μ -calculus is

$$\phi ::= x \mid \mathbf{true} \mid \mathbf{false} \mid (a)\phi \mid \mu x.\phi \mid \nu x.\phi \mid \neg\phi \mid \phi \vee \phi \mid \phi \wedge \phi$$

Here, a ranges over arbitrary transition labels; in particular, we can let a range over events $ev(c)$. The semantics of the linear μ -calculus is defined in Figure 18. This semantics is defined over potentially infinite traces $\theta^\infty \in \Theta^\infty$ that, unlike sets of traces θ , may not be prefix-closed. V denotes a mapping from μ -calculus variables to sets of infinite traces, \emptyset denotes the empty mapping, and $\llbracket \phi \rrbracket$ is shorthand for $\llbracket \phi \rrbracket_\emptyset$. Several formulae are not given in this figure because they can be defined in terms of the others:

$$\phi_1 \vee \phi_2 \iff \neg(\neg\phi_1 \wedge \neg\phi_2) \quad \mathbf{false} \iff \neg\mathbf{true} \quad \mu x.\phi \iff \neg(\nu x.\neg\phi)$$

Since our trace effect semantics is prefix-closed, we will explicitly prefix-close $\llbracket \phi \rrbracket$ so the two sets are compatible.

$$\begin{aligned}
 \llbracket \text{true} \rrbracket_V &= \Theta^\infty \\
 \llbracket x \rrbracket_V &= V(x) \\
 \llbracket \neg\phi \rrbracket_V &= \Theta^\infty - \llbracket \phi \rrbracket_V \\
 \llbracket \phi_1 \wedge \phi_2 \rrbracket_V &= \llbracket \phi_1 \rrbracket_V \cap \llbracket \phi_2 \rrbracket_V \\
 \llbracket (a)\phi \rrbracket_V &= \{ \theta^\infty \in \Theta^\infty \mid \theta^\infty = a; \theta_1^\infty \text{ and } \theta_1^\infty \in \llbracket \phi \rrbracket_V \} \\
 \llbracket \nu x.\phi \rrbracket_V &= \bigcup \{ W \subseteq \Theta^\infty \mid W \subseteq \llbracket \phi \rrbracket_{V[x \rightarrow W]} \}
 \end{aligned}$$

Fig. 18. Semantics of the linear-time μ -calculus.

Definition 9.4

The prefix closure $\Theta^\infty \downarrow$ of a set of infinite traces Θ^∞ is:

$$\begin{aligned}
 &\{ \theta \mid \theta \text{ is a prefix of some } \theta^\infty \in \Theta^\infty \} \\
 &\cup \\
 &\{ \theta \downarrow \mid \theta \text{ is finite and } \theta \in \Theta^\infty \}
 \end{aligned}$$

Validity of a formulae with respect to a trace effect is defined via set containment of the trace interpretations of each:

Definition 9.5

The formula ϕ is valid for H , written $H \Vdash \phi$, iff $\llbracket H \rrbracket \subseteq \llbracket \phi \rrbracket \downarrow$.

This relation is decidable by known model-checking results and the above equivalence of BPA processes and trace effects:

Lemma 9.1

The relation is $H \Vdash \phi$ is decidable.

Proof

Model-checking linear μ -calculus formulae over BPAs are shown decidable in Esparza (1994). The key in this construction is that the valid traces of each linear μ -calculus formula ϕ can be characterized by a finite automaton and Büchi automaton that accept the finite and infinite traces of $\neg\phi$, respectively. Histories can be translated to equivalent BPAs by Corollary 9.1. \square

9.3 Relating trace effect and trace runtime properties

We now instantiate the logic of trace checks in λ_{trace} with linear μ -calculus formulae ϕ . In Lemma 9.1, the relation $H \Vdash \phi$ was shown to be decidable, so this will make a natural foundation for trace effect verification.

One important requirement of this logic is that formulae must have truth values for a given trace effect H , and for a trace runtime η . The meaning of a formula ϕ under a run-time trace η is taken by verifying $\hat{\eta} \in \llbracket \phi \rrbracket \downarrow$. We will define the trace check interpretation function Π of Figure 2 in terms of this relation (Definition 9.6).

In Corollary 4.1, trace effects were shown to approximate dynamic histories—in particular, if $\epsilon, e \rightarrow^* \eta, e'$ and $H, \Gamma \vdash e : \tau$ is derivable, then $\hat{\eta} \in \llbracket H \rrbracket$. The key result linking the static and dynamic histories is the following:

Lemma 9.2

If $\hat{\eta} \in \llbracket H \rrbracket$ and $H \Vdash \phi$, then $\hat{\eta} \in \llbracket \phi \rrbracket \downarrow$.

Proof

Immediate from Definitions 9.5 and 3.3. \square

Corollary 4.1 ensures that by the above Lemma, verifying a trace effect entails verifying all trace checks that may occur at run-time, meaning a combination of type inference and automated verification yields a sound static analysis for λ_{trace} , as we state formally in Theorem 9.1.

9.4 Soundness of verification

We have almost completed the definition of our logical framework for trace checks and trace effect verification. However, a few small points remain.

First, in λ_{trace} , we are able to parameterize checks with constants with the syntax $\phi(c)$. To implement this, we specify a distinguished variable χ that may occur free in trace checks, which is assumed to be instantiated with a trace check parameter during verification.

Second, checks ϕ should express expected trace patterns that may occur up to the point of the check. This is the same as requiring that if an event $ev_{\phi}(c)$ occurs in a trace (resp. is predicted statically), the subtrace η that immediately precedes it (resp. any trace that may precede it as predicted by typing) must exhibit the pattern specified by ϕ . However, there is an infinite regress lurking here: a property ϕ that mentions an event $ev_{\phi}(c)$ suggests circularity in the syntax of ϕ . Thus, we introduce a distinguished label *Now*, that in any formula ϕ represents the relevant checkpoint of ϕ . This label is interpreted appropriately during verification.

The mechanics of χ and *Now* are fully specified in Definition 9.6 below.

We now stand ready to instantiate our logic and verification framework. In the dynamic system, this is accomplished by defining the language of trace checks, and by defining the implementation of Π , our previously abstract representation of trace check verification:

Definition 9.6 (Definition of Π)

The framework of Section 2 is instantiated to let ϕ range over linear μ -calculus formulae, where labels a in ϕ are events $ev(c)$. Furthermore, we distinguish one event *Now* to represent checkpoints within the check specification (hence avoiding circularity), and letting $\chi \in \mathcal{V}_s$ be a distinguished variable for parameterising constants c in ϕ , we define Π as follows:

$$\Pi(\phi(c), \theta) \iff \theta \in \llbracket \phi[c/\chi][ev_{\phi}(c)/\text{Now}] \rrbracket \downarrow$$

Now, we specify what it means for a trace effect to be verified; intuitively, it means that if a trace effect H predicts the occurrence of a check event $ev_{\phi}(c)$, then H semantically entails ϕ instantiated with c . Formally:

Definition 9.7 (Trace Effect Verification)

A trace effect H is *verifiable* iff for all subterms $ev_{\phi}(c)$ of H it is the case that

$$H \Vdash \phi[c/\chi][ev_{\phi}(c)/\text{Now}]$$

Here is a simple example that demonstrates the system in action:

Example 9.1

Assume we want to define a policy whereby if a particular check $\phi(c)$ is encountered at run-time, the only allowable previous events are $ev_1(c)$ or $ev_2(c)$, and $ev_2(c)$ must be the most recent event prior to the check. For brevity, let

$$\begin{aligned}(ev_{1|2}(\chi))\phi &\triangleq (ev_1(\chi))\phi \vee (ev_2(\chi))\phi \\ (ev_{1|2}(\chi)^*)\phi &\triangleq \mu x.(ev_{1|2}(\chi))x \vee \phi\end{aligned}$$

The above policy can then be expressed in our scheme with the following μ -calculus formula ϕ :

$$\phi \triangleq (ev_{1|2}(\chi)^*)(ev_2(\chi))(\text{Now})\mathbf{true}$$

Now, assuming $\emptyset, \epsilon \vdash e' : \text{bool}$, let

$$e \triangleq \text{if } e' \text{ then } ev_1(c); ev_2(c) \text{ else } ev_2(c)$$

Thus, depending on the valuation of e' , either:

$$\epsilon, e; \phi(c) \rightarrow^* ev_1(c); ev_2(c), \phi(c)$$

or

$$\epsilon, e; \phi(c) \rightarrow^* ev_2(c), \phi(c)$$

and we see that the check in the program $(e; \phi(c))$ will always succeed at run-time, since these relations hold:

$$\begin{aligned}ev_2(c)ev_\phi(c) &\in \llbracket \phi[c/\chi][ev_\phi(c)/\text{Now}] \rrbracket \downarrow \\ ev_1(c)ev_2(c)ev_\phi(c) &\in \llbracket \phi[c/\chi][ev_\phi(c)/\text{Now}] \rrbracket \downarrow\end{aligned}$$

We also have that the following judgement is derivable:

$$\emptyset, (ev_2(c)|ev_1(c); ev_2(c)); ev_\phi(c) \vdash e; \phi(c) : \text{unit}$$

and this entailment holds:

$$(ev_2(c)|(ev_1(c); ev_2(c))); ev_\phi(c) \Vdash \phi[c/\chi][ev_\phi(c)/\text{Now}]$$

Therefore, the judgement is valid, predicting the run-time safety of the program $(e; \phi(c))$.

The preceding construction completes the definition of the framework. Lemma 9.2 and definition of Π together yield the desired formal property for trace effect verification as a Corollary:

Corollary 9.2 (Verification Soundness)

If H is verified, then H is valid.

Proof

Immediate by Definition 3.4, Definition 9.6, Definition 9.7, and Lemma 9.2. \square

This result can then be composed with Lemma 7.22 and Theorem 6.2 to yield a type safety property for the complete automated analysis.

Theorem 9.1

Suppose $\emptyset, H \vdash_{\bar{\beta}} e : \tau/C$ is derivable, $close(C)$ is consistent, and $hextract C H$ is verifiable. Then e does not go wrong.

9.5 Incompleteness of verification

Completeness of verification is a matter different from soundness, and the reader may notice that our verification technique is approximate. In particular, verification of a trace effect H (Definition 9.7) with respect to a predicate ϕ always ensures that $\theta \in \llbracket H \rrbracket$ implies $\theta \in \llbracket \phi \rrbracket$, regardless of whether ϕ occurs in θ , whereas validity (Definition 3.4) only requires that (roughly speaking) $\theta e_{\phi} \in \llbracket H \rrbracket$ implies $\theta e_{\phi} \in \llbracket \phi \rrbracket$, i.e., validity with respect to a given ϕ is concerned only with traces that end in e_{ϕ} .

Example 9.2

With ϕ defined as in Example 9.1, the effect $(e_{1(c)}; \phi(c))$ is both valid and verifiable, as is the effect $e_{3(c)}$ since it contains no checks. However, the join of the two is not verifiable, since $e_{3(c)}$ does not satisfy $\phi(c)$, even though it is valid since $e_{3(c)}$ precedes no checks and hence is irrelevant to validity:

$$e_{3(c)}|(e_{1(c)}; \phi(c)) \text{ is valid but not verifiable}$$

While this is sound, more precise and complete verification can be obtained with some modifications. We now show how to enhance verification with a technique that will remove from consideration those traces containing no occurrence of a given formula.

We begin by defining a convenient μ -calculus formula abbreviation, using familiar notation. This abbreviation will be reused in Section 10:

Definition 9.8

Since verification of predicates ϕ is always with respect to a particular trace effect H , and any event $e(c)$ occurring in $\llbracket H \rrbracket$ must occur in H for the class of trace effects we consider for verification (see Section 9.1), we can easily obtain the set of possible events in a trace set $\llbracket H \rrbracket$ as the set of events in H . Letting $e_{1(c_1)}, \dots, e_{n(c_n)}$ be the events for given H , we define the following abbreviations:

$$\begin{aligned} (\cdot)\phi &\triangleq (e_{1(c_1)})\phi \vee \dots \vee (e_{n(c_n)})\phi \\ (\cdot*)\phi &\triangleq \mu x.(\cdot)x \vee \phi \end{aligned}$$

A refined, more complete verification can then be defined as follows.

Definition 9.9 (Refined Verification)

Let the *occurs* predicate be defined as follows: a trace θ is in the interpretation of $\text{occurs}(e(c))$ just in case $e(c)$ occurs in θ :

$$\text{occurs}(e(c)) \triangleq (\cdot*)(e(c))(\cdot*)\mathbf{true}$$

A trace effect H is then *verified* iff for all subterms $e_{\phi}(c)$ of H , letting

$$\phi' = \phi[c/\chi][e_{\phi}(c)/\text{Now}]$$

it is the case that

$$H \Vdash \phi' \vee \neg \text{occurs}(\phi')$$

The soundness results in the previous section can then be replayed in the presence of this refined verification.

10 Applications to language-based security

In this section we show that our program logic is sufficiently expressive to be useful in practice applications including security. In particular, we show how both stack inspection and history-based security paradigms can be statically enforced in our system. Singleton types allow a precise typing of resource parameters in the model.

In our examples we will be interested in unparameterized events and checks; in such cases we will write ev and ϕ for $ev(c_{\text{dummy}})$ and $ev_{\phi}(c_{\text{dummy}})$ respectively, where c_{dummy} is a distinguished dummy constant. Also, we will abbreviate events ev_i by their subscripts i , and the notation:

$$(\vee \{ev_1(c_1), \dots, ev_n(c_n)\})\phi \triangleq ev_1(c_1)\phi \vee \dots \vee ev_n(c_n)\phi$$

will be convenient.

10.1 Stack inspection with parameterized privileges

Java stack inspection Wallach & Felten (1998); Skalka & Smith (2000); Pottier *et al.* (2001) is a language-based security mechanism that uses an underlying security model of principals and resources—all code is annotated with a principal identifier p , and a local ACL policy \mathcal{A} mapping principals p to resources $r(c)$ for which they are authorized is taken as given. An event ev_p is issued whenever a codebase annotated with p is entered. One additional feature of Java is that any principal may also explicitly *enable* a resource for which it is authorized. When a function is activated, its associated principal identifier is pushed on the stack, along with any resource enablings that occur in its body. Stack inspection for a particular resource $r(c)$ then checks the stack for an enabling of $r(c)$, searching frames from most to least recent, and failing if a principal unauthorized for $r(c)$ is encountered before an enabling of $r(c)$, or if no such enabling is encountered.

Stack inspection can be modeled in the stack-based variant of our programming logic defined in Section 8. Rather than defining the general encoding, we develop one particular example that illustrates all issues. Consider the following function checkit:

$$\text{checkit} \triangleq \lambda x.p:\text{system}; \phi_{\text{inspect},r:\text{filew}}(x)$$

Every function upon execution first issues an owner (principal) event, in this case $p:\text{system}$ indicating “system” is the principal p that owns checkit. The function takes a parameter x (a file name) and inspects the stack for the “filew” resource with parameter x , via embedded μ -calculus assertion $\phi_{\text{inspect},r:\text{filew}}(x)$. This assertion is defined below; it enforces the fact that all functions on the call stack back to the nearest enable must be owned by principals p that according to ACL \mathcal{A} are authorized for the $r:\text{filew}(x)$ resource.

To model resource enabling, we use a check event $\phi_{\text{enable},r}(x)$; the use of a check event allows to simultaneously mark a temporary enabling, and to check that enabling is performed by an authorized principal. We illustrate the use of explicit enabling via an example “wrapper” function enableit, owned by the “accountant”

$$\begin{aligned}
(p_{\neg r(c)})\phi &\triangleq (\bigvee \{p|r(c) \notin \mathcal{A}(p)\})\phi \\
(\bar{p})\phi &\triangleq (\bigvee (\{e_{v_1}(c_1), \dots, e_{v_n}(c_n)\} \setminus \text{dom}(\mathcal{A})))\phi \\
(\neg e(c))\phi &\triangleq (\bigvee (\{e_{v_1}(c_1), \dots, e_{v_n}(c_n)\} \setminus \{e(c)\}))\phi \\
(a^*)\phi &\triangleq \mu x.(a)x \vee \phi \quad \text{for } a \in \{\bar{p}, \neg e(c)\} \\
\phi_{\text{enable},r}(c) &\triangleq \neg((\cdot)(p_{\neg r(c)})(\bar{p}^*)(\text{Now})\text{true}) \\
\phi_{\text{amplify},r}(c) &\triangleq \phi_{\text{enable},r}(c) \\
\phi_{\text{inspect},r}(c) &\triangleq \neg((\neg\phi_{\text{enable},r}(c)^*)(\text{Now})\text{true}) \wedge \\
&\quad \neg((\cdot)(p_{\neg r(c)})(\neg\phi_{\text{enable},r}(c)^*)(\text{Now})\text{true}) \\
\phi_{\text{demand1},r}(c) &\triangleq \neg((\cdot)(p_{\neg r(c)})(\cdot)(\text{Now})\text{true}) \\
\phi_{\text{demand2},r}(c) &\triangleq \neg((\neg\phi_{\text{amplify},r}(c)^*)(\text{Now})\text{true}) \wedge \\
&\quad \neg((\cdot)(p_{\neg r(c)})(\neg\phi_{\text{amplify},r}(c)^*)(\text{Now})\text{true})
\end{aligned}$$

Fig. 19. Definitions of $\phi_{\text{demand},r}$ and $\phi_{\text{inspect},r}$.

principal $p:\text{acct}$, that takes a function f and a constant x , and enables $r:\text{filew}(x)$ for the application of f to x :

$$\text{enableit} \triangleq \lambda f.p:\text{acct}; (\lambda x.p:\text{acct}; \phi_{\text{enable},r:\text{filew}(x)}; \text{let } y = f(x) \text{ in } y)$$

The definition of $\phi_{\text{inspect},r}(c)$, for fixed $r(c)$, is generalized over parameterized resources $r(c)$. For trace effect H containing only the events $e_{v_1}(c_1), \dots, e_{v_n}(c_n)$, and parameterized resource $r(c)$, Figure 19 gives the definition of $\phi_{\text{inspect},r}(c)$. The check ensures that it is neither the case that the check occurs without previous enablings of $r(c)$, nor that a principal unauthorized for the privilege is interposed between the check and the most recent enabling; this is logically equivalent to stack inspection.

Returning to our previous example expressions `checkit` and `enableit`, the following most general types are inferred in our system:

$$\begin{aligned}
\text{checkit} &: \forall \alpha. \{\alpha\} \xrightarrow{p:\text{system}; \phi_{\text{inspect},r:\text{filew}(x)}} \text{unit} \\
\text{enableit} &: \forall \alpha h t. (\{\alpha\} \xrightarrow{h} t) \xrightarrow{p:\text{acct}} \{\alpha\} \xrightarrow{p:\text{acct}; \phi_{\text{enable},r:\text{filew}(x)}; h} t
\end{aligned}$$

The stackification of the *hextracted* effect of the application:

$$\text{enableit checkit } (/ \text{accts/ledger.txt})$$

will yield the effect $p:\text{acct}|H$, where

$$H = p:\text{acct}; \phi_{\text{enable},r:\text{filew}(/\text{accts/ledger.txt})}; p:\text{system}; \phi_{\text{inspect},r:\text{filew}(/\text{accts/ledger.txt})}$$

This reflects that the call and return of the application (`enableit checkit`) is assigned the effect $p:\text{acct}$ while the subsequent application to `/accts/ledger.txt` is assigned effect H . Assuming that both $p:\text{system}$ and $p:\text{acct}$ are authorized for $r:\text{filew}(/\text{accts/ledger.txt})$ in \mathcal{A} , verification will clearly succeed on this expression. On the other hand, stackification of the application `checkit(/accts/ledger.txt)` will generate the following trace effect:

$$p:\text{system}; \phi_{\text{inspect},r:\text{filew}(/\text{accts/ledger.txt})}$$

for which verification will fail: there is no required $\phi_{\text{enable},r:\text{filew}}(/accts/ledger.txt)$ on the stack.

10.2 History-based access control

History-based access control is a generalization of Java's notion of stack inspection that takes into account all past events, not just those on the stack Abadi & Fournet (2003). Our language is well suited to the static typechecking of such security policies. In the basic history model of Abadi and Fournet (2003), some initial *current rights* are given, and with every new activation the static rights of that activation are automatically intersected with the current rights to generate the new current rights. Unlike stack inspection, removal of the activation does not return the current rights to its state prior to the activation.

We assume the same underlying model of principals as above. A demand of a resource r with parameter c , denoted $\phi_{\text{demand},r}(c)$, requires that all invoked functions possess the right for that resource. This general check may be expressed in our logic as in Figure 19, where we assume given for verification some trace effect H containing events $ev_1(c_1), \dots, ev_n(c_n)$. For example, validity of the following type judgement requires $r(c) \in \mathcal{A}(p_1) \cap \mathcal{A}(p_2)$:

$$\Gamma, p_1; p_2; \phi_{\text{demand},r}(c) \vdash p_1; (\lambda x.p_2; \phi_{\text{demand},r}(x)) c : \text{unit}$$

The model in Abadi and Fournet (2003) also allows for a combination of stack- and history-based properties, by allowing the *amplification* of a right on the stack. When a right is amplified, it remains active after function return regardless of the current rights generated by the invocation, provided that amplification is performed by an authorized principal. Such assertions can be expressed in our framework using a combination of stack- and history-based assertions.

Taking the same strategy as for enabling in the stack inspection model, we can define amplification as a check event $\phi_{\text{amplify},r}$; in fact, this can be defined equivalent to $\phi_{\text{enable},r}$, and enforced by stackifying effects before verifying. In other words, the semantics of amplification are identical to the semantics of enabling. Then, making the simplifying assumption that by convention all resources are initially amplified, the demand predicate can be redefined in the presence of amplification as $\phi_{\text{demand},r}$ in Figure 19. Observe that this predicate is syntactically identical to $\phi_{\text{inspect},r}$, except that effects should *not* be stackified before verifying $\phi_{\text{demand},r}$. The check $\phi_{\text{demand},r}(c)$ ensures that an amplification of $r(c)$ occurs prior to it, without any interposed principal that is unauthorised for $r(c)$.

Given that history-based access control requires an interplay of stack- and history-based checks, we argue that a contribution of our model is its ability to deal with these variations in a uniform manner, within one rigorous formal framework.

10.3 Discussion

The examples applications of our system developed in this section illustrate its flexibility and power. The distinction of traces and predicates on traces allows a general class of security properties to be defined. Trace effects themselves are

amenable to interpretation as abstractions of either history traces or stack traces. The examples, stack inspection and history-based access control, show particular applications in stack trace and history trace settings respectively, but any property expressible in a temporal logic could be used to specify desired history or stack trace behavior. This follows the JDK security model Gong *et al.* (1997) in which stack inspection is the default predicate on stack traces, but other predicates can be programmed and used instead. In practice this offers a great deal of flexibility in the kinds of security properties that can be enforced and the ability to build security “libraries” for various domains, such as the Java sandboxing model for mobile code. Again, we should emphasize our system places stack- and history-based mechanisms on a single unified foundation.

A limitation of our system is that event and check parameters must be syntactic constants. In Java, for example, privilege parameters can be dynamically generated; e.g., “file read” privileges may be parameterized by a new filehandle. Other researchers have addressed this issue in the similar static program verification contexts, in particular Igarishi and Kobayashi (2002), and Hamlen *et al.* (2006) who consider reference monitors in the .NET security model. In Higuchi and Ohori (2007) the authors deal specifically with stack inspection and define privilege parameters to be either “may” sets of literal values, or a distinguished “unknown” value. Rather than generate privilege constraints that must all be statically satisfied, run time checks are inserted into the program at checks that *may* fail. A similar soft-typing approach is not precluded by our system, and extending it to treat dynamically generated constants in this manner is an interesting direction for future work.

11 Conclusion

In this section, we conclude with a discussion of related work and a summary of the paper.

11.1 Related work

Previous work relevant to the application of trace-based security models has been noted in Section 1. Here we discuss related theories and systems designed to enforce trace-based properties of program execution.

11.1.1 Compile-time versus run-time verification

Perhaps the principal division between previous approaches to the enforcement of trace-based program properties is between those systems that detect errors at run-time (Nierstrasz 1993; Edjlali *et al.* 1998; Rossie 1998; Schneider 2000; Bauer *et al.* 2002a, 2002b; Abadi & Fournet 2003) versus those that detect errors at compile-time (Schmidt 1998; Ball & Rajamani 2000; Besson *et al.* 2001; Chen & Wagner 2002). Run-time approaches are more accurate since a compile-time analysis must conservatively approximate what events could occur; the compile-time analysis will also reject some safe programs, due to the need to be conservative. On the other

hand, liveness properties (“all SSL sockets are eventually closed”) can be verified at compile-time but not at run-time.

11.1.2 Run-time verification

Some early work in the area was carried out by Nierstrasz and other object-oriented language designers (Nierstrasz 1993; Rossie 1998). Their goal was to specify only the legal patterns of messages that could be passed to a particular object. The SSL example in Section 1 could be cast in that view by taking the events to be messages to some SSL control object. Regular expressions and finite automata were used for the specification logic in this work. Security automata (Schneider 2000; Bauer *et al.* 2002a, 2002b) also use finite automata for specification. Security automata are targeted at specification and verification of security policies. All of these run-time approaches assume that there is a run-time monitoring process that will either abort or raise an exception if a bad trace is realized. Some automata transitions can be optimized away using simple static program analyses Erlingsson & Schneider (2000), but the analysis is still mostly dynamic: errors will not be realized until run-time.

Several groups (Edjlali *et al.* 1998; Abadi & Fournet 2003) have proposed using event traces for access control. A given access is permitted only under a restricted set of previous circumstances, recorded in the history. These projects propose no particular specification language—the policies are directly coded. The stack inspection algorithm used in the .Net CLR and Java Security Architecture is another specific access control policy based on execution history (Gong *et al.* 1997).

11.1.3 Compile-time verification

The MOPS system (Chen & Wagner 2002) compiles C programs to Push-Down Automata (PDA) reflecting the program control flow, where transitions are program transitions and the automaton stack abstracts the program call stack. Similarly, the ESP system (Das *et al.* 2002) checks temporal properties of programs (encoded as finite state machines) by simulating an approximation of property transitions that will occur at run-time. Jensen *et al.* (1999) and Besson *et al.* (2001) assume that some (undefined) algorithm has already converted a program to a control flow graph, expressed as a form of PDA. The Vault system (DeLine & Fähndrich 2001) enforces resource usage constraints with *type guards*, temporal specifications of when an operator may be applied—similar in spirit to the trace checks presented here—and verification is performed via type checking that ensures well-typed programs satisfy their guard specifications.

These aforementioned abstractions work well for procedural programs, but are not powerful enough to fully address advanced language features such as higher order functions. Our approach, based on type and effect theory (Talpin & Jouvelot 1992; Amtoft *et al.* 1999), does allow abstract interpretation of higher order programs. Trace effects yielded by the analysis provide a conservative approximation of trace behavior via a Labelled Transition System (LTS) interpretation. Program assertions can be expressed as temporal logical formulae that can be automatically verified by model-checking (Steffen & Burkart 1992).

Systems have also been developed for statically verifying correctness of security automata using dependent types (Walker 2000), and in a more general form as refinement types Mandelbaum *et al.* (2003). These systems do not extract any abstract interpretations, and so are in a somewhat different category than the others. They also lack type inference and do not use a logic that can be automatically verified, so user intervention is needed.

11.1.4 Trace specification logics and model checking

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

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

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

11.1.5 Local and global reasoning

Logical assertions can be local, concerning a particular program point, or global, defining the whole behavior required. However, access control systems (Edjlali *et al.*

1998; Wallach & Felten 1998; Abadi & Fournet 2003) use local checks. Since we are interested in the static enforcement of access control mechanisms (see Section 10), the focus in this paper is on local, compile-time checkable assertions, although in principle the verification of global properties is also possible in our system. Related work has also modified our basic approach to enforce “policy framings” that support so-called local liveness properties (Bartoletti *et al.* 2005b).

11.2 Summary

This paper establishes a foundation for the enforcement of local well-formedness properties on program event traces. Our language model λ_{trace} extends a functional core with constructs to label program points, *aka* events, and with local assertions over program event traces up to the point of the assertion during evaluation. The language model is parameterized by a logic of event trace assertions and has been instantiated here with the linear μ -calculus.

We have developed a static type and effect theory for λ_{trace} programs, whereby computational traces arising at run-time can be conservatively approximated by trace effects, components of the type language that are interpreted as model-checkable LTSs. We have presented two variations on the system: the first uses a unified term syntax, while the second uses constraint subtyping. Both include parametric let-polymorphism over types and effects. The unified representation allows for weakening of trace effects through trace containment, whereas the constraint form allows for more expressive type subsumption. The constraint system is shown to be a conservative extension of both the unified term system and the underlying effect-free ML-style type system. Since trace effects can be model-checked, verification of trace effect approximations obtained by typing guarantees that all trace assertions will hold during execution of a given program. Thus, our type safety result guarantees that statically well-typed programs do not contain assertions that will fail dynamically.

Type reconstruction algorithms have been defined for both type systems. Type inference for constraint subtyping has been shown to be sound and complete, the latter result entailing a principal types property. Inference is decidable, despite the general undecidability of trace effect equality and containment, due to a restricted form of constraints generated by inference. Term forms of trace effects are extractable from effect constraints via the *hextract* algorithm, obtaining a compliant form for existing model checking techniques.

We have shown that trace effects can be transformed to reflect a stack discipline for function calls and returns via the *stackify* algorithm, allowing the expression of stack-based policies such as stack inspection and rights amplification for history-based access control. We have also shown how these policies can be defined in our system, as particular applications of a general framework for enforcing trace-based properties expressible in temporal logics. A primary strength of our approach is that the analysis is fully automated—no program abstraction or type annotations are presupposed—and each component of this automation has been rigorously justified. Our static analysis and verification techniques for trace-based properties in higher order languages provide well-founded, flexible, and general tools for enforcing program safety and security.

References

- Abadi, M. & Fournet, C. (2003) Access control based on execution history. In *Proceedings of the 10th Annual Network and Distributed System Security Symposium (NDSS'03)*.
- Amtoft, T., Nielson, F. & Nielson, H. R. (1999) *Type and Effect Systems: Behaviours for Concurrency*. London: Imperial College Press.
- Ball, T. & Rajamani, S. K. (2000) Bebop: A symbolic model checker for boolean programs. In proceedings of the 7th International SPIN Workshop on SPIN Model Checking and Software Verification. *Lecture Notes in Computer Science* vol. 1885, pp. 113–130.
- Bartoletti, M., Degano, P. & Ferrari, G. L. (2005a) Enforcing secure service composition. In *CSFW*. IEEE Computer Society, pp. 211–223.
- Bartoletti, M., Degano, P. & Ferrari, G. L. (2005b) History-based access control with local policies. In *FoSSaCS*, Vladimiro, S. (ed), *Lecture Notes in Computer Science*, vol. 3441. Springer, pp. 316–332
- Bauer, L., Ligatti, J. & Walker, D. (2002a) More enforceable security policies. In *Proceedings of the Foundations of Computer Security Workshop*.
- Bauer, L., Ligatti, J. & Walker, D. (2002b) Types and effects for non-interfering program monitors. In *Proceedings of the International Symposium on Software Security*.
- Besson, F., Jensen, T., Métayer, D. Le & Thorn, T. (2001) Model checking security properties of control flow graphs. *J. Comput. Secur.* **9**, 217–250.
- Besson, F., de Grenier de Latour, T. & Jensen, T. (2002) Secure calling contexts for stack inspection. In *PPDP '02: Proceedings of the 4th ACM SIGPLAN International Conference on Principles and Practice of Declarative Programming*, pp. 76–87.
- Burkart, O., Caucal, D., Moller, F. & Steffen, B. (2001) Verification on infinite structures. In *Handbook on Process Algebra*, Bergstra, J., Pons, A., & Smolka, S. (eds), Elsevier, pp. 545–623.
- Chen, H. & Wagner, D. (2002) MOPS: An infrastructure for examining security properties of software. In *CCS '02: Proceedings of the 9th ACM Conference on Computer and Communications Security*, pp. 235–244.
- Colcombet, T. & Fradet, P. (2000) Enforcing trace properties by program transformation. In *POPL '00: Proceedings of the 27th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pp. 54–66.
- Das, M., Lerner, S. & Seigle, M. (2002) ESP: path-sensitive program verification in polynomial time. In *PLDI '02: Proceedings of the ACM SIGPLAN 2002 Conference on Programming Language Design and Implementation*, pp. 57–68.
- DeLine, R. & Fähndrich, M. (2001) Enforcing high-level protocols in low-level software. In *PLDI '01: Proceedings of the ACM SIGPLAN 2001 Conference on Programming Language Design and Implementation*, pp. 59–69.
- Edjlali, G., Acharya, A. & Chaudhary, V. (1998) History-based access control for mobile code. In *ACM Conference on Computer and Communications Security*, pp. 38–48.
- Eifrig, J., Smith, S. & Trifonov, V. (1995) Type inference for recursively constrained types and its application to OOP. In *Proceedings of the 1995 Mathematical Foundations of Programming Semantics Conference*. *Electronic Notes in Computer Science*, vol. 1. Elsevier.
- Erlingsson, Ú. & Schneider, F. B. (2000) SASI enforcement of security policies: A retrospective. In *NSPW '99: Proceedings of the 1999 Workshop on New Security Paradigms*, pp. 87–95.
- Esparza, J. (1994) On the decidability of model checking for several mu-calculi and Petri nets. In *Proceeding of CAAP '94*. *Lecture Notes in Computer Science*, vol. 787.
- Esparza, J., Kucera, A. & Schwoon, S. (2001) Model-checking LTL with regular valuations for pushdown systems. In *TACS: 4th International Conference on Theoretical Aspects of Computer Software*.

- Foster, J. S., Terauchi, T. & Aiken, A. (2002) Flow-sensitive type qualifiers. In *Proceedings of the 2002 ACM SIGPLAN Conference on Programming Language Design and Implementation*, pp. 1–12.
- Gong, L., Mueller, M., Prafullchandra, H. & Schemers, R. (1997) Going beyond the sandbox: An overview of the new security architecture in the Java Development Kit 1.2. In *USENIX Symposium on Internet Technologies and Systems*, pp. 103–112.
- Hamlen, K. W., Morrisett, G. & Schneider, F. B. (2006) Certified in-lined reference monitoring on .NET. In *Plas '06: Proceedings of the 2006 Workshop on Programming Languages and Analysis for Security*. New York: ACM Press, pp. 7–16.
- Higuchi, Tomoyuki & Ohori, Atsushi. (2007) A static type system for JVM access control. *ACM Trans. Program. Lang. Syst.* **29**(1), 4.
- Igarashi, A. & Kobayashi, N. (2002) Resource usage analysis. In *Conference Record of POPL'02: The 29th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pp. 331–342.
- Jensen, T., Métayer, D. Le & Thorn, T. (1999) Verification of control flow based security properties. In *Proceedings of the 1999 IEEE Symposium on Security and Privacy*.
- K. Marriott, P. J. Stuckey & Sulzmann, M. (2003) Resource usage verification. In *Proceedings of First Asian Programming Languages Symposium, APLAS 2003*.
- Kozen, D. (1983) Results on the propositional mu-calculus. *Theor. Comput. Sci.*, **27**, 333–354.
- Mandelbaum, Y., Walker, D. & Harper, R. (2003) An effective theory of type refinements. In *Proceedings of the Eighth ACM SIGPLAN International Conference on Functional Programming (ICFP'03)*.
- Nierstrasz, O. (1993) Regular types for active objects. In *Proceedings of the OOPSL '93 Conference on Object-Oriented Programming Systems, Languages, and Applications*, pp. 1–15.
- Palsberg, J. & O'Keefe, P. (1995) A type system equivalent to flow analysis. *ACM Trans. Program. Lang. Syst.*, **17**(4), 576–599.
- Pierce, B. C. (2002) *Types and Programming Languages*. The MIT Press, Chap. 22.
- Pottier, F., Skalka, C. & Smith, S. (2001) A systematic approach to static access control. In *Proceedings of the 10th European Symposium on Programming (ESOP'01)*, Sands, D. (ed), Lecture Notes in Computer Science, vol. 2028. Springer-Verlag, pp. 30–45.
- Rossie, J. G., Jr. (1998) Logical observable entities. In *OOPSLA '98: Proceedings of the 13th ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications*, pp. 154–165.
- Schmidt, D. A. (1998) Trace-based abstract interpretation of operational semantics. *Lisp Symbol. Comput.*, **10**(3), 237–271.
- Schneider, F. B. (2000) Enforceable security policies. *Inform. Syst. Secur.* **3**(1), 30–50.
- Skalka, C. (2005) Trace effects and object orientation. In *PPDP '05: Proceedings of the 7th ACM SIGPLAN International Conference on Principles and Practice of Declarative Programming*, pp. 139–150.
- Skalka, C. & Pottier, F. (2003) Syntactic type soundness for HM(X). *Electro. Notes Theor. Comput. Sci.*, **75**, pp. 61–74.
- Skalka, C. & Smith, S. (2000) Static enforcement of security with types. In *Proceedings of the the Fifth ACM SIGPLAN International Conference on Functional Programming (ICFP'00)*, pp. 34–45.
- Skalka, C. & Smith, S. (2004) History effects and verification. In *Asian Programming Languages Symposium*.
- Skalka, C., Smith, S. & Van Horn, D. (2005, January) A type and effect system for flexible abstract interpretation of Java. In *Proceedings of the ACM Workshop on Abstract*

- Interpretation of Object Oriented Languages*. Electronic Notes in Theoretical Computer Science.
- Steffen, B. & Burkart, O. (1992) Model checking for context-free processes. In *CONCUR'92, Stony Brook (NY)*. Lecture Notes in Computer Science (LNCS), vol. 630. Heidelberg, Germany: Springer-Verlag, pp. 123–137.
- Stone, C. (2000) *Singleton Types and Singleton Kinds*. Tech. Rept. CMU-CS-00-153. Carnegie Mellon University.
- Sulzmann, M. (2001) A general type inference framework for Hindley/Milner style systems. In *International Symposium on Functional and Logic Programming*. Lecture Notes in Computer Science, vol. 2024. Springer-Verlag, pp. 246–263.
- Talpin, J.-P. & Jouvelot, P. (1992) The type and effect discipline. In *Seventh Annual IEEE Symposium on Logic in Computer Science, Santa Cruz, California*. Los Alamitos, CA: IEEE Computer Society Press, pp. 162–173.
- Tarski, A. (1955) A lattice-theoretical fixpoint theorem and its applications. *Pac. J. Math.* 5(2), 285–309.
- Trifonov, V. & Smith, S. (1996) Subtyping constrained types. In *Proceedings of the Third International Static Analysis Symposium*, vol. 1145. Springer-Verlag, pp. 349–365.
- Van Horn, D. (2006a) *Algorithmic Trace Effect Analysis*. M.Phil. thesis, University of Vermont.
- Van Horn, D. (2006b) Trace effect analysis, OCaml implementation. Available at: <http://www.cs.uvm.edu/~dvanhorn/trace/>. Accessed June 2007.
- Walker, D. (2000) A type system for expressive security policies. In *Conference Record of POPL'00: The 27th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pp. 254–267.
- Wallach, D. S. & Felten, E. (1998) Understanding Java stack inspection. In *Proceedings of the 1998 IEEE Symposium on Security and Privacy*.