

# Implicit Self-Adjusting Computation for Purely Functional Programs

Yan Chen   Joshua Dunfield   Matthew A. Hammer   Umut A. Acar

Max Planck Institute for Software Systems  
{chenyan, joshua, hammer, umut}@mpi-sws.org

## Abstract

Computational problems that involve dynamic data, such as physics simulations and program development environments, have been an important subject of study in programming languages. Building on this work, recent advances in self-adjusting computation have developed techniques that enable programs to respond automatically and efficiently to dynamic changes in their inputs. Self-adjusting programs have been shown to be efficient for a reasonably broad range of problems but the approach still requires an explicit programming style, where the programmer must use specific monadic types and primitives to identify, create and operate on data that can change over time.

We describe techniques for automatically translating purely functional programs into self-adjusting programs. In this implicit approach, the programmer need only annotate the (top-level) input types of the programs to be translated. Type inference finds all other types, and a type-directed translation rewrites the source program into an explicitly self-adjusting target program. The type system is related to information-flow type systems and enjoys decidable type inference via constraint solving. We prove that the translation outputs well-typed self-adjusting programs and preserves the source program's input-output behavior, guaranteeing that translated programs respond correctly to all changes to their data. Using a cost semantics, we also prove that the translation preserves the asymptotic complexity of the source program.

**Categories and Subject Descriptors** D.1.1 [Programming Techniques]: Applicative (Functional) Programming; F.3.3 [Logics and Meanings of Programs]: Studies of Program Constructs

**General Terms** Algorithms, Languages, Performance

## 1. Introduction

Dynamic changes are pervasive in computational problems: physics simulations often involve moving objects; robots interact with dynamic environments; compilers must respond to slight modifications in their input programs. Such dynamic changes are often small, or *incremental*, and result in only slightly different output, so computations can often respond to them asymptotically faster than performing a complete re-computation. Such asymptotic improvements can lead to massive speedup in practice but tradition-

ally require careful algorithm design and analysis (e.g., Chiang and Tamassia [1992]; Guibas [2004]; Demetrescu et al. [2005]), which can be challenging even for seemingly simple problems.

Motivated by this problem, researchers have developed language-based techniques that enable computations to respond to dynamic data changes automatically and efficiently (see Ramalingam and Reps [1993] for a survey). This line of research, traditionally known as *incremental computation*, aims to reduce dynamic problems to static (conventional or batch) problems by developing compilers that automatically generate code for dynamic responses. This is challenging, because the compiler-generated code aims to handle changes asymptotically faster than the source code. Early proposals [Demers et al. 1981; Pugh and Teitelbaum 1989; Field and Teitelbaum 1990] were limited to certain classes of applications (e.g., attribute grammars), allowed limited forms of data changes, and/or yielded suboptimal efficiency. Some of these approaches, however, had the important advantage of being *implicit*: they required little or no change to the program code to support dynamic change—conventional programs could be compiled to executables that respond automatically to dynamic changes.

Recent work based on *self-adjusting computation* made progress towards achieving efficient incremental computation by providing algorithmic language abstractions to express computations that respond automatically to changes to their data [Ley-Wild et al. 2008; Acar et al. 2009]. Self-adjusting computation can deliver asymptotically efficient updates in a reasonably broad range of problem domains [Acar et al. 2007, 2010a], and have even helped solve challenging open problems [Acar et al. 2010b]. Existing self-adjusting computation techniques, however, require the programmer to program *explicitly* by using a certain set of primitives [Carlsson 2002; Ley-Wild et al. 2008; Acar et al. 2009]. Specifically the programmer must manually distinguish *stable data*, which remains the same, from *changeable data*, which can change over time, and operate on changeable data via a special set of primitives. As a result, rewriting a conventional program into a self-adjusting program requires extensive changes to the code. For example, a purely functional program will need to be rewritten in imperative style using write-once, monadic references.

In this paper, we present techniques for *implicit* self-adjusting computation that allow conventional programs to be translated automatically into efficient self-adjusting programs. Our approach consists of a type system for inferring self-adjusting computation types from purely functional programs and a type-guided translation algorithm that rewrites purely functional programs into self-adjusting programs.

The type system hinges on a key observation connecting self-adjusting computation to information flow [Pottier and Simonet 2003; Sabelfeld and Myers 2003]: both involve tracking data dependencies (of changeable data and sensitive data, respectively) as well as dependencies between expressions and data. Specifi-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ICFP'11, September 19–21, 2011, Tokyo, Japan.  
Copyright © 2011 ACM 978-1-4503-0865-6/11/09...\$10.00

cally, we show that a type system that encodes the changeability of data and expressions in self-adjusting computation as secrecy of information suffices to statically enforce the invariants needed by self-adjusting computation. The type system uses polymorphism to capture stable and changeable uses of the same data or expression. Our type system admits a constraint-based formulation where the constraints are a strict subset of those needed by traditional information-flow type systems. Consequently, as with information flow, our type system admits an HM(X) inference algorithm [Odersky et al. 1999] that can infer all type annotations from top-level type specifications on the input of a program.


For this work, determination of types via type inference is not an end unto itself but a means for translating purely functional programs into self-adjusting programs. To achieve this, we first present a set of compositional, non-deterministic transformation rules. Guided by the types, the rules identify the set of all changeable expressions that operate on changeable data and rewrite them into the self-adjusting target language. We then present a deterministic translation algorithm that applies the compositional rules judiciously, considering the type and context (enclosing expressions) of each translated subexpression, to generate a well-typed self-adjusting target program.

Taken together, the type system, its inference algorithm, and the translation algorithm enable translating purely functional source programs to self-adjusting target programs using top-level type annotations on the input type of the source program. These top-level type annotations simply mark what part of the input data is subject to change. Figure 1 illustrates how source programs written in Level ML, a purely functional subset of ML with *level types*, can be translated to self-adjusting programs in the target language AFL, a language for self-adjusting computation with explicit primitives [Acar et al. 2006]. We prove three critical properties of the approach.

- **Type soundness.** On source code of a given type, the translation algorithm produces well-typed self-adjusting code of a corresponding target type (Theorem 6.1).
- **Observational equivalence.** The translated self-adjusting program, when evaluated, produces the same value as the source program (Theorem 6.5).
- **Asymptotic complexity.** The time to evaluate the translated program is asymptotically the same as the time to evaluate the source program (Theorem 6.9).

Type soundness and observational equivalence together imply a critical consistency property: that self-adjusting programs respond correctly to changing data (via the consistency of the target self-adjusting language [Acar et al. 2006]). The third property shows that the translated program takes asymptotically as long to evaluate (from scratch) as the corresponding source program. In addition, it places a worst-case bound on the time taken to self-adjust via change propagation, which can and often does take significantly less time when data changes are small. To prove this complexity result, we use a cost semantics [Sands 1990; Sansom and Peyton Jones 1995] that enables precise reasoning about the complexity of the evaluation time. We do not, however, prove tighter bounds on the complexity of self-adjustments; this would be beyond the scope of this paper.

We intend to complete an implementation of our approach as an extension of Standard ML and the MLton compiler [MLton]. However, we expect the proposed approach could be implemented in other languages such as Haskell, where self-adjusting libraries also exist [Carlsson 2002]. In general, since our approach simply generates target code, it is agnostic to implementation details of



**Figure 1.** Visualizing the translation between the source language Level ML and the target language AFL, and related properties.

the explicit self-adjusting-computation mechanisms employed in the target language and thus can be applied broadly.

**Paper guide.** We find it better to give an overview of the proposed approach by focusing on the translation problem and working back to the type system in a “top-down” manner (Section 2). The details of the translation algorithm and our theorems, however, rely on the type system. We therefore take a more “bottom-up” approach in the rest of the paper: we first present the static semantics (the syntax and the type system) (Sections 3 and 4), and then describe the target language AFL (Section 5) and the translation (Section 6). Finally, we discuss related work (Section 7) and conclude. Due to space restrictions, we include all the proofs in the appendix [Chen et al. 2011].

## 2. Overview

We present an informal overview of our approach via examples. First we briefly describe explicit self-adjusting computation, as laid out in previous work, and which we use as a target language. Then we outline our proposed approach.

### 2.1 Explicit Self-Adjusting Computation

The key concept behind explicit approaches is the notion of a *modifiable (reference)*, which stores *changeable* values that can change over time [Acar et al. 2006]. The programmer operates on modifiabls with **mod**, **read**, and **write** constructs to create, read from, and write into modifiabls. The run-time system of a self-adjusting language uses these constructs to represent the execution as a graph, enabling efficient *change propagation* when the data changes in small amounts.

As an example, consider a trivial program that computes  $x^2 + y$ :

```
squareplus: int * int → int
fun squareplus (x, y) =
  let x2 = x * x in
  let r = x2 + y in
  r
```

To make this program self-adjusting with respect to changes in  $y$ , while leaving  $x$  unchanging or *stable*, we assign  $y$  the type `int mod` (of modifiabls containing integers) and **read** the contents of the modifiable. The body of the **read** is a *changeable expression* ending with a **write**. This function has a changeable arrow type  $\overrightarrow{\cdot}$ :

```
squareplus_SC: int * int mod  $\overrightarrow{\cdot}$  int
fun squareplus_SC (x, y) =
  let x2 = x * x in
  read y as y' in
  let r = x2 + y' in
  write(r)
```

The **read** operation delineates the code that depends on the changeable value  $y$ , and the changeable arrow type ensures a critical consistency property:  $\overrightarrow{\cdot}$ -functions can only be called within

the context of a changeable expression. If we change the value of  $y$ , change propagation can update the result, re-executing only the read and its body, reusing the computation of the square  $x2$ .

Suppose we wish to make  $x$  changeable while leaving  $y$  stable. We need to read  $x$  and place  $x2$  into a modifiable (because we can only read within the context of a changeable expression), and immediately read back  $x2$  and finish by writing the sum. (To avoid creating this modifiable would require further structural changes to the code.)

```
squareplus_CS: int mod * int  $\xrightarrow{c}$  int
fun squareplus_CS (x, y) =
  let x2 = mod (read x as x' in write(x' * x')) in
    read x2 as x2' in
      let r = x2' + y in
        write(r)
```

As this example shows, rewriting even a trivial program can require modifications to the code, and different choices about what is or is not changeable lead to different code. Moreover, if we need `squareplus_SC` and `squareplus_CS`—for instance, if we want to pass `squareplus` to various higher-order functions—we must write, and maintain, both versions.

Conservatively treating all data as changeable would require writing just one version, but treating all data as modifiable can introduce unacceptably high overhead. At the other extreme, making everything stable requires no rewriting, but forgoes the benefits of change propagation. Instead, we take an approach where data is modifiable only where necessary.

## 2.2 Implicit Self-Adjusting Computation

To make self-adjusting computation implicit, we use type information to insert **reads**, **writes**, and **mods** automatically. The user annotates the input type of the program; we infer types for all expressions, and use this information to guide a translation algorithm. The translation algorithm returns well-typed self-adjusting target programs. The translation requires no expression-level annotations. For the example function `squareplus` above, we can automatically derive `squareplus_SC` and `squareplus_CS` from just the type of the function (expressed in a slightly different form, as we discuss next).

**Level types.** To uniformly describe source functions (more generally, expressions) that differ only in their “changeability”, we need a more general type system than that of the target language. This type system refines types with *levels*  $\mathbb{S}$  (stable) and  $\mathbb{C}$  (changeable). The type  $\mathbf{int}^\delta$  is an integer whose level is  $\delta$ ; for example, to get `squareplus_CS` we can annotate `squareplus`’s argument with the type  $\mathbf{int}^C \times \mathbf{int}^S$ .

Level types are an important connection between information-flow types [Pottier and Simonet 2003] and those needed for our translation: high-security secret data (level  $H$ ) behaves like changeable data (level  $\mathbb{C}$ ), and low-security public data (level  $L$ ) behaves like stable data (level  $\mathbb{S}$ ). In information flow, data that depends on secret data must be secret; in self-adjusting computation, data that depends on changeable data must be changeable. Building on this connection, we develop a type system with several features and mechanisms similar to information flow. Among these is level polymorphism; our type system assigns level-polymorphic types to expressions that accommodate various “changeabilities”. (As with ML’s polymorphism over types, our level polymorphism is prenex.) Another similarity is evident in our constraint-based type inference system, where the constraints are a strict subset of those in Pottier and Simonet [2003]. As a corollary, our system admits a constraint-based type inference algorithm [Odersky et al. 1999].

**Translation.** The main purpose of our type system is to support translation. Given a source expression and its type, translation in-

```
datatype  $\alpha$  list = nil | cons of  $\alpha$  *  $\alpha$  list
```

```
inc : int  $\rightarrow$  int
fun inc (x) = x+1
```

```
map : ( $\alpha \rightarrow \beta$ )  $\rightarrow$   $\alpha$  list  $\rightarrow$   $\beta$  list
fun map f l =
  case l of
    nil  $\Rightarrow$  nil
  | cons(h,t)  $\Rightarrow$  cons(f h, map f t)
```

```
mapPair : (int list * int list)  $\rightarrow$  (int list * int list)
fun mapPair (l,a) = (map inc l, map inc a)
```

Figure 2. Function `mapPair` in ML

```
datatype  $\alpha$  list $^\delta$  = nil | cons of  $\alpha$  * ( $\alpha$  list $^\delta$ )
```

```
mapPair : ((int $^S$  list $^C$ ) * (int $^C$  list $^S$ ))
 $\xrightarrow{S}$  ((int $^S$  list $^C$ ) * (int $^C$  list $^S$ ))
```

... (\*inc, map, mapPair same as in Figure 1. \*)

Figure 3. Function `mapPair` in Level ML, with level types

serts the appropriate `mod`, `read`, and `write` primitives and restructures the code to produce an expression that is well-typed in the target language. The type system of the target language, which is explicitly self-adjusting, is monomorphic in the levels or changeability, while the implicitly self-adjusting source language is polymorphic over levels. Consequently, translation also needs to monomorphize the source code. Our translation generates code that is well-typed, has the same input-output behavior as the source program, and is, at worst, a constant factor slower than the source program. Since the source and target languages differ, proving these properties is nontrivial; in fact, the proofs critically guided our formulation of the type system and translation algorithm.

**A more detailed example: `mapPair`.** To illustrate how our translation works, consider a function `mapPair` that takes two integer lists and increments the elements in both lists. This function can be written by applying the standard higher-order `map` over lists. Figure 2 shows the purely functional code in an ML-like language for an implementation of `mapPair`, with a datatype  $\alpha$  list, an increment function `inc`, and a polymorphic `map` function. Type signatures give the types of functions.

To obtain a self-adjusting `mapPair`, we first decide how we wish to allow the input to change. Suppose that we want to allow insertion and deletion of elements in the first list, but we expect the length of the second list to remain constant, with only its elements changing. We can express this with the versions of the list type with different changeability:

- $\alpha$  list $^C$  for lists of  $\alpha$  with changeable tails;
- $\alpha$  list $^S$  for lists of  $\alpha$  with stable tails.

Then a list of integers allowing insertion and deletion has type  $\mathbf{int}^S$  list $^C$ , and one with unchanging length has type  $\mathbf{int}^C$  list $^S$ . Now we can write the type annotation on `mapPair` shown in Figure 3. Given only that annotation, type inference can find appropriate types for `inc` and `map` and our translation algorithm generates self-adjusting code from these annotations. Note that to obtain a self-adjusting program, we only had to provide types for the function. We call this language with level types Level ML.

**Target code for `mapPair`.** Translating the code in Figure 3 produces the self-adjusting target code in Figure 4. Note that `inc` and `map` have *level-polymorphic* types. In `map inc l` we increment sta-

```

datatype  $\alpha$  list_S = nil | cons of  $\alpha$  *  $\alpha$  list_S
datatype  $\alpha$  list_C = nil | cons of  $\alpha$  * ( $\alpha$  list_C) mod

inc_S : int  $\xrightarrow{\mathbb{S}}$  int    (* 'inc' specialized for stable data *)
funS inc_S (x) = x+1

inc_C : int  $\xrightarrow{\mathbb{C}}$  int    (* 'inc' specialized for changeable data *)
funC inc_C (x) = read x as x' in write (x'+1)

inc :  $\forall \delta. \text{int}^\delta \xrightarrow{\delta} \text{int}^\delta$ 
val inc = select { $\delta=\mathbb{S} \Rightarrow \text{inc}_S$ 
                 |  $\delta=\mathbb{C} \Rightarrow \text{inc}_C$ }

map_SC : ( $\alpha \xrightarrow{\mathbb{S}} \beta$ )  $\xrightarrow{\mathbb{S}}$  ( $\alpha$  list_C) mod  $\xrightarrow{\mathbb{S}}$  ( $\beta$  list_C) mod
funS map_SC f l = (* 'map' for stable heads, changeable tails *)
mod (read l as x in
  case x of
    nil  $\Rightarrow$  write nil
  | cons(h,t)  $\Rightarrow$  write (cons(f h, map_SC f t)))

map_CS : ( $\alpha \xrightarrow{\mathbb{C}} \beta$ )  $\xrightarrow{\mathbb{S}}$  ( $\alpha$  list_S)  $\xrightarrow{\mathbb{S}}$  ( $\beta$  list_S)
funS map_CS f l = (* 'map' for changeable heads, stable tails *)
case l of
  nil  $\Rightarrow$  nil
  | cons(h,t)  $\Rightarrow$  let val h' = mod (f h)
  in cons(h', map_CS f t)

map :  $\forall \delta_h, \delta_t. (\alpha \xrightarrow{\delta_h} \beta) \xrightarrow{\mathbb{S}} \alpha \text{list}^{\delta_h} \xrightarrow{\mathbb{S}} \beta \text{list}^{\delta_t}$ 
val map = select { $\delta_h=\mathbb{S}, \delta_t=\mathbb{C} \Rightarrow \text{map}_SC$ 
                 |  $\delta_h=\mathbb{C}, \delta_t=\mathbb{S} \Rightarrow \text{map}_CS$ }

mapPair : ((int list_C) mod * (int mod) list_S)
 $\xrightarrow{\mathbb{S}}$  ((int list_C) mod * (int mod) list_S)
funS mapPair (l, a) = (map[ $\delta_h=\mathbb{S}, \delta_t=\mathbb{C}$ ] inc[ $\delta=\mathbb{S}$ ] l,
                      map[ $\delta_h=\mathbb{C}, \delta_t=\mathbb{S}$ ] inc[ $\delta=\mathbb{C}$ ] a)

```

**Figure 4.** Translated mapPair with **mod** types and explicit level polymorphism.

ble integers, and in map inc a we increment changeable integers, so the type inferred for inc must be generic:  $\forall \delta. \text{int}^\delta \xrightarrow{\delta} \text{int}^\delta$ . Our translation produces two implementations of inc, one per instantiation ( $\delta=\mathbb{S}$  and  $\delta=\mathbb{C}$ ): inc\_S and inc\_C (in Figure 4). Since we want to use inc with the higher-order function map, we need to generate a “selector” function that takes an instantiation and picks out the appropriate implementation:

```

inc :  $\forall \delta. \text{int}^\delta \xrightarrow{\delta} \text{int}^\delta$ 
val inc = select { $\delta=\mathbb{S} \Rightarrow \text{inc}_S$ 
                 |  $\delta=\mathbb{C} \Rightarrow \text{inc}_C$ }

```

In mapPair itself, we pass a level instantiation to the selector: inc<sup>[ $\delta=\mathbb{S}$ ]</sup>. (This instantiation is known statically, so it could be replaced with inc\_S at compile time.)

Observe how the single annotation on mapPair led to duplication of the two functions it uses. While inc\_S is the same as the original inc, the changeable version inc\_C adds a **read** and a **write**. Note also that the two generated versions of map are both different from the original.

**The interplay of type inference and translation.** Given user annotations on the input, type inference finds a satisfying type assignment, which then guides our translation algorithm to produce self-adjusting code. In many cases, multiple type assignments could satisfy the annotations; for example, subsumption allows any stable type to be promoted to a changeable type. Translation yields target code that satisfies the crucial type soundness, operational equivalence, and complexity properties under any satisfying assignment.

<i>Levels</i>	$\delta, \varepsilon ::= \mathbb{S} \mid \mathbb{C} \mid \alpha$
<i>Types</i>	$\tau ::= \text{int}^\delta \mid (\tau_1 \times \tau_2)^\delta \mid (\tau_1 + \tau_2)^\delta \mid (\tau_1 \xrightarrow{\varepsilon} \tau_2)^\delta$
<i>Constraints</i>	$C, D ::= \text{true} \mid \text{false} \mid \exists \bar{\alpha}. C \mid C \wedge D \mid$ $\alpha = \beta \mid \alpha \leq \beta \mid \delta \triangleleft \tau$
<i>Type schemes</i>	$\sigma ::= \tau \mid \forall \bar{\alpha} [D]. \tau$

**Figure 5.** Levels, constraints, types, and type schemes

But some type assignments are preferable, especially when one considers constant factors. Choosing  $\mathbb{C}$  levels whenever possible is always a viable strategy, but treating all data as changeable results in more overhead. As in information flow, where we want to consider data secret only when absolutely necessary, inference yields principal typings that are minimally changeable, always preferring  $\mathbb{S}$  over  $\mathbb{C}$ .

### 3. From Information Flow Types to SAC

Self-adjusting computation separates the computation and data into two parts: stable and changeable. Changeable data refers to data that can change over time; all non-changeable data is stable. Similarly, changeable expressions refers to expressions that operate (via elimination forms) on changeable data; all non-changeable expressions are stable. Evaluation of changeable expressions (that is, changeable computations) can change as the data that they operate on changes: changes in data cause changes in control flow. These distinctions are critical to effective self-adjustment: previous work shows that it suffices to track and remember changeable data and evaluations of changeable expressions because stable data and evaluations of stable expressions remain invariant over time. Previous work therefore presents languages that enable the programmer to separate stable and changeable data, and type systems that enforce the correct usage of these constructs.

In this section, we describe the self-adjusting computation types that we infer for purely functional programs. A key insight behind our approach is that in information-flow type systems, secret (high-security) data is infectious: any data that depends on secret data itself must be secret. This corresponds to self-adjusting computation: data that depends on changeable data must itself be changeable. In addition, self-adjusting computation requires expressions that inspect changeable data—elimination forms—to be changeable. To encode this invariant, we extend function types with a *mode*, which is either stable or changeable; only changeable functions can inspect changeable data. This additional structure preserves the spirit of information flow-based type systems, and, moreover, supports constraint-based type inference in a similar style.

The starting point for our formulation is Pottier and Simonet [2003]. Our types include two (*security*) levels, stable and changeable. We generally follow their approach and notation. The two key differences are that (1) since Level ML is purely functional, we need no “program counter” level “pc”; (2) we need a mode  $\varepsilon$  on function types.

**Levels.** The levels  $\mathbb{S}$  (*stable*) and  $\mathbb{C}$  (*changeable*) have a total order:

$$\overline{\mathbb{S}} \leq \overline{\mathbb{S}} \qquad \overline{\mathbb{C}} \leq \overline{\mathbb{C}} \qquad \overline{\mathbb{S}} \leq \overline{\mathbb{C}}$$

To support polymorphism and enable type inference, we allow *level variables*  $\alpha, \beta$  to appear in types.

**Types.** Types consist of integers tagged with their level, products<sup>1</sup> and sums with an associated level, and arrow (function) types. Function types  $(\tau_1 \xrightarrow{\varepsilon} \tau_2)^\delta$  carry two level annotations  $\varepsilon$  and  $\delta$ .

<sup>1</sup> In Pottier and Simonet [2003], product types are low-security (stable) because pairing adds no extra information. In our setting, changeable products give more control over the granularity of change propagation.

$$\begin{array}{c}
\frac{\delta \leq \delta'}{\mathbf{int}^\delta <: \mathbf{int}^{\delta'}} \text{ (subInt)} \quad \frac{\tau_1 <: \tau_1' \quad \tau_2 <: \tau_2' \quad \delta \leq \delta'}{(\tau_1 \times \tau_2)^\delta <: (\tau_1' \times \tau_2')^{\delta'}} \text{ (subProd)} \\
\frac{\tau_1 <: \tau_1' \quad \tau_2 <: \tau_2' \quad \delta \leq \delta'}{(\tau_1 + \tau_2)^\delta <: (\tau_1' + \tau_2')^{\delta'}} \text{ (subSum)} \\
\frac{\varepsilon = \varepsilon' \quad \delta \leq \delta' \quad \tau_1' <: \tau_1 \quad \tau_2 <: \tau_2'}{(\tau_1 \xrightarrow{\varepsilon} \tau_2)^\delta <: (\tau_1' \xrightarrow{\varepsilon'} \tau_2)^{\delta'}} \text{ (subArrow)}
\end{array}$$

**Figure 6.** Subtyping

$$\begin{array}{c}
\frac{\delta \leq \delta'}{\delta < \mathbf{int}^{\delta'}} \text{ (<-Int)} \quad \frac{\delta \leq \delta'}{\delta < (\tau_1 \times \tau_2)^{\delta'}} \text{ (<-Prod)} \\
\frac{\delta \leq \delta'}{\delta < (\tau_1 \xrightarrow{\varepsilon} \tau_2)^{\delta'}} \text{ (<-Arrow)} \quad \frac{\delta \leq \delta'}{\delta < (\tau_1 + \tau_2)^{\delta'}} \text{ (<-Sum)}
\end{array}$$

**Figure 7.** Lower bound of a type

$$\begin{array}{c}
\overline{\mathbf{int}^{\mathbb{S}} \text{ O.S.}} \quad \overline{(\tau_1 \xrightarrow{\varepsilon} \tau_2)^{\mathbb{S}} \text{ O.S.}} \quad \overline{(\tau_1 \times \tau_2)^{\mathbb{S}} \text{ O.S.}} \quad \overline{(\tau_1 + \tau_2)^{\mathbb{S}} \text{ O.S.}} \\
\overline{\mathbf{int}^{\mathbb{C}} \text{ O.C.}} \quad \overline{(\tau_1 \xrightarrow{\varepsilon} \tau_2)^{\mathbb{C}} \text{ O.C.}} \quad \overline{(\tau_1 \times \tau_2)^{\mathbb{C}} \text{ O.C.}} \quad \overline{(\tau_1 + \tau_2)^{\mathbb{C}} \text{ O.C.}} \\
\mathbf{int}^{\delta_1} \doteq \mathbf{int}^{\delta_2} \quad (\tau_1 \times \tau_2)^{\delta_1} \doteq (\tau_1 \times \tau_2)^{\delta_2} \quad (\tau_1 + \tau_2)^{\delta_1} \doteq (\tau_1 + \tau_2)^{\delta_2} \\
(\tau_1 \xrightarrow{\varepsilon} \tau_2)^{\delta_1} \doteq (\tau_1 \xrightarrow{\varepsilon} \tau_2)^{\delta_2}
\end{array}$$

**Figure 8.** Outer-stable and outer-changeable types, and equality up to outer levels

The *mode*  $\varepsilon$  is the level of the computation encapsulated by the function. This mode determines how a function can manipulate changeable values: a function in stable mode cannot directly manipulate changeable values; it can only pass them around. By contrast, a changeable-mode function can directly manipulate changeable values. The outer level  $\delta$  is the level of the function itself, as a value. We say that a type is *ground* if it contains no level variables.

**Subtyping.** Figure 6 shows the subtyping relation  $\tau <: \tau'$ , which is standard except for the levels. It requires that the outer level of the subtype is smaller than the outer level of the supertype and that the modes match in the case of functions: a stable-mode function is never a subtype or supertype of a changeable-mode function. (It would be sound to make stable-mode functions subtypes of changeable-mode functions, but changeable mode functions are more expensive; silent coercion would make performance less predictable.)

**Levels and types.** We rely on several relations between levels and types to ascertain various invariants. A type  $\tau$  is *higher than*  $\delta$ , written  $\delta < \tau$ , if the outer level of the type is at least  $\delta$ . In other words,  $\delta$  is a lower bound of the outer level(s) of  $\tau$ . Figure 7 defines this relation. We distinguish between outer-stable and outer-changeable types (Figure 8). We write  $\tau$  O.S. if the outer level of  $\tau$  is  $\mathbb{S}$ . Similarly, we write  $\tau$  O.C. if the outer level of  $\tau$  is  $\mathbb{C}$ . Finally, two types  $\tau_1$  and  $\tau_2$  are *equal up to their outer levels*, written  $\tau_1 \doteq \tau_2$ , if  $\tau_1 = \tau_2$  or they differ only in their outer levels.

**Constraints.** To perform type inference, we extend levels with level variables  $\alpha$  and  $\beta$ , and use a constraint solver to find solu-

*Values*  $v ::= n \mid x \mid (v_1, v_2) \mid \mathbf{inl} \ v \mid \mathbf{inr} \ v \mid \mathbf{fun} \ f(x) = e$   
*Expr.'s*  $e ::= v \mid \oplus(x_1, x_2) \mid \mathbf{fst} \ x \mid \mathbf{snd} \ x \mid$   
 $\mathbf{case} \ x \ \mathbf{of} \ \{x_1 \Rightarrow e_1, x_2 \Rightarrow e_2\} \mid$   
 $\mathbf{apply}(x_1, x_2) \mid \mathbf{let} \ x = e_1 \ \mathbf{in} \ e_2$

**Figure 9.** Abstract syntax of the source language Level ML

tions for the variables. Our constraints  $C, D$  include level-variable comparisons  $\leq$  and level-type comparisons  $\delta < \tau$ , which type inference composes into conjunctions of satisfiability predicates  $\exists \vec{\alpha}. C$ .

The subtyping and lower bound relations defined in Figures 6 and 7 consider closed types only. For type inference, we can extend these with a constraint to allow non-closed types.

A (*ground*) *assignment*, written  $\phi$ , substitutes concrete levels  $\mathbb{S}$  and  $\mathbb{C}$  for level variables. An assignment  $\phi$  satisfies a constraint  $C$ , written  $\phi \vdash C$ , if and only if  $C$  holds true after the substitution of variables to ground types as specified by  $\phi$ . We say that  $C$  *entails*  $D$ , written  $C \Vdash D$ , if and only if every assignment  $\phi$  that satisfies  $C$  also satisfies  $D$ . We write  $\phi(\alpha)$  for the solution of  $\alpha$  in  $\phi$ , and  $[\phi]\tau$  for the usual substitution operation on types. For example, if  $\phi(\alpha) = \mathbb{S}$  then  $[\phi]((\mathbf{int}^\alpha + \mathbf{int}^{\mathbb{C}})^\alpha) = (\mathbf{int}^{\mathbb{S}} + \mathbf{int}^{\mathbb{C}})^{\mathbb{S}}$ .

**Type schemes.** A *type scheme*  $\sigma$  is a type with universally quantified level variables:  $\sigma = \forall \vec{\alpha}[D]. \tau$ . We say that the variables  $\vec{\alpha}$  are bound by  $\sigma$ . The type scheme is bounded by the constraint  $D$ , which specifies the conditions that must hold on the variables. As usual, we consider type schemes equivalent under capture-avoiding renaming of their bound variables. Ground types can be written as type schemes, e.g.  $\mathbf{int}^{\mathbb{C}}$  as  $\forall \emptyset[\mathbf{true}]. \mathbf{int}^{\mathbb{C}}$ .

## 4. Source Language

### 4.1 Static Semantics

**Syntax.** Figure 9 shows the syntax for our source language Level ML, a purely functional language with integers (as base types), products, and sums. The expressions consist of values (integers, pairs, tagged values, recursive functions), projections, case expressions, function applications, and let bindings. For convenience, we consider only expressions in A-normal form, which names intermediate results. A-normal form simplifies some technical issues, while maintaining expressiveness.

**Constraint-based type system.** We could define types as

$$\tau ::= \mathbf{int} \mid \tau_1 \times \tau_2 \mid \tau_1 + \tau_2 \mid \tau_1 \rightarrow \tau_2$$

Such a type system would be completely standard. Instead, we use a richer type system that allows us to directly translate Level ML programs into self-adjusting programs in AFL. This constraint-based type system has the level-decorated types, constraints, and type schemes in Figure 5 and described in Section 3. After discussing the rules themselves, we will look at type inference (Section 4.2).

Typing takes place in the context of a constraint formula  $C$  and a typing environment  $\Gamma$  that maps variables to type schemes:  $\Gamma ::= \cdot \mid \Gamma, x : \sigma$ . The typing judgment  $C; \Gamma \vdash_\varepsilon e : \tau$  has a constraint  $C$  and typing environment  $\Gamma$ , and infers type  $\tau$  for expression  $e$  in mode  $\varepsilon$ . Beyond the usual typing concerns, there are three important aspects of the typing rules: the determination of modes and levels, level polymorphism, and constraints. To help separate concerns, we discuss constraints later in the section—at this time, the reader can ignore the constraints in the rules and read  $C; \Gamma \vdash_\varepsilon e : \tau$  as  $\Gamma \vdash_\varepsilon e : \tau$ , read  $C \Vdash \delta < \tau_2$  as  $\delta < \tau_2$ , and so on.

The mode of each typing judgment affects the types that can be used “directly” by the expression being typed. Specifically, the mode discipline prevents the elimination forms from being applied

$$\boxed{C; \Gamma \vdash_{\varepsilon} e : \tau} \quad \text{Under constraint } C \text{ and source typing environment } \Gamma, \text{ source expression } e \text{ has type } \tau$$

$$\frac{}{C; \Gamma \vdash_{\varepsilon} n : \mathbf{int}^{\mathbb{S}}} \text{(SInt)} \quad \frac{\Gamma(x) = \forall \vec{\alpha}[D]. \tau \quad C \Vdash \exists \vec{\beta}. [\vec{\beta}/\vec{\alpha}]D}{C \wedge [\vec{\beta}/\vec{\alpha}]D; \Gamma \vdash_{\varepsilon} x : [\vec{\beta}/\vec{\alpha}]\tau} \text{(SVar)}$$

$$\frac{C; \Gamma \vdash_{\varepsilon} v_1 : \tau_1 \quad C; \Gamma \vdash_{\varepsilon} v_2 : \tau_2}{C; \Gamma \vdash_{\varepsilon} (v_1, v_2) : (\tau_1 \times \tau_2)^{\mathbb{S}}} \text{(SPair)}$$

$$\frac{C; \Gamma \vdash_{\varepsilon} v : \tau_1}{C; \Gamma \vdash_{\varepsilon} \mathbf{inl} v : (\tau_1 + \tau_2)^{\mathbb{S}}} \text{(SSum)}$$

$$\frac{C; \Gamma, x : \tau_1, f : (\tau_1 \xrightarrow{\varepsilon} \tau_2)^{\mathbb{S}} \vdash_{\varepsilon} e : \tau_2 \quad C \Vdash \varepsilon \triangleleft \tau_2}{C; \Gamma \vdash_{\varepsilon'} (\mathbf{fun} f(x) = e) : (\tau_1 \xrightarrow{\varepsilon} \tau_2)^{\mathbb{S}}} \text{(SFun)}$$

$$\frac{C; \Gamma \vdash_{\mathbb{S}} x_1 : \mathbf{int}^{\delta_1} \quad C \Vdash \delta_1 = \delta_2 \quad C; \Gamma \vdash_{\mathbb{S}} x_2 : \mathbf{int}^{\delta_2} \quad C \Vdash \delta_1 \leq \varepsilon \quad \oplus : \mathbf{int} \times \mathbf{int} \rightarrow \mathbf{int}}{C; \Gamma \vdash_{\varepsilon} \oplus(x_1, x_2) : \mathbf{int}^{\delta_1}} \text{(SPrim)}$$

$$\frac{C; \Gamma \vdash_{\mathbb{S}} x : (\tau_1 \times \tau_2)^{\delta} \quad C \Vdash \delta \leq \varepsilon}{C; \Gamma \vdash_{\varepsilon} \mathbf{fst} x : \tau_1} \text{(SFst)}$$

$$\frac{C; \Gamma \vdash_{\varepsilon'} e_1 : \tau' \quad C; \Gamma, x : \tau'' \vdash_{\varepsilon} e_2 : \tau \quad C \Vdash \tau' <: \tau'' \quad C \Vdash \tau' \doteq \tau''}{C; \Gamma \vdash_{\varepsilon} \mathbf{let} x = e_1 \mathbf{in} e_2 : \tau} \text{(SLetE)}$$

$$\frac{C \wedge D; \Gamma \vdash_{\mathbb{S}} v_1 : \tau' \quad C; \Gamma, x : \forall \vec{\alpha}[D]. \tau'' \vdash_{\varepsilon} e_2 : \tau \quad \vec{\alpha} \cap FV(C, \Gamma) = \emptyset \quad C \Vdash \tau' <: \tau'' \quad C \Vdash \tau' \doteq \tau''}{C \wedge \exists \vec{\alpha}. D; \Gamma \vdash_{\varepsilon} \mathbf{let} x = v_1 \mathbf{in} e_2 : \tau} \text{(SLetV)}$$

$$\frac{C; \Gamma \vdash_{\mathbb{S}} x_1 : (\tau_1 \xrightarrow{\varepsilon'} \tau_2)^{\delta} \quad C \Vdash \varepsilon' = \varepsilon \quad C; \Gamma \vdash_{\mathbb{S}} x_2 : \tau_1 \quad C \Vdash \delta \triangleleft \tau_2}{C; \Gamma \vdash_{\varepsilon} \mathbf{apply}(x_1, x_2) : \tau_2} \text{(SApp)}$$

$$\frac{C; \Gamma \vdash_{\mathbb{S}} x : (\tau_1 + \tau_2)^{\delta} \quad C; \Gamma, x_1 : \tau_1 \vdash_{\varepsilon} e_1 : \tau \quad C \Vdash \delta \leq \varepsilon \quad C \Vdash \delta \triangleleft \tau \quad C; \Gamma, x_2 : \tau_2 \vdash_{\varepsilon} e_2 : \tau}{C; \Gamma \vdash_{\varepsilon} \mathbf{case} x \mathbf{of} \{x_1 \Rightarrow e_1, x_2 \Rightarrow e_2\} : \tau} \text{(SCase)}$$

Figure 10. Typing rules for Level ML

to changeable values in the stable mode. This is a key principle of the type system.

No computation happens in values, so they can be typed in either mode. The typing rules for variables (SVar), integers (SInt), pairs (SPair), and sums (SSum) are otherwise standard (we omit the symmetric judgment  $\mathbf{inr} v$ ). Rule (SVar) instantiates a variable’s polymorphic type. For clarity, we also make explicit the renaming of the quantified type variables  $\vec{\alpha}$  to some fresh  $\vec{\beta}$  (which will be instantiated later by constraint solving). To type a function (SFun), we type the body in the mode  $\varepsilon$  specified by the function type  $(\tau_1 \xrightarrow{\varepsilon} \tau_2)^{\delta}$ , and require the result type  $\tau_2$  to be higher than the mode,  $\varepsilon \triangleleft \tau_2$ . As a result, a changeable-mode function must have a changeable return type. This captures the idea that a changeable-mode function is a computation that depends on changeable data, and thus its result must accommodate changes to that data. Primitive operators  $\oplus$  take two stable integers and return a stable integer result.

As is common in Damas-Milner-style systems, when typing  $\mathbf{let}$  we can generalize variables in types (in our system, level variables) to yield a polymorphic value only when the bound expression is a value. This *value restriction* is not essential because Level ML is pure, but facilitates adding side effects at a later date. In the first

case (SLetE), the expression bound may be a non-value, so we do not generalize and simply type the body in the same mode as the whole  $\mathbf{let}$ , assuming that the bound expression has the specified type in any mode  $\varepsilon'$ .<sup>2</sup> We allow subsumption only when the subtype and supertype are equal up to their outer levels, e.g. from a bound expression  $e_1$  of subtype  $\mathbf{int}^{\mathbb{S}}$  to an assumption  $x : \mathbf{int}^{\mathbb{C}}$ . This simplifies the translation, with no loss of expressiveness: to handle “deep” subsumption, such as  $(\mathbf{int}^{\mathbb{S}} \xrightarrow{\mathbb{S}} \mathbf{int}^{\mathbb{S}})^{\mathbb{S}} <: (\mathbf{int}^{\mathbb{S}} \xrightarrow{\mathbb{S}} \mathbf{int}^{\mathbb{C}})^{\mathbb{C}}$ , we can insert *coercions* into the source program before typing it with these rules. (This process could easily be automated.)

In the second case (SLetV), when the expression bound is a value, we type the let expression in mode  $\varepsilon$  by typing the body in the same mode  $\varepsilon$ , assuming that the value bound is typed in the stable mode (the mode is ignored in the rules typing values). As in (SLetE), we allow subsumption on the bound value only when the types are equal up to their outer level. Because we are binding a value, we generalize its type by quantifying over the type’s free level variables.

Function application,  $\oplus$ ,  $\mathbf{fst}$ , and  $\mathbf{case}$  are the forms that eliminate values of changeable type. An application is typed in the mode  $\varepsilon'$  of the function being applied because changeable functions can operate on changeable values; the typing mode must match ( $\varepsilon' = \varepsilon$ ). Furthermore, the result of the function must be higher than the function’s level: if a function is itself changeable,  $(\tau_1 \xrightarrow{\varepsilon} \tau_2)^{\mathbb{C}}$ , then it could be replaced by another function and thus the result of this application must be changeable. (Due to  $\mathbf{let}$ -subsumption, checking this in (SFun) alone is not enough.)

The rule (SCase) types a case expression, in either mode  $\varepsilon$ , by typing each branch in  $\varepsilon$ . The mode  $\varepsilon$  must be higher than the level  $\delta$  of the scrutinee to ensure that a changeable sum type is not inspected at the stable mode. Furthermore, the level of the result  $\tau$  must also be higher than  $\delta$ : if the scrutinee changes, we may take the other branch, requiring a changeable result.

Rule (SFst) enforces a condition, similar to (SCase), that we can project out of a changeable tuple of type  $(\tau_1 \times \tau_2)^{\mathbb{C}}$  only in changeable mode. We omit the symmetric rule for  $\mathbf{snd}$ .

Our premises on variables, such as the scrutinee of (SCase), are stable-mode ( $\vdash_{\mathbb{S}}$ ), but this was an arbitrary decision; since (SVar) is the only rule that can derive such premises, their mode is irrelevant.

## 4.2 Constraints and Type Inference

Many of the rules simply pass around the constraint  $C$ . An implementation of rules with constraint-based premises, such as (SFun), implicitly adds those premises to the constraint, so that  $C = \dots \wedge (\varepsilon \triangleleft \tau_2)$ . Rule (SLetV) generalizes level variables instead of type variables, with the “occurs check”  $\vec{\alpha} \cap FV(C, \Gamma) = \emptyset$ .

Standard techniques in the tradition of Damas and Milner [1982] can infer types for Level ML. In particular, our rules and constraints fall within the HM(X) framework [Odersky et al. 1999], permitting inference of principal types via constraint solving. As always, we cannot infer the types of polymorphically recursive functions.

Using a constraint solver that, given the choice between assigning  $\mathbb{S}$  or  $\mathbb{C}$  to some level variable, prefers  $\mathbb{S}$ , inference finds principal typings that are *minimally* changeable. Thus, data and computations will only be made changeable—and incur tracking overhead—where necessary to satisfy the programmer’s annotation. This corresponds to preferring a lower security level in information flow [Pottier and Simonet 2003].

<sup>2</sup>In the target language, bound expressions must be stable-mode, but the translation puts changeable bound expressions inside a  $\mathbf{mod}$ , yielding a stable-mode bound expression.

<i>Levels</i>	$\delta, \varepsilon ::= \mathbb{S} \mid \mathbb{C}$
<i>Types</i>	$\tau ::= \mathbf{int} \mid \tau \mathbf{mod} \mid \tau_1 \times \tau_2 \mid \tau_1 + \tau_2 \mid \tau_1 \xrightarrow{\varepsilon} \tau_2$
<i>Type schemes</i>	$\sigma ::= \Pi \vec{\alpha}[D]. \tau$
<i>Typing environments</i>	$\Gamma ::= \cdot \mid \Gamma, x : \sigma \mid \Gamma, x : \tau$
<i>Variables</i>	$\underline{x} ::= x \mid x[\vec{\alpha} = \vec{\delta}]$
<i>Values</i>	$v ::= n \mid \underline{x} \mid \ell \mid (v_1, v_2) \mid \mathbf{inl} \ v \mid \mathbf{inr} \ v \mid$ $\mathbf{fun}^{\mathbb{S}} f(x) = e^{\mathbb{S}} \mid \mathbf{fun}^{\mathbb{C}} f(x) = e^{\mathbb{C}} \mid$ $\mathbf{select} \{(\vec{\alpha}_i = \vec{\delta}_i) \Rightarrow e_i\}_i$
<i>Expressions</i>	$e ::= e^{\mathbb{S}} \mid e^{\mathbb{C}}$
<i>Stable expressions</i>	$e^{\mathbb{S}} ::= v \mid \oplus(x_1, x_2) \mid \mathbf{fst} \ \underline{x} \mid \mathbf{snd} \ \underline{x} \mid$ $\mathbf{apply}^{\mathbb{S}}(x_1, x_2) \mid \mathbf{let} \ x = e^{\mathbb{S}} \ \mathbf{in} \ e^{\mathbb{S}} \mid$ $\mathbf{case} \ \underline{x} \ \mathbf{of} \ \{x_1 \Rightarrow e^{\mathbb{S}}, x_2 \Rightarrow e^{\mathbb{S}}\} \mid$ $\mathbf{mod} \ e^{\mathbb{C}}$
<i>Changeable expressions</i>	$e^{\mathbb{C}} ::= \mathbf{apply}^{\mathbb{C}}(x_1, x_2) \mid \mathbf{let} \ x = e^{\mathbb{S}} \ \mathbf{in} \ e^{\mathbb{C}} \mid$ $\mathbf{case} \ \underline{x} \ \mathbf{of} \ \{x_1 \Rightarrow e^{\mathbb{C}}, x_2 \Rightarrow e^{\mathbb{C}}\} \mid$ $\mathbf{read} \ \underline{x} \ \mathbf{as} \ y \ \mathbf{in} \ e^{\mathbb{C}} \mid \mathbf{write}(x)$

**Figure 11.** Types and expressions in the target language AFL

### 4.3 Dynamic Semantics

The call-by-value semantics of source programs is defined by a big-step judgment  $e \Downarrow v$ , read “ $e$  evaluates to value  $v$ ”. Our rules in Figure 13 are standard; we write  $[v/x]e$  for capture-avoiding substitution of  $v$  for the variable  $x$  in  $e$ .

## 5. Target Language

The target language AFL (Figure 11) is a self-adjusting language with modifiabiles. In addition to integers, products, and sums, the target type system makes a modal distinction between ordinary types (e.g.  $\mathbf{int}$ ) and modifiable types (e.g.  $\mathbf{int} \mathbf{mod}$ ). It also distinguishes stable-mode and changeable-mode functions. Level polymorphism is supported through an explicit  $\mathbf{select}$  construct and an explicit polymorphic instantiation. In Section 6, we describe how polymorphic source expressions become  $\mathbf{select}$ s in AFL.

The values of the language are integers, variables, polymorphic variable instantiation  $x[\vec{\alpha} = \vec{\delta}]$ , locations  $\ell$  (which appear only at runtime), pairs, tagged values, stable and changeable functions, and the  $\mathbf{select}$  construct, which acts as a function and case expression on levels: if  $x$  is bound to  $\mathbf{select} \{(\alpha = \mathbb{S}) \Rightarrow e_1 \mid (\alpha = \mathbb{C}) \Rightarrow e_2\}$  then  $x[\alpha = \mathbb{S}]$  yields  $e_1$ . The symbol  $\underline{x}$  stands for a bare variable  $x$  or an instantiation  $x[\vec{\alpha} = \vec{\delta}]$ .

We distinguish stable expressions  $e^{\mathbb{S}}$  from changeable expressions  $e^{\mathbb{C}}$ . Stable expressions create purely functional values;  $\mathbf{apply}^{\mathbb{S}}$  applies a stable-mode function. The  $\mathbf{mod}$  construct evaluates a changeable expression and writes the output value to a modifiable, yielding a location, which is a stable expression. Changeable expressions are computations that end in a  $\mathbf{write}$  of a pure value. Changeable-mode application  $\mathbf{apply}^{\mathbb{C}}$  applies a changeable-mode function.

The  $\mathbf{let}$  construct is either stable or changeable according to its body. When the body is a changeable expression,  $\mathbf{let}$  enables a changeable computation to evaluate a stable expression and bind its result to a variable. The  $\mathbf{case}$  expression is likewise stable or changeable, according to its case arms. The  $\mathbf{read}$  expression binds the contents of a modifiable  $\underline{x}$  to a variable  $y$  and evaluates the body of the  $\mathbf{read}$ .

The typing rules in Figure 12 follow the structure of the expressions. Rule (TSelect) checks that each monomorphized expression  $e_i$  in a  $\mathbf{select}$  has type  $\llbracket [\vec{\delta}/\vec{\alpha}] \tau \rrbracket$ , where  $[\vec{\delta}/\vec{\alpha}] \tau$  is a source-level

$\Lambda; \Gamma \vdash_{\varepsilon} v : \sigma$	Under store typing $\Lambda$ and target typing environment $\Gamma$ , target value $v$ has type scheme $\sigma$
$\Lambda; \Gamma \vdash_{\mathbb{S}} \mathbf{select} \{ \vec{\delta}_i \Rightarrow e_i \}_i : \Pi \vec{\alpha}[D]. \tau$	for all $\vec{\delta}_i$ such that $\vec{\alpha} = \vec{\delta}_i \Vdash D$ $\Lambda; \Gamma \vdash_{\mathbb{S}} e_i : \llbracket [\vec{\delta}_i/\vec{\alpha}] \tau \rrbracket$ (TSelect)
$\Lambda; \Gamma \vdash_{\varepsilon} e^{\varepsilon} : \tau$	Under store typing $\Lambda$ and target typing environment $\Gamma$ , target expression $e^{\varepsilon}$ has target type $\tau$
$\Lambda(\ell) = \tau$ $\Lambda; \Gamma \vdash_{\mathbb{S}} \ell : \tau$	(TLoc)
$\Lambda; \Gamma \vdash_{\mathbb{S}} n : \mathbf{int}$	(TInt)
$\Gamma(x) = \tau$ $\Lambda; \Gamma \vdash_{\mathbb{S}} x : \tau$	(TPVar)
$\Gamma(x) = \Pi \vec{\alpha}[D]. \tau$ $\Lambda; \Gamma \vdash_{\mathbb{S}} x[\vec{\alpha} = \vec{\delta}] : \llbracket [\vec{\delta}/\vec{\alpha}] \tau \rrbracket$	(TVar)
$\Lambda; \Gamma \vdash_{\mathbb{S}} v_1 : \tau_1$ $\Lambda; \Gamma \vdash_{\mathbb{S}} v_2 : \tau_2$	(TPair)
$\Lambda; \Gamma, x : \tau_1, f : (\tau_1 \xrightarrow{\varepsilon} \tau_2) \vdash_{\varepsilon} e : \tau_2$	(TFun)
$\Lambda; \Gamma \vdash_{\mathbb{S}} v : \tau_1$	(TSum)
$\Lambda; \Gamma \vdash_{\mathbb{S}} \underline{x} : \tau_1 \times \tau_2$	(TFst)
$\Lambda; \Gamma \vdash_{\mathbb{S}} \underline{x}_1 : \mathbf{int}$ $\Lambda; \Gamma \vdash_{\mathbb{S}} \underline{x}_2 : \mathbf{int}$	$\vdash \oplus : \mathbf{int} \times \mathbf{int} \rightarrow \mathbf{int}$ (TPrim)
$\Lambda; \Gamma \vdash_{\varepsilon} e_1^{\mathbb{S}} : \sigma$ $\Lambda; \Gamma, x : \sigma \vdash_{\varepsilon} e_2 : \tau'$	$\Lambda; \Gamma \vdash_{\varepsilon} \mathbf{let} \ x = e_1^{\mathbb{S}} \ \mathbf{in} \ e_2 : \tau'$ (TLet)
$\Lambda; \Gamma \vdash_{\mathbb{S}} \underline{x}_1 : (\tau_1 \xrightarrow{\varepsilon} \tau_2)$ $\Lambda; \Gamma \vdash_{\mathbb{S}} \underline{x}_2 : \tau_1$	$\Lambda; \Gamma \vdash_{\varepsilon} \mathbf{apply}^{\varepsilon}(x_1, x_2) : \tau_2$ (TApp)
$\Lambda; \Gamma \vdash_{\mathbb{S}} \underline{x} : \tau_1 + \tau_2$ $\Lambda; \Gamma, x_1 : \tau_1 \vdash_{\varepsilon} e_1 : \tau$ $\Lambda; \Gamma, x_2 : \tau_2 \vdash_{\varepsilon} e_2 : \tau$	$\Lambda; \Gamma \vdash_{\varepsilon} \mathbf{case} \ \underline{x} \ \mathbf{of} \ \{x_1 \Rightarrow e_1, x_2 \Rightarrow e_2\} : \tau$ (TCase)
$\Lambda; \Gamma \vdash_{\mathbb{C}} e : \tau$	(TMod)
$\Lambda; \Gamma \vdash_{\mathbb{S}} \underline{x} : \tau$	$\Lambda; \Gamma \vdash_{\mathbb{C}} \mathbf{write}(\underline{x}) : \tau$ (TWrite)
$\Lambda; \Gamma \vdash_{\mathbb{S}} \underline{x}_1 : \tau_1 \mathbf{mod}$ $\Lambda; \Gamma, x : \tau_1 \vdash_{\mathbb{C}} e_2 : \tau_2$	$\Lambda; \Gamma \vdash_{\mathbb{C}} \mathbf{read} \ \underline{x}_1 \ \mathbf{as} \ x \ \mathbf{in} \ e_2 : \tau_2$ (TRead)

**Figure 12.** Typing rules of the target language AFL

polymorphic type with the levels  $\vec{\delta}$  substituted for the variables  $\vec{\alpha}$ , and  $\llbracket - \rrbracket$  translates source types to target types (see Section 6.1). Rule (TPVar) is a standard rule for variables of monomorphic type, but rule (TVar) gives the instantiation  $x[\vec{\alpha} = \vec{\delta}]$ , of a variable  $x$  of polymorphic type, the type  $\llbracket [\vec{\delta}/\vec{\alpha}] \tau \rrbracket$ —matching the monomorphic expression from the  $\mathbf{select}$  to which  $x$  is bound.

### 5.1 Dynamic Semantics

For the source language, our big-step evaluation rules (Figure 13) are standard. In the target language AFL, our rules (Figure 14) model the evaluation of a first run of the program: modifiabiles are created, written to (once), and read from (any number of times), but never updated to reflect changes to the program input. Both sets of rules permit expressions that are not in A-normal form, enabling

$$\boxed{e \Downarrow v} \text{ Source expression } e \text{ evaluates to } v$$

$$\frac{}{v \Downarrow v} \text{ (SEvValue)} \quad \frac{e_1 \Downarrow v_1 \quad e_2 \Downarrow v_2}{(e_1, e_2) \Downarrow (v_1, v_2)} \text{ (SEvPair)}$$

$$\frac{e \Downarrow v}{\mathbf{inl} \ e \Downarrow \mathbf{inl} \ v} \text{ (SEvSum)} \quad \frac{e_1 \Downarrow v_1 \quad e_2 \Downarrow v_2 \quad \oplus(v_1, v_2) = v'}{\oplus(e_1, e_2) \Downarrow v'} \text{ (SEvPrimop)}$$

$$\frac{e \Downarrow (v_1, v_2)}{\mathbf{fst} \ e \Downarrow v_1} \text{ (SEvFst)} \quad \frac{e_1 \Downarrow v_1 \quad [v_1/x]e_2 \Downarrow v_2}{\mathbf{let} \ x = e_1 \ \mathbf{in} \ e_2 \Downarrow v_2} \text{ (SEvLet)}$$

$$\frac{e \Downarrow \mathbf{inl} \ v_1 \quad [v_1/x_1]e_1 \Downarrow v}{\mathbf{case} \ e \ \mathbf{of} \ \{x_1 \Rightarrow e_1, x_2 \Rightarrow e_2\} \Downarrow v} \text{ (SEvCaseLeft)}$$

$$\frac{e_1 \Downarrow \mathbf{fun} \ f(x) = e \quad e_2 \Downarrow v_2 \quad [(\mathbf{fun} \ f(x) = e)/f][v_2/x]e \Downarrow v}{\mathbf{apply}(e_1, e_2) \Downarrow v} \text{ (SEvApply)}$$

**Figure 13.** Dynamic semantics of source Level ML programs

standard capture-avoiding substitution. To simplify the translation, we call instantiations  $x[\vec{\alpha} = \vec{\delta}]$  values, even though  $x[\vec{\alpha} = \vec{\delta}]$  does not evaluate to itself. So we distinguish *machine values*  $w$ —which do evaluate to themselves—from values  $v$ . The only difference is that machine values do not include  $x[\vec{\alpha} = \vec{\delta}]$ .

$$\begin{array}{l}
\text{Machine } w ::= n \mid x \mid \ell \mid (w, w) \mid \mathbf{inl} \ w \mid \mathbf{inr} \ w \mid \\
\text{values} \quad \quad \quad \mathbf{fun}^\varepsilon \ f(x) = e^\varepsilon \mid \mathbf{select} \ \{(\vec{\alpha}_i = \vec{\delta}_i) \Rightarrow e_i\}_i
\end{array}$$

## 6. Translation

We specify the translation from Level ML to the target language AFL by a set of a rules. Because AFL is a modal language that distinguishes stable and changeable expressions, with a corresponding type system (Section 5), the translation is also modal: the translation in the stable mode  $\xrightarrow{\mathbb{S}}$  produces a stable AFL expression  $e^{\mathbb{S}}$ , and the translation in the changeable mode  $\xrightarrow{\mathbb{C}}$  produces a changeable expression  $e^{\mathbb{C}}$ .

It is not enough to generate AFL expressions of the right syntactic form; they must also have the right type. To achieve this, the rules are type-directed: we translate a source expression  $e$  at type  $\tau$ . But we are transforming expressions from one language to another, where each language has its own type system; translating some  $e : \tau$  cannot produce some  $e' : \tau$ , but some  $e' : \tau'$  where  $\tau'$  is a target type that *corresponds* to  $\tau$ . To express this vital property, we need to translate types, as well as expressions. We developed the translation of expressions and types together (along with the proof that the property holds); the translation of types was instrumental in getting the translation of expressions right. To understand how to translate expressions, it is helpful to first understand how we translate types.

### 6.1 Translating Types

Figure 15 defines the translation of types via two mutually recursive functions from Level ML types to AFL types. The first function,  $\|\tau\|$ , tells us what type the target expression  $e^{\mathbb{S}}$  should have when we translate  $e$  in the stable mode,  $e : \tau \xrightarrow{\mathbb{S}} e^{\mathbb{S}}$ . We also use it to translate the types in the environment  $\Gamma$ . The second function,  $\|\tau\|^{\mathbb{C}}$ , makes sense in two related situations: translating the type  $\tau$  of an expression  $e$  in the changeable mode ( $e : \tau \xrightarrow{\mathbb{C}} e^{\mathbb{C}}$ ) and translating the codomain of changeable functions.

In the stable mode, values of stable type can be used and created directly, so the “stable” translation  $\|\mathbf{int}^{\mathbb{S}}\|$  of a stable integer is

$$\boxed{\rho \vdash e \Downarrow (\rho' \vdash w)} \text{ In the store } \rho, \text{ target expression } e \text{ evaluates to } w \text{ with updated store } \rho'$$

$$\frac{}{\rho \vdash w \Downarrow (\rho \vdash w)} \text{ (TEvMachineValue)}$$

$$\frac{\rho \vdash e_1 \Downarrow (\rho_1 \vdash w_1) \quad \rho_1 \vdash e_2 \Downarrow (\rho_2 \vdash w_2)}{\rho \vdash (e_1, e_2) \Downarrow (\rho_2 \vdash (w_1, w_2))} \text{ (TEvPair)}$$

$$\frac{\rho \vdash e \Downarrow (\rho' \vdash w)}{\rho \vdash \mathbf{inl} \ e \Downarrow (\rho' \vdash \mathbf{inl} \ w)} \text{ (TEvSum)}$$

$$\frac{\rho \vdash e_1 \Downarrow (\rho_1 \vdash w_1) \quad \rho_1 \vdash e_2 \Downarrow (\rho_2 \vdash w_2) \quad \oplus(w_1, w_2) = w'}{\rho \vdash \oplus(e_1, e_2) \Downarrow (\rho_2 \vdash w')} \text{ (TEvPrimop)}$$

$$\frac{\rho \vdash e \Downarrow (\rho' \vdash (w_1, w_2))}{\rho \vdash \mathbf{fst} \ e \Downarrow (\rho' \vdash w_1)} \text{ (TEvFst)}$$

$$\frac{\rho \vdash e_1 \Downarrow (\rho_1 \vdash w_1) \quad \rho_1 \vdash [w_1/x]e_2 \Downarrow (\rho_2 \vdash w_2)}{\rho \vdash \mathbf{let} \ x = e_1 \ \mathbf{in} \ e_2 \Downarrow (\rho_2 \vdash w_2)} \text{ (TEvLet)}$$

$$\frac{\rho \vdash e \Downarrow (\rho_1 \vdash \mathbf{inl} \ w_1) \quad \rho_1 \vdash [w_1/x_1]e_1 \Downarrow (\rho_2 \vdash w)}{\rho \vdash \mathbf{case} \ e \ \mathbf{of} \ \{x_1 \Rightarrow e_1, x_2 \Rightarrow e_2\} \Downarrow (\rho_2 \vdash w)} \text{ (TEvCaseLeft)}$$

$$\frac{\rho \vdash e_1^{\mathbb{S}} \Downarrow (\rho_1 \vdash \mathbf{fun}^\varepsilon \ f(x) = e^\varepsilon) \quad \rho_1 \vdash e_2^{\mathbb{S}} \Downarrow (\rho_2 \vdash w_2) \quad \rho_2 \vdash [(\mathbf{fun}^\varepsilon \ f(x) = e^\varepsilon)/f][w_2/x]e^{\mathbb{S}} \Downarrow (\rho_3 \vdash w)}{\rho \vdash \mathbf{apply}^\varepsilon(e_1^{\mathbb{S}}, e_2^{\mathbb{S}}) \Downarrow (\rho_3 \vdash w)} \text{ (TEvApply)}$$

$$\frac{\rho \vdash e \Downarrow (\rho' \vdash w)}{\rho \vdash \mathbf{write}(e) \Downarrow (\rho' \vdash w)} \text{ (TEvWrite)}$$

$$\frac{\rho \vdash e_1 \Downarrow (\rho_1 \vdash \ell) \quad \rho_1 \vdash [\rho_1(\ell)/x']e^{\mathbb{C}} \Downarrow (\rho_2 \vdash w)}{\rho \vdash \mathbf{read} \ e_1 \ \mathbf{as} \ x' \ \mathbf{in} \ e^{\mathbb{C}} \Downarrow (\rho_2 \vdash w)} \text{ (TEvRead)}$$

$$\frac{\rho_1 \vdash e_1 \Downarrow (\rho' \vdash w)}{\rho \vdash (\mathbf{select} \ \{\dots, \vec{\delta} \Rightarrow e_1, \dots\})[\vec{\alpha} = \vec{\delta}] \Downarrow (\rho' \vdash w)} \text{ (TEvSelectE)}$$

$$\frac{\rho \vdash e^{\mathbb{C}} \Downarrow (\rho' \vdash w)}{\rho \vdash \mathbf{mod} \ e^{\mathbb{C}} \Downarrow ((\rho', \ell \mapsto w) \vdash \ell)} \text{ (TEvMod)}$$

**Figure 14.** Dynamic semantics for first runs of AFL programs

just **int**. In contrast, a changeable integer cannot be inspected or directly created in stable mode, but must be placed into a modifiable:  $\|\mathbf{int}^{\mathbb{C}}\| = \mathbf{int} \ \mathbf{mod}$ . The remaining parts of the definition follow this pattern: the target type is wrapped with **mod** if and only if the outer level of the source type is  $\mathbb{C}$ . When we translate a changeable-mode function type (with  $\mathbb{C}$  below the arrow), its codomain is translated “output-changeable”:  $\|(\tau_1 \xrightarrow{\mathbb{C}} \tau_2)^{\mathbb{S}}\| = \|\tau_1\| \xrightarrow{\mathbb{C}} \|\tau_2\|^{\mathbb{C}}$ . The reason is that a changeable-mode function can only be applied in the changeable mode; the function result is not placed into a modifiable until we return to the stable mode, so putting a **mod** on the codomain would not match the dynamic semantics of AFL.

The second function  $\|\tau\|^{\mathbb{C}}$  defines the type of a changeable expression  $e$  that writes to a modifiable containing  $\tau$ , yielding a changeable target expression  $e^{\mathbb{C}}$ . The source type has an outer  $\mathbb{C}$ , so when the value is written, it will be placed into a modifiable and have **mod** type. But while evaluating  $e^{\mathbb{C}}$ , there is no outer **mod**. Thus the translation  $\|\tau\|^{\mathbb{C}}$  ignores the outer level (using the function  $\|-^{\mathbb{S}}\|$ , which replaces an outer level  $\mathbb{C}$  with  $\mathbb{S}$ ), and never



$$\begin{aligned}
|\mathbf{int}^{\mathbf{C}}|^{\mathbf{S}} &= \mathbf{int}^{\mathbf{S}} \\
|(\tau_1 \times \tau_2)^{\mathbf{C}}|^{\mathbf{S}} &= (\tau_1 \times \tau_2)^{\mathbf{S}} \\
|(\tau_1 + \tau_2)^{\mathbf{C}}|^{\mathbf{S}} &= (\tau_1 + \tau_2)^{\mathbf{S}} \\
|(\tau_1 \xrightarrow{\varepsilon} \tau_2)^{\mathbf{C}}|^{\mathbf{S}} &= (\tau_1 \xrightarrow{\varepsilon} \tau_2)^{\mathbf{S}} \\
\|\mathbf{int}^{\mathbf{C}}\| &= \mathbf{int} \\
\|\mathbf{int}^{\mathbf{C}}\| &= \mathbf{int} \mathbf{mod} \\
\|(\tau_1 \xrightarrow{\mathcal{S}} \tau_2)^{\mathbf{S}}\| &= \|\tau_1\| \xrightarrow{\mathcal{S}} \|\tau_2\| \\
\|(\tau_1 \xrightarrow{\mathcal{S}} \tau_2)^{\mathbf{C}}\| &= \left( \|\tau_1\| \xrightarrow{\mathcal{S}} \|\tau_2\| \right) \mathbf{mod} \\
\|(\tau_1 \xrightarrow{\mathcal{C}} \tau_2)^{\mathbf{S}}\| &= \|\tau_1\| \xrightarrow{\mathcal{C}} \|\tau_2\|^{\mathbf{C}} \\
\|(\tau_1 \xrightarrow{\mathcal{C}} \tau_2)^{\mathbf{C}}\| &= \left( \|\tau_1\| \xrightarrow{\mathcal{C}} \|\tau_2\|^{\mathbf{C}} \right) \mathbf{mod} \\
\|(\tau_1 \times \tau_2)^{\mathbf{S}}\| &= \|\tau_1\| \times \|\tau_2\| \\
\|(\tau_1 \times \tau_2)^{\mathbf{C}}\| &= \left( \|\tau_1\| \times \|\tau_2\| \right) \mathbf{mod} \\
\|(\tau_1 + \tau_2)^{\mathbf{S}}\| &= \|\tau_1\| + \|\tau_2\| \\
\|(\tau_1 + \tau_2)^{\mathbf{C}}\| &= \left( \|\tau_1\| + \|\tau_2\| \right) \mathbf{mod} \\
\|\tau\|^{\mathbf{C}} &= \begin{cases} \|\tau\|^{\mathbf{S}} & \text{if } \tau \text{ O.C.} \\ \|\tau\| & \text{if } \tau \text{ O.S.} \end{cases} \\
\|\cdot\| &= \cdot & \|\tau\|_{\phi} &= \|\phi\| \tau \\
\|\Gamma, x : \forall \emptyset[\mathbf{true}]. \tau\| &= \|\Gamma\|, x : \|\tau\| & \|\tau\|_{\phi}^{\mathbf{C}} &= \|\phi\| \tau^{\mathbf{C}} \\
\|\Gamma, x : \forall \bar{\alpha}[D]. \tau\| &= \|\Gamma\|, x : \Pi \bar{\alpha}[D]. \tau & \|\Gamma\|_{\phi} &= \|\phi\| \Gamma
\end{aligned}$$

**Figure 15.** Stabilization of types  $|\tau|^{\mathbf{S}}$ ; translations  $\|\tau\|$  and  $\|\tau\|^{\mathbf{C}}$  of types; translation of typing environments  $\|\Gamma\|$

returns a type of the form  $(\dots \mathbf{mod})$ . However, since the value being returned may contain subexpressions that will be placed into modifiabls, we use  $\|\cdot\|$  for the inner types. For instance,  $\|(\tau_1 + \tau_2)^{\mathbf{S}}\|^{\mathbf{C}} = \|\tau_1\| + \|\tau_2\|$ .

These functions are defined on closed types—types with no free level variables. Before applying one of these functions to a type found by the constraint typing rules, we always need to apply the satisfying assignment  $\phi$  to the type, so for convenience we write  $\|\tau\|_{\phi}$  for  $\|\phi\| \tau$ , and so on. Because the translation only makes sense for closed types, type schemes  $\forall \bar{\alpha}[D]. \tau$  cannot be translated. The translation  $\|\Gamma\|$  therefore translates only monomorphic types  $\tau$ ; type schemes are left alone (except for replacing the symbol  $\forall$  with  $\Pi$ ) until instantiation. Once instantiated, the type scheme is an ordinary closed source type, and can be translated by rule (TVar).

## 6.2 Translating Expressions

We define the translation of expressions as a set of type-directed rules. Given (1) a derivation of  $C; \Gamma \vdash_e e : \tau$  in the constraint-based typing system and (2) a satisfying assignment  $\phi$  for  $C$ , it is always possible to produce a correctly typed stable target expression  $e^{\mathbf{S}}$  and a correctly typed changeable target expression  $e^{\mathbf{C}}$  (see Theorem 6.1 below). The environment  $\Gamma$  in the translation rules is a source typing environment, but must have no free level variables. Given an environment  $\Gamma$  from the constraint typing, we apply the satisfying assignment  $\phi$  to eliminate its free level variables before using it for the translation:  $[\phi]\Gamma$ . With the environment closed, we need not refer to  $C$ .

Many of the rules in Figure 17 are purely syntax-directed and are similar to the constraint-based rules. One exception is the (Var) rule, which needs the source type to know how to instantiate the level variables in the type scheme. For example, given the polymorphic  $x : \forall \alpha[\mathbf{true}]. (\mathbf{int}^{\alpha} \xrightarrow{\alpha} \mathbf{int}^{\alpha})^{\mathbf{S}}$ , we need the type from  $C; \Gamma \vdash_e x : (\mathbf{int}^{\mathbf{C}} \xrightarrow{\mathbf{C}} \mathbf{int}^{\mathbf{C}})^{\mathbf{S}}$  so we can instantiate  $\alpha$  in the translated term  $x[\alpha = \mathbf{C}]$ .

$$\begin{array}{c}
\boxed{\Gamma \vdash e \rightsquigarrow (x \gg x' : \tau \vdash e')} \quad \text{Under source typing } \Gamma, \\
\text{renaming the “head” } x \text{ in } e \\
\text{to } x' : \tau \text{ yields expression } e' \\
\hline
\frac{\Gamma \vdash x_1 : \tau}{\Gamma \vdash \oplus(x_1, x_2) \rightsquigarrow (x_1 \gg x'_1 : \tau \vdash \oplus(x'_1, x_2))} \text{ (LPrimop1)} \\
\frac{\Gamma \vdash x_2 : \tau}{\Gamma \vdash \oplus(x_1, x_2) \rightsquigarrow (x_2 \gg x'_2 : \tau \vdash \oplus(x_1, x'_2))} \text{ (LPrimop2)} \\
\frac{\Gamma \vdash x : \tau}{\Gamma \vdash \mathbf{fst} x \rightsquigarrow (x \gg x' : \tau \vdash \mathbf{fst} x')} \text{ (LFst)} \\
\frac{\Gamma \vdash x_1 : \tau}{\Gamma \vdash \mathbf{apply}(x_1, x_2) \rightsquigarrow (x_1 \gg x' : \tau \vdash \mathbf{apply}(x', x_2))} \text{ (LApply)} \\
\frac{\Gamma \vdash x : \tau}{\Gamma \vdash \mathbf{case} x \mathbf{of} \{x_1 \Rightarrow e_1, x_2 \Rightarrow e_2\} \rightsquigarrow (x \gg x' : \tau \vdash \mathbf{case} x' \mathbf{of} \{x_1 \Rightarrow e_1, x_2 \Rightarrow e_2\})} \text{ (LCase)}
\end{array}$$

**Figure 16.** Renaming the variable to be read (elimination forms)

Our rules are nondeterministic, avoiding the need to “decorate” them with context-sensitive details. Our algorithm in Section 6.3 resolves the nondeterminism through type information.

**Stable rules.** The rules (Int), (Var), (Pair), (Fun), (Sum), (Fst) and (Prim) can only translate in the stable mode. To translate to a changeable expression, use a rule that shifts to changeable mode.

**Shifting to changeable mode.** Given a translation of  $e$  in the stable mode to some  $e^{\mathbf{S}}$ , the rules (Write) and (ReadWrite) at the bottom of Figure 17 translate  $e$  in the changeable mode, producing an  $e^{\mathbf{C}}$ . If the expression’s type  $\tau$  is outer stable (say,  $\mathbf{int}^{\mathbf{S}}$ ), the (Write) rule simply binds it to a variable and then writes that variable. If  $\tau$  is outer changeable (say,  $\mathbf{int}^{\mathbf{C}}$ ) it will be in a modifiable at runtime, so we read it into  $r'$  and then write it. (The let-bindings merely satisfy the requirements of A-normal form.)

**Shifting to stable mode.** To generate a stable expression  $e^{\mathbf{S}}$  based on a changeable expression  $e^{\mathbf{C}}$ , we have the (Lift) and (Mod) rules. These rules require the source type  $\tau$  to be outer changeable: in (Lift), the premise  $|\tau|^{\mathbf{S}} = \tau'$  requires that  $|\tau|^{\mathbf{S}}$  is defined, and it is defined only for outer changeable  $\tau$ ; in (Mod), the requirement is explicit:  $\vdash \tau$  O.C.

(Mod) is the simpler of the two: if  $e$  translates to  $e^{\mathbf{C}}$  at type  $\tau$ , then  $e$  translates to the stable expression  $\mathbf{mod} e^{\mathbf{C}}$  at type  $\tau$ . In (Lift), the expression is translated not at the given type  $\tau$  but at its *stabilized*  $|\tau|^{\mathbf{S}}$ , capturing the “shallow subsumption” in the constraint typing rules (SLetE) and (SLetV): a bound expression of type  $\tau_0^{\mathbf{S}}$  can be translated at type  $\tau_0^{\mathbf{S}}$  to  $e^{\mathbf{S}}$ , and then “promoted” to type  $\tau_0^{\mathbf{C}}$  by placing it inside a **mod**.

**Reading from changeable data.** To use an expression of changeable type in a context where a stable value is needed—such as passing some  $x : \mathbf{int}^{\mathbf{C}}$  to a function expecting  $\mathbf{int}^{\mathbf{S}}$ —the (Read) rule generates a target expression that reads the value out of  $x : \mathbf{int}^{\mathbf{C}}$  into a variable  $x' : \mathbf{int}^{\mathbf{S}}$ . The variable-renaming judgment  $\Gamma \vdash e \rightsquigarrow (x \gg x' : \tau \vdash e')$  takes the expression  $e$ , finds a variable  $x$  about to be used, and yields an expression  $e'$  with that occurrence replaced by  $x'$ . For example,  $\Gamma \vdash \mathbf{case} x \mathbf{of} \dots \rightsquigarrow (x \gg x' : \tau \vdash \mathbf{case} x' \mathbf{of} \dots)$ . This judgment is derivable only for **apply**, **case**, **fst**, and  $\oplus$ , because these are the elimination forms for outer-changeable data. For  $\oplus(x_1, x_2)$ , we need to read both variables, so we have one rule for each. The rules are given in Figure 16.

$\Gamma \vdash e : \tau \xrightarrow{\varepsilon} e^\varepsilon$

Under closed source typing environment  $\Gamma$ , source expression  $e$  is translated at type  $\tau$  in mode  $\varepsilon$  to target expression  $e^\varepsilon$

$$\frac{}{\Gamma \vdash n : \mathbf{int}^\delta \xrightarrow{\mathbb{S}} n} \text{(Int)} \quad \frac{\Gamma(x) = \forall \vec{\alpha}[D]. \tau}{\Gamma \vdash x : [\vec{\delta}/\vec{\alpha}]\tau \xrightarrow{\mathbb{S}} x[\vec{\alpha} = \vec{\delta}]} \text{(Var)}$$

$$\frac{\Gamma \vdash v_1 : \tau_1 \xrightarrow{\mathbb{S}} v'_1 \quad \Gamma \vdash v_2 : \tau_2 \xrightarrow{\mathbb{S}} v'_2}{\Gamma \vdash (v_1, v_2) : (\tau_1 \times \tau_2)^\mathbb{S} \xrightarrow{\mathbb{S}} (v'_1, v'_2)} \text{(Pair)}$$

$$\frac{\Gamma, x : \tau_1, f : (\tau_1 \xrightarrow{\varepsilon} \tau_2)^\mathbb{S} \vdash e : \tau_2 \xrightarrow{\varepsilon} e^\varepsilon}{\Gamma \vdash \mathbf{fun} f(x) = e : (\tau_1 \xrightarrow{\varepsilon} \tau_2)^\mathbb{S} \xrightarrow{\mathbb{S}} \mathbf{fun}^\varepsilon f(x) = e^\varepsilon} \text{(Fun)}$$

$$\frac{\Gamma \vdash v : \tau_1 \xrightarrow{\mathbb{S}} v'}{\Gamma \vdash \mathbf{inl} v : (\tau_1 + \tau_2)^\mathbb{S} \xrightarrow{\mathbb{S}} \mathbf{inl} v'} \text{(Sum)}$$

$$\frac{\Gamma \vdash x : (\tau_1 \times \tau_2)^\mathbb{S} \xrightarrow{\mathbb{S}} \underline{x}}{\Gamma \vdash \mathbf{fst} x : \tau_1 \xrightarrow{\mathbb{S}} \mathbf{fst} \underline{x}} \text{(Fst)}$$

$$\frac{\Gamma \vdash x_1 : \mathbf{int}^\mathbb{S} \xrightarrow{\mathbb{S}} \underline{x_1} \quad \Gamma \vdash x_2 : \mathbf{int}^\mathbb{S} \xrightarrow{\mathbb{S}} \underline{x_2}}{\Gamma \vdash \oplus(x_1, x_2) : \mathbf{int}^\delta \xrightarrow{\mathbb{S}} \oplus(\underline{x_1}, \underline{x_2})} \text{(Prim)}$$

$$\frac{\Gamma \vdash x_1 : (\tau_1 \xrightarrow{\varepsilon} \tau_2)^\mathbb{S} \xrightarrow{\mathbb{S}} \underline{x_1} \quad \Gamma \vdash x_2 : \tau_1 \xrightarrow{\mathbb{S}} \underline{x_2}}{\Gamma \vdash \mathbf{apply}(x_1, x_2) : \tau_2 \xrightarrow{\varepsilon} \mathbf{apply}^\varepsilon(\underline{x_1}, \underline{x_2})} \text{(App)}$$

$$\frac{\Gamma \vdash x : (\tau_1 + \tau_2)^\mathbb{S} \xrightarrow{\mathbb{S}} \underline{x} \quad \begin{array}{l} \Gamma, x_1 : \tau_1 \vdash e_1 : \tau \xrightarrow{\varepsilon} e'_1 \\ \Gamma, x_2 : \tau_2 \vdash e_2 : \tau \xrightarrow{\varepsilon} e'_2 \end{array}}{\Gamma \vdash \mathbf{case} x \mathbf{of} \{x_1 \Rightarrow e_1, x_2 \Rightarrow e_2\} : \tau \xrightarrow{\varepsilon} \mathbf{case} x \mathbf{of} \{x_1 \Rightarrow e'_1, x_2 \Rightarrow e'_2\}} \text{(Case)}$$

$$\frac{\Gamma \vdash e : \tau \xrightarrow{\mathbb{C}} e^C \quad |\tau|^\mathbb{S} = \tau'}{\Gamma \vdash e : \tau \xrightarrow{\mathbb{S}} \mathbf{mod} e^C} \text{(Lift)} \quad \frac{\Gamma \vdash e : \tau \xrightarrow{\mathbb{C}} e^C \quad \tau \text{ O.C.}}{\Gamma \vdash e : \tau \xrightarrow{\mathbb{S}} \mathbf{mod} e^C} \text{(Mod)}$$

$$\frac{\Gamma \vdash e_1 : \tau' \xrightarrow{\mathbb{S}} e^\mathbb{S} \quad \Gamma, x : \tau' \vdash e_2 : \tau \xrightarrow{\varepsilon} e'_2}{\Gamma \vdash \mathbf{let} x = e_1 \mathbf{in} e_2 : \tau \xrightarrow{\varepsilon} \mathbf{let} x = e^\mathbb{S} \mathbf{in} e'_2} \text{(LetE)}$$

$$\frac{\Gamma, x : \forall \vec{\alpha}[D]. \tau' \vdash e : \tau \xrightarrow{\varepsilon} e' \quad \begin{array}{l} \text{For all } \vec{\delta}_i \text{ s.t. } \vec{\alpha} = \vec{\delta}_i \Vdash D, \\ \Gamma \vdash v : [\vec{\delta}_i/\vec{\alpha}]\tau' \xrightarrow{\mathbb{S}} e'_i \end{array}}{\Gamma \vdash \mathbf{let} x = v \mathbf{in} e : \tau \xrightarrow{\varepsilon} \mathbf{let} x = \mathbf{select} \{(\vec{\alpha} = \vec{\delta}_i) \Rightarrow e'_i\}_i \mathbf{in} e'} \text{(LetV)}$$

$$\frac{\begin{array}{l} \Gamma \vdash e \rightsquigarrow (x \gg x' : \tau' \vdash e') \quad \tau' \text{ O.C.} \\ \Gamma, x' : |\tau'|^\mathbb{S} \vdash e' : \tau \xrightarrow{\mathbb{C}} e^C \quad \Gamma \vdash x : \tau' \xrightarrow{\mathbb{S}} \underline{x} \end{array}}{\Gamma \vdash e : \tau \xrightarrow{\mathbb{C}} \mathbf{read} \underline{x} \mathbf{as} x' \mathbf{in} e^C} \text{(Read)}$$

$$\frac{\Gamma \vdash e : \tau \xrightarrow{\mathbb{S}} e^\mathbb{S} \quad \tau \text{ O.S.}}{\Gamma \vdash e : \tau \xrightarrow{\mathbb{C}} \mathbf{let} r = e^\mathbb{S} \mathbf{in} \mathbf{write}(r)} \text{(Write)}$$

$$\frac{\Gamma \vdash e : \tau \xrightarrow{\mathbb{S}} e^\mathbb{S} \quad \tau \text{ O.C.}}{\Gamma \vdash e : \tau \xrightarrow{\mathbb{C}} \mathbf{let} r = e^\mathbb{S} \mathbf{in} \mathbf{read} r \mathbf{as} r' \mathbf{in} \mathbf{write}(r')} \text{(ReadWrite)}$$

**Figure 17.** Monomorphizing translation

**Monomorphization.** A polymorphic source expression has no directly corresponding target expression: the `map` function from Section 2 corresponds to the two functions `map_SC` and `map_CS`. Given a polymorphic source value  $v : \forall \vec{\alpha}[D]. \tau'$ , the (LetV) rule translates  $v$  once for each instantiation  $\vec{\delta}_i$  that satisfies the constraint  $D$  (each  $\vec{\delta}_i$  such that  $\vec{\alpha} = \vec{\delta}_i \Vdash D$ ). That is, we translate the value at source type  $[\vec{\delta}_i/\vec{\alpha}]\tau'$ . This yields a sequence of source expressions  $e_1, \dots, e_n$  for the  $n$  possible instances. For example, given  $\forall \alpha[\text{true}]. \tau'$ , we translate the value at type  $[\mathbb{S}/\alpha]\tau'$  yielding  $e_1$  and at type  $[\mathbb{C}/\alpha]\tau'$  yielding  $e_2$ . Finally, the rule produces a **select** expression, which acts as a function that takes the desired instance  $\vec{\delta}_i$  and returns the appropriate  $e_i$ .

Since (LetV) generates one function for each satisfying  $\vec{\delta}_i$ , it can create up to  $2^n$  instances for  $n$  variables. However, dead-code elimination can remove functions that are not used. Moreover, the functions that *are* used would have been handwritten in an explicit setting, so while the code size is exponential in the worst case, the saved effort is as well.

### 6.3 Algorithm

The system of translation rules in Figure 17 is not deterministic. In fact, if the wrong choices are made it can produce painfully inefficient code. Suppose we have  $2 : \mathbf{int}^C$ , and want to translate it to a stable target expression. Choosing rule (Int) yields the target expression 2. But we could use (Int), then (ReadWrite)—which generates an  $e^C$  with a **let**, a **read** and a **write**—then (Mod), which wraps that  $e^C$  in a **mod**. Clearly, we should have stopped with (Int).

To resolve this nondeterminism in the rules would complicate them further. Instead, we give the algorithm in Figure 18, which examines the source expression  $e$  and, using type information, applies the rules necessary to produce an expression of mode  $\varepsilon$ .

### 6.4 Properties

Given a constraint-based source typing derivation and assignment  $\phi$  for some term  $e$ , there are translations from  $e$  to (1) a stable  $e^\mathbb{S}$  and (2) a changeable  $e^C$ , with appropriate target types:

**Theorem 6.1** (Translation Type Soundness).

If  $C; \Gamma \vdash_\varepsilon e : \tau$  and  $\phi$  is a satisfying assignment for  $C$  then

- (1) there exists  $e^\mathbb{S}$  such that  $[\phi]\Gamma \vdash e : [\phi]\tau \xrightarrow{\mathbb{S}} e^\mathbb{S}$  and  $;\|\Gamma\|_\phi \vdash_\mathbb{S} e^\mathbb{S} : \|\tau\|_\phi$  and, if  $e$  is a value, then  $e^\mathbb{S}$  is a value;
- (2) there exists  $e^C$  such that  $[\phi]\Gamma \vdash e : [\phi]\tau \xrightarrow{\mathbb{C}} e^C$  and  $;\|\Gamma\|_\phi \vdash_\mathbb{C} e^C : \|\tau\|_\phi^C$ .

The proof (in the appendix [Chen et al. 2011]) is by induction on the height of the given derivation of  $C; \Gamma \vdash_\varepsilon e : \tau$ . If the concluding rule was (SLetE), we use a substitution property (Lemma A.2) for each  $\vec{\delta}_i$  to get a monomorphic constraint typing derivation; that derivation is not larger than the input derivation, so we can apply the induction hypothesis to get a translated  $e'_i$ . The proof constructs the same translation derivations as the algorithm in Figure 18 (in fact, we extracted the algorithm from the proof).

We also prove that running a translated program gives the same result as running the source program. Theorem 6.5 states that in an initially empty store  $\cdot$ , if evaluating the translated program  $e'$  yields  $v'$  with new store  $\rho'$ , then  $e$  evaluates to  $v$  where  $v$  corresponds to  $[\rho']v'$  (the result of substituting values in the store  $\rho'$  for locations appearing in  $v'$ ).

To define this correspondence, we use a device somewhat similar to logical relations: a relation  $\rightsquigarrow$  on source and target expressions, allowing us to show that if  $e : \tau \rightsquigarrow e' : \tau'$  then  $v : \tau \rightsquigarrow [\rho]v' : \tau'$ . Both  $e$  and  $e'$  must be closed. Our definition is weaker than the equivalence relations used in logical relations proofs: **apply**(id, 4)  $\not\rightsquigarrow$  4, for example. It does not attempt to equate all programs that have the same meaning, but only particu-

```

function trans (e, ε) = case (e, ε) of
| (n, S) ⇒ Int
| (x, S) ⇒ Var
| ((v1, v2), S) ⇒ Pair(trans(v1, S), trans(v2, S))
| (fun f(x) = e' : (τ1  $\xrightarrow{\varepsilon}$  τ2)S, S) ⇒ Fun(trans(e', ε'))
| (inl v, S) ⇒ Sum(trans(v, S))
| (fst (x : (τ1 × τ2)δ, ε) ⇒ case (δ, ε) of
| (S, S) ⇒ Fst(trans(x, S))
| (S, C) ⇒ if τ1 O.S. then Write(trans(e, S))
else ReadWrite(trans(e, S))
| (C, C) ⇒ Read(LFst, trans(fst (x' : (τ1 × τ2)S), C),
trans(x, S))
| (⊕(x1 : intS, x2 : intS), S) ⇒
Prim(trans(x1, S), trans(x2, S))
| (⊕(x1 : intS, x2 : intS), C) ⇒ Write(trans(e, S))
| (⊕(x1 : intC, x2 : intC), C) ⇒
Read(LPrimop1,
Read(LPrimop2, Write(trans(⊕(x'1, x'2), S))))
| (let x : τ'' = e1 : τ' in e2, ε) ⇒
LetE(if τ'' O.S. then trans(e1, S)
else (if τ' = τ'' then Mod(trans(e1, C))
else Lift(trans(e1, C))),
trans(e2, ε))
| (let x : ∀α[D]. τ'' = v1 : τ' in e2, ε) ⇒
let variants = all δi such that α = δi ⊢ D in
let f = λs. if τ'' O.S. then trans(v1, S)
else (if τ' = τ'' then Mod(trans(v1, C))
else Lift(trans(v1, C))) in
LetV(map f variants, trans(e2, ε))
| (apply(x1 : (τ1  $\xrightarrow{\varepsilon}$  τ2)δ, x2), ε) ⇒ case (ε', δ, ε) of
| (S, S, S) ⇒ App(trans(x1, S), trans(x2, S))
| (C, S, C) ⇒ App(trans(x1, S), trans(x2, S))
| (S, S, C) ⇒ if τ2 O.S. then Write(trans(e, S))
else ReadWrite(trans(e, S))
| (ε', C, C) ⇒ Read(LApply,
trans(apply(x' : (τ1  $\xrightarrow{\varepsilon}$  τ2)S, x2), C),
trans(x1, S))
| (C, S, S) ⇒ Mod(trans(e, C))
| (ε', C, S) ⇒ Mod(trans(e, C))
| (case x : τ of {x1 ⇒ e1, x2 ⇒ e2}, ε) ⇒
if τ O.S. then
Case(trans(x, S), trans(e1, ε), trans(e2, ε))
else Read(LCase,
trans(case x' : τS of {x1 ⇒ e1, x2 ⇒ e2}, C),
trans(x, S))
| (x : τ, C) ⇒ if τ O.S. then Write(trans(e, S))
else ReadWrite(trans(e, S))
| (fun f(x) = e', C) | (inl v, C)
| (n, C) | ((v1, v2), C) ⇒ Write(trans(e, S))

```

Figure 18. Translation algorithm

lar Level ML terms to AFL terms that are similarly structured, but have overhead (**mod**, **write**, etc.). Thus, integers are related to integers, pairs are related if their components are related, and so forth. The definition essentially ignores **mod** and **write** and ignores the mode in **apply**<sup>ε</sup>. Since translated programs can have “extra” **read** and **let** expressions, these are “substituted out” in the relation, so that  $3 \mapsto \mathbf{let} \ x = 3 \ \mathbf{in} \ x$ . Functions are related if, given related arguments, they produce related results. Note that we will not induct over this relation; neither the term, nor the type, gets smaller.

We also relate substitutions: a source substitution  $s$  and a target substitution  $\underline{s}$ ,

$$s = v_1/x_1, \dots, v_n/x_n$$

$$\underline{s} = w_1/x_1, \dots, w_n/x_n$$

are related at their environments  $\Gamma = x_1 : \tau_1, \dots, x_n : \tau_n$  and  $\Gamma' = x_1 : \tau'_1, \dots, x_n : \tau'_n$ , written

$$(v_1/x_1, \dots, v_n/x_n) : (x_1 : \tau_1, \dots, x_n : \tau_n)$$

$$\mapsto (w_1/x_1, \dots, w_n/x_n) : (x_1 : \tau'_1, \dots, x_n : \tau'_n)$$

if, for all  $k$  from 1 to  $n$ , we have  $v_k : \tau_k \mapsto w_k : \tau'_k$ .

The key lemma (Lemma 6.3) is that if  $e \xrightarrow{\varepsilon} e'$ , the target program  $e'$  is related to  $e$ . Combined with Theorem 6.4, which shows that related programs evaluate to related values, this means that the translated program  $e'$  evaluates to the same value that  $e$  does. (Actually,  $e'$  is a value *related to* that value; only at **int**<sup>δ</sup>/**int**, and products thereof, are they identical.)

We begin by defining a *store substitution* operation:

**Definition 6.2.** The *store substitution*  $[\rho]e$  is defined as an ordinary substitution, except for  $e = \ell$ , in which case  $[\rho]\ell = [\rho](\mathbf{mod} \ \rho(\ell))$ .

For example,  $[\ell_1 \mapsto 1, \ell_2 \mapsto 2](\ell_1, \ell_2) = (\mathbf{mod} \ 1, \mathbf{mod} \ 2)$ .

Proofs and several other lemmas can be found in the appendix [Chen et al. 2011].

**Lemma 6.3** (Relation of Translation). *If  $\Gamma \vdash e : \tau \xrightarrow{\varepsilon} e'$  and  $\cdot \vdash s : \Gamma$  and  $\cdot \vdash \underline{s} : \|\Gamma\|$  and  $s : \Gamma \mapsto \underline{s} : \|\Gamma\|$  then*

$$[s]e : \tau \mapsto [s]e' : \tau'$$

where  $\tau' = \|\tau\|$  if  $\varepsilon = S$ , and  $\tau' = \|\tau\|^C$  if  $\varepsilon = C$ .

**Theorem 6.4** (Generalized Translation Soundness).

*If  $e : \sigma \mapsto [\rho]e' : \sigma'$  and  $\mathcal{D} :: \rho \vdash e \Downarrow (\rho' \vdash w)$  then  $\mathcal{D}' :: e \Downarrow v$  where  $v : \sigma \mapsto [\rho']w : \sigma'$ .*

Translation soundness now follows from Lemma 6.3 and Theorem 6.4:

**Theorem 6.5** (Translation Soundness). *If  $\cdot \vdash e : \tau \xrightarrow{\varepsilon} e'$  and  $\cdot \vdash e' \Downarrow (\rho' \vdash w)$ , then  $e \Downarrow v$  where  $v : \tau \mapsto [\rho']w : \tau'$ .*

Finally, we extend Theorem 6.5 to further show that the size  $W(\mathcal{D})$  of the derivation of the target-language evaluation is within a constant factor of the size  $W(\mathcal{D}')$  of the derivation of  $e \Downarrow v$ . We need a few definitions and intermediate results, which can be found in the appendix. The proof hinges on classifying the keywords added by the translation, such as **write**, as “dirty”: a dirty keyword will lead to applications of the dirty rule (TEvWrite) in the evaluation derivation; such applications have no equivalent in the source-language evaluation.

We then define the “head cost”  $HC$  of terms and derivations, which counts the number of dirty rules applied near the root of the term, or the root of the derivation, without passing through clean parts of the term or derivation. Just counting *all* the dirty keywords in a term would not rule out a  $\beta$ -reduction duplicating a particularly dirty part of the term. By defining head cost and proving that the translation generates terms with bounded head cost—including for all subterms—we ensure that *no* part of the term is too dirty; consequently, substituting a subterm during evaluation yields terms that are not too dirty.

**Definition 6.6.** A term  $e$  is shallowly  $k$ -bounded if  $HC(e) \leq k$ . A term  $e$  is deeply  $k$ -bounded if every subterm of  $e$  (including  $e$  itself) is shallowly  $k$ -bounded. Similarly, a derivation  $\mathcal{D}$  is shallowly  $k$ -bounded if  $HC(\mathcal{D}) \leq k$ , and deeply  $k$ -bounded if all its subderivations are shallowly  $k$ -bounded.

**Theorem 6.7.** *If  $\mathbf{trans} \ (e, \varepsilon) = e'$  then  $e'$  is deeply 6-bounded.*

$e : \sigma \mapsto e' : \underline{\sigma}'$ Source expression $e$ at type [schema] $\sigma$ is related to target expression $e'$ at type [schema] $\underline{\sigma}'$	
$n : \mathbf{int}^\delta \mapsto n : \mathbf{int}$ $(e_1, e_2) : (\tau_1 \times \tau_2)^\delta \mapsto (e'_1, e'_2) : \tau'_1 \times \tau'_2$ $(\mathbf{inl} e) : (\tau_1 + \tau_2)^\delta \mapsto (\mathbf{inl} e') : \tau'_1 + \tau'_2$ $(\mathbf{fst} e) : \tau \mapsto (\mathbf{fst} e') : \tau'$ $\oplus(e_1, e_2) : \tau \mapsto \oplus(e'_1, e'_2) : \tau'$ $e : \tau \mapsto (\mathbf{mod} e^C) : \tau' \mathbf{mod}$ $e : \tau \mapsto (\mathbf{write}(e^S)) : \tau'$ $[e_1/x]e : \tau_2 \mapsto (\mathbf{read} e^S \mathbf{as} x \mathbf{in} e^C) : \tau'_2$ $\mathbf{apply}(e_1, e_2) : \tau_2 \mapsto \mathbf{apply}^\varepsilon(e'_1, e'_2) : \tau'_2$ $e : [\delta/\bar{\alpha}]\tau \mapsto e'[\bar{\alpha} = \bar{\delta}] : \ \delta/\bar{\alpha}\ \tau\ $ $v : \forall \bar{\alpha}[D]. \tau \mapsto \mathbf{select} \{\bar{\delta}_i \Rightarrow e_i\}_i : \Pi \bar{\alpha}[D]. \tau$ $(\mathbf{fun} f(x) = e) : (\tau_1 \xrightarrow{\varepsilon} \tau_2)^\delta \mapsto (\mathbf{fun}^\varepsilon f(x) = e^\varepsilon) : \tau'_1 \xrightarrow{\varepsilon} \tau'_2$ $(\mathbf{let} x = e_1 \mathbf{in} e_2) : \tau_2 \mapsto (\mathbf{let} x = e'_1 \mathbf{in} e'_2) : \tau'_2$ $[e_1/x]e_2 : \tau_2 \mapsto (\mathbf{let} x = e'_1 \mathbf{in} e'_2) : \tau'_2$ $(\mathbf{case} e \mathbf{of} \{x_1 \Rightarrow e_1, x_2 \Rightarrow e_2\}) : \tau \mapsto (\mathbf{case} e' \mathbf{of} \{x_1 \Rightarrow e'_1, x_2 \Rightarrow e'_2\}) : \tau'$	$\text{if } e_1 : \tau_1 \mapsto e'_1 : \tau'_1 \text{ and } e_2 : \tau_2 \mapsto e'_2 : \tau'_2$ $\text{if } e : \tau_1 \mapsto e' : \tau'_1$ $\text{if } e : (\tau \times \tau_2)^\delta \mapsto e' : \tau' \times \tau'_2$ $\text{if } e_1 : \tau_1 \mapsto e'_1 : \tau'_1 \text{ and } e_2 : \tau_2 \mapsto e'_2 : \tau'_2$ $\text{if } e : \tau \mapsto e^C : \tau'$ $\text{if } e : \tau \mapsto e^S : \tau'$ $\text{if } e_1 : \tau_1 \mapsto e^S : \tau'_1 \mathbf{mod} \text{ and for all } v : \tau_1 \mapsto w : \tau'_1,$ $\text{we have } [v/x]e : \tau_2 \mapsto [w/x]e^C : \tau'_2$ $\text{if } e_1 : (\tau_1 \xrightarrow{\varepsilon} \tau_2)^\delta \mapsto e'_1 : \tau'_1 \xrightarrow{\varepsilon} \tau'_2 \text{ and } e_2 : \tau_1 \mapsto e'_2 : \tau'_1$ $\text{if } e : \forall \bar{\alpha}[D]. \tau \mapsto e' : \Pi \bar{\alpha}[D]. \tau$ $\text{if for all } i \text{ we have } v : [\delta/\bar{\alpha}]\tau \mapsto e_i : \ \delta/\bar{\alpha}\ \tau\ $ $\text{if for all } v : \tau_1 \mapsto w : \tau'_1,$ $((\mathbf{fun} f(x) = e)/f)[v/x]e : \tau_2$ $\mapsto ((\mathbf{fun}^\varepsilon f(x) = e^\varepsilon)/f)[v/x]e^\varepsilon : \tau'_2$ $\text{if } e_1 : \tau_1 \mapsto e'_1 : \tau'_1 \text{ and}$ $\text{for all } v : \tau_1 \mapsto w : \tau'_1, [v/x]e_2 : \tau_2 \mapsto [w/x]e'_2 : \tau'_2$ $\text{if } e_1 : \tau_1 \mapsto e'_1 : \tau'_1 \text{ and}$ $\text{for all } v : \tau_1 \mapsto w : \tau'_1, [v/x]e_2 : \tau_2 \mapsto [w/x]e'_2 : \tau'_2$ $\text{if } e : (\tau_1 + \tau_2)^\delta \mapsto e' : \tau'_1 + \tau'_2 \text{ and}$ $\text{for all } v : \tau_k \mapsto w : \tau'_k,$ $[v/x_k]e_k : \tau \mapsto [w/x_k]e'_k : \tau' \text{ for } k \in \{1, 2\}$

**Figure 19.** Correspondence of source and target expressions, used in Theorem 6.5 and other results

**Theorem 6.8** (Cost Result). *Given  $\mathcal{D} :: \rho \vdash e' \Downarrow (\rho' \vdash w)$  where for every subderivation  $\mathcal{D}^* :: \rho'_1 \vdash e^* \Downarrow (\rho'_2 \vdash w^*)$  of  $\mathcal{D}$  (including  $\mathcal{D}$ ),  $HC(\mathcal{D}^*) \leq k$ , then the number of dirty rule applications in  $\mathcal{D}$  is at most  $\frac{k}{k+1}W(\mathcal{D})$ .*

The cost theorem follows from Theorem C.5 (in the appendix)—a generalization of Theorem 6.4—and Theorem 6.8:

**Theorem 6.9.** *If  $\mathcal{D}$  derives  $\cdot \vdash \mathbf{trans}(e, \varepsilon) \Downarrow (\rho' \vdash w)$  then  $\mathcal{D}'$  derives  $e \Downarrow v$  where  $v : \tau \mapsto [\rho']w : \tau'$  and  $W(\mathcal{D}) \leq 7W(\mathcal{D}')$ .*

Acar et al. [2006] proved that given a well-typed AFL program, change propagation updates the output consistently with an initial run. Using Theorems 6.1 and 6.5, this implies that change propagation is consistent with an initial run of the source program.

## 7. Related Work

**Incremental computation.** Self-adjusting computation provides an approach to incremental computation, which has been studied extensively [Ramalingam and Reps 1993; Demers et al. 1981; Pugh and Teitelbaum 1989; Abadi et al. 1996]. Key techniques behind self-adjusting computation include dynamic dependence graphs, which allows a fully general change propagation mechanism [Acar et al. 2006], and a form of memoization that allows inexact computations to be reused via memoized computations that are (recursively) self-adjusting [Acar et al. 2009]. Programming-language features allow writing self-adjusting programs but these require syntactically separating stable and changeable data, as well as code that operates on such data [Acar et al. 2006, 2009; Ley-Wild et al. 2008; Hammer et al. 2009]. DITTO [Shankar and Bodik 2007] shows the benefits of eliminating user annotations. By customizing dependency tracking for invariant checking programs, DITTO provides a fully automatic incremental invariant checker. The ap-

proach, however, is domain-specific and only works for certain programs (e.g., functions cannot return arbitrary values): it is unsound in general.

**Information flow and constraint-based type inference.** A number of information flow type systems have been developed to check security properties, including the SLam calculus [Heintze and Riecke 1998], JFlow [Myers 1999] and a monadic system [Crary et al. 2005]. Our type system uses many ideas from Pottier and Simonet [2003], including a form of constraint-based type inference [Odersky et al. 1999], and is also broadly similar to other systems that use subtyping constraints [Simonet 2003; Foster et al. 2006].

**Cost semantics.** To prove that our translation yields efficient self-adjusting target programs, we use a simple cost semantics. The idea of instrumenting evaluations with cost information goes back to the early '90s [Sands 1990]. Cost semantics is particularly important in lazy [Sands 1990; Sansom and Peyton Jones 1995] and parallel languages [Spoonhower et al. 2008] where it is especially difficult to relate execution time to the source code, as well as in self-adjusting computation [Ley-Wild et al. 2009].

## 8. Conclusion

This paper presents techniques for translating purely functional programs to programs that can automatically self-adjust in response to dynamic changes to their data. Our contributions include a constraint-based type system for inferring self-adjusting-computation types from purely functional programs, a type-directed translation algorithm that rewrites purely functional programs into self-adjusting programs, and proofs of critical properties of the translation: type soundness and observational equivalence, as well as the intrinsic property of time complexity. Perhaps unsurprisingly,

the theorems and their proofs were critical to the determination of the type systems and the translation algorithm: many of our initial attempts at the problem resulted in target programs that were not type sound, that did not ensure observational equivalence, or were asymptotically slower than the source.

These results take an important step towards the development of languages and compilers that can generate code that can respond automatically to dynamically changing data correctly and asymptotically optimally, without substantial programming effort. Remaining open problems include generalization to imperative programs with references, techniques and proofs to determine or improve the asymptotic complexity of dynamic responses, and a complete and careful implementation and its evaluation.

### Acknowledgments

We thank the anonymous ICFP reviewers, as well as Arthur Charguéraud, for their useful comments on the submitted version of this paper.

### References

- M. Abadi, B. W. Lampson, and J.-J. Lévy. Analysis and caching of dependencies. In *International Conference on Functional Programming*, pages 83–91, 1996.
- U. A. Acar, G. E. Blelloch, and R. Harper. Adaptive functional programming. *ACM Trans. Prog. Lang. Sys.*, 28(6):990–1034, 2006.
- U. A. Acar, A. Ihler, R. Mettu, and O. Sümer. Adaptive Bayesian inference. In *Neural Information Processing Systems (NIPS)*, 2007.
- U. A. Acar, G. E. Blelloch, M. Blume, R. Harper, and K. Tangwongsan. An experimental analysis of self-adjusting computation. *ACM Trans. Prog. Lang. Sys.*, 32(1):3:1–53, 2009.
- U. A. Acar, G. E. Blelloch, R. Ley-Wild, K. Tangwongsan, and D. Türkoğlu. Traceable data types for self-adjusting computation. In *Programming Language Design and Implementation*, 2010a.
- U. A. Acar, A. Cotter, B. Hudson, and D. Türkoğlu. Dynamic well-spaced point sets. In *Symposium on Computational Geometry*, 2010b.
- M. Carlsson. Monads for incremental computing. In *International Conference on Functional Programming*, pages 26–35, 2002.
- Y. Chen, J. Dunfield, M. A. Hammer, and U. A. Acar. Online appendix to *Implicit Self-Adjusting Computation for Purely Functional Programs*, 2011. <http://www.mpi-sws.org/~joshua/Chen11appendix.pdf>.
- Y.-J. Chiang and R. Tamassia. Dynamic algorithms in computational geometry. *Proceedings of the IEEE*, 80(9):1412–1434, 1992.
- K. Crary, A. Kligler, and F. Pfenning. A monadic analysis of information flow security with mutable state. *Journal of Functional Programming*, 15(2):249–291, Mar. 2005.
- L. Damas and R. Milner. Principal type-schemes for functional programs. In *Principles of Programming Languages*, pages 207–212. ACM, 1982.
- A. Demers, T. Reps, and T. Teitelbaum. Incremental evaluation of attribute grammars with application to syntax-directed editors. In *Principles of Programming Languages*, pages 105–116, 1981.
- C. Demetrescu, I. Finocchi, and G. Italiano. *Handbook on Data Structures and Applications*, chapter 36: Dynamic Graphs. CRC Press, 2005.
- J. Field and T. Teitelbaum. Incremental reduction in the lambda calculus. In *ACM Conf. LISP and Functional Programming*, pages 307–322, 1990.
- J. S. Foster, R. Johnson, J. Kodumal, and A. Aiken. Flow-insensitive type qualifiers. *ACM Trans. Prog. Lang. Sys.*, 28:1035–1087, 2006.
- L. Guibas. Modeling motion. In J. Goodman and J. O’Rourke, editors, *Handbook of Discrete and Computational Geometry*, pages 1117–1134. Chapman and Hall/CRC, 2nd edition, 2004.
- M. A. Hammer, U. A. Acar, and Y. Chen. CEAL: a C-based language for self-adjusting computation. In *ACM SIGPLAN Conference on Programming Language Design and Implementation*, 2009.
- N. Heintze and J. G. Riecke. The SLam calculus: programming with secrecy and integrity. In *Principles of Programming Languages (POPL ’98)*, pages 365–377, 1998.
- R. Ley-Wild, M. Fluet, and U. A. Acar. Compiling self-adjusting programs with continuations. In *Int’l Conference on Functional Programming*, 2008.
- R. Ley-Wild, U. A. Acar, and M. Fluet. A cost semantics for self-adjusting computation. In *Proceedings of the 26th Annual ACM Symposium on Principles of Programming Languages*, 2009.
- MLton. MLton web site. <http://www.mlton.org>.
- A. C. Myers. JFlow: practical mostly-static information flow control. In *Principles of Programming Languages*, pages 228–241, 1999.
- M. Odersky, M. Sulzmann, and M. Wehr. Type inference with constrained types. *Theory and Practice of Object Systems*, 5(1):35–55, 1999.
- F. Pottier and V. Simonet. Information flow inference for ML. *ACM Trans. Prog. Lang. Sys.*, 25(1):117–158, Jan. 2003.
- W. Pugh and T. Teitelbaum. Incremental computation via function caching. In *Principles of Programming Languages*, pages 315–328, 1989.
- G. Ramalingam and T. Reps. A categorized bibliography on incremental computation. In *Principles of Programming Languages*, pages 502–510, 1993.
- A. Sabelfeld and A. C. Myers. Language-based information-flow security. *IEEE J. Selected Areas in Communications*, 21(1), 2003.
- D. Sands. *Calculi for Time Analysis of Functional Programs*. PhD thesis, University of London, Imperial College, Sept. 1990.
- P. M. Sansom and S. L. Peyton Jones. Time and space profiling for non-strict, higher-order functional languages. In *Principles of Programming Languages*, pages 355–366, 1995.
- A. Shankar and R. Bodik. DITTO: Automatic incrementalization of data structure invariant checks (in Java). In *Programming Language Design and Implementation*, 2007.
- V. Simonet. Type inference with structural subtyping: A faithful formalization of an efficient constraint solver. In *APLAS*, pages 283–302, 2003.
- D. Spoonhower, G. E. Blelloch, R. Harper, and P. B. Gibbons. Space profiling for parallel functional programs. In *International Conference on Functional Programming*, 2008.