

Typed Memory Management via Static Capabilities

DAVID WALKER

Carnegie Mellon University

KARL CRARY

Carnegie Mellon University

and

GREG MORRISETT

Cornell University

Region-based memory management is an alternative to standard tracing garbage collection that makes operations such as memory deallocation explicit but verifiably safe. In this article, we present a new compiler intermediate language, called the Capability Language (CL), that supports region-based memory management and enjoys a provably safe type system. Unlike previous region-based type systems, region lifetimes need not be lexically scoped, and yet the language may be checked for safety without complex analyses. Therefore, our type system may be deployed in settings such as extensible operating systems where both the performance and safety of untrusted code is important. The central novelty of the language is the use of static capabilities to specify the permissibility of various operations, such as memory access and deallocation. In order to ensure capabilities are relinquished properly, the type system tracks aliasing information using a form of bounded quantification. Moreover, unlike previous work on region-based type systems, the proof of soundness of our type system is relatively simple, employing only standard syntactic techniques. In order to show how our language may be used in practice, we show how to translate a variant of Tofte and Talpin's high-level type-and-effects system for region-based memory management into our language. When combined with known region inference algorithms, this translation provides a way to compile source-level languages to CL.

Categories and Subject Descriptors: D.3.1 [**Programming Languages**]: Formal Definitions and Theory—*Semantics, Syntax*; D.3.4 [**Programming Languages**]: Processors—*Compilers*; F.3.2 [**Logics and Meanings of Programs**]: Semantics of Programming Languages—*Operational Semantics*; F.3.3 [**Logics and Meanings of Programs**]: Studies of Program Constructs—*Type Structure*

General Terms: Languages, Theory, Verification

Additional Key Words and Phrases: Certified code, region-based memory management, type-directed compilation, typed intermediate languages

This material is based on work supported in part by the AFOSR grant F49620-97-1-0013 and the National Science Foundation under Grant No. EIA 97-03470. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not reflect the views of these agencies. This research was performed while the first and second authors were at Cornell University.

Authors' addresses: David Walker, Karl Crary, Carnegie Mellon University, Computer Science Department, 5000 Forbes Avenue, Pittsburgh, PA 15213. Greg Morrisett, Cornell University, Computer Science Department, Upson Hall, Ithaca, NY 14853.

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 2000 ACM 0164-0925/00/0500-0431 \$5.00

1. MOTIVATION AND BACKGROUND

A current trend in systems software is to allow untrusted extensions to be installed in protected services, relying upon language technology to protect the integrity of the service instead of hardware-based protection mechanisms [Lindholm and Yellin 1996; Wahbe et al. 1993; Bershad et al. 1995; Necula 1997; Morrisett et al. 1998; Kozen 1998; Hawblitzel et al. 1998]. For example, the SPIN project [Bershad et al. 1995] relies upon the Modula-3 type system to protect an operating system kernel from erroneous extensions. Similarly, Web browsers rely upon the Java Virtual Machine byte-code verifier [Lindholm and Yellin 1996] to protect users from malicious applets. In both situations, the goal is to eliminate expensive communications or boundary crossings by allowing extensions to directly access the resources they require.

Recently, Necula and Lee [Necula and Lee 1996; Necula 1997] have proposed Proof-Carrying Code (PCC) and Morrisett et al. [1999; 1998] have suggested Typed Assembly Language (TAL) as language technologies that provide the security advantages of high-level languages, but without the overheads of interpretation or just-in-time compilation. In both systems, low-level machine code can be heavily optimized, by hand or by compiler, and yet be automatically verified through proof- or type-checking.

However, in all of these systems (SPIN, JVM, TAL, and Touchstone [Necula and Lee 1998], a compiler that generates PCC), there is one aspect over which programmers and optimizing compilers have little or no control: memory management. In particular, their soundness depends on memory being reclaimed by a trusted garbage collector. Hence, applets or kernel extensions may not perform their own optimized memory management. Furthermore, as garbage collectors tend to be large, complicated pieces of unverified software, the degree of trust in language-based protection mechanisms is diminished.

The goal of this work is to provide a high degree of control over memory management for programmers and compilers, but as in the PCC and TAL frameworks, make verification of the safety of programs a straightforward task.

1.1 Regions

Tofte and Talpin [1994; 1997] suggest a type-and-effects system for verifying the soundness of *region-based* memory management. In later work, Tofte and others show how to infer region types and lifetimes and how to implement their theory [Tofte and Birkedal 1998; Birkedal et al. 1993; Birkedal et al. 1996]. There are several advantages to region-based memory management; from our point of view, the two most important are:

- (1) Region-based memory management can be implemented using relatively simple constant-time routines.
- (2) All memory operations are explicit in the program text, but safety is guaranteed.

The first advantage has several implications. If regions are used in a secure system then the simplicity of the implementation leads to a smaller trusted computing base. Moreover, it may be possible to formally verify that the region operations are implemented correctly. In contrast, a formal analysis of a garbage collector would

be a much more onerous task. Second, because region operations are constant-time and do not trace the structure of the heap, programs do not suffer from the pauses that are associated with conventional garbage collectors. Consequently, region-based memory management systems may be a practical alternative to real-time garbage collectors [Baker 1978; Wilson 1992].

The second advantage gives programmers greater control over memory use. By using a region-profiler [Birkedal et al. 1993], programmers can quickly identify the memory regions that are causing performance problems in their applications. Next, because allocation and deallocation operations are explicit in the program text, programmers can use the profiling data to accurately relate the run-time behavior of programs to their static representation. In other words, given information about the ways regions are used at run time, it is often straightforward to examine program code, identify memory-intensive routines, and reason about the lifetimes of the data structures allocated there. Once the trouble spots have been identified, programmers can concentrate their optimization efforts on a small portion of the code. Most importantly, throughout the programming process, a type checker guarantees that all memory operations are safe. More specifically, it prevents dereferencing a pointer to an object that has been deallocated so programmers do not have to worry about programs crashing due to memory faults. It also ensures that every memory region that is allocated in a program is later deallocated (assuming the program does not enter an infinite loop).

In order to ensure that regions are used safely, the Tofte-Talpin language includes a lexically scoped expression “`letregion r in e end`” that delimits the lifetime of a region r . A region is allocated when control enters the scope of the `letregion` construct and is deallocated when control leaves the scope. Programs may allocate values into live regions using the notation `v at r`. These values may be used until the region is deallocated. For example,

$$\text{Region lifetime} \left\{ \begin{array}{ll} \vdots & \\ \text{letregion } r \text{ in} & \% \text{ Allocate region } r \\ \quad \text{let } x = v \text{ at } r \text{ in} & \% \text{ Allocate value } v \text{ in } r \\ \quad \quad f(r, x) & \% \text{ Function } f \text{ may access } r \\ \text{end} & \% \text{ Deallocate } r \text{ (and } v) \\ \vdots & \end{array} \right.$$

Tofte and Talpin ensure that deallocated values are not accessed unsafely using a type-and-effects system. Informally, whenever an expression uses a value in region r , the type system expresses this fact using the effect $\text{access}(r)$. However, effects occurring within the scope of the `letregion` construct are masked. More specifically, if the expression e has effects $\text{access}(r) \cup \psi$ (for some set of effects ψ) then the overall effect of the expression `letregion r in e end` is simply ψ . Hence, if there is no overall effect for an entire program then every region access must have occurred within the scope of the corresponding `letregion` construct. In other words, values in region r are used only during the lifetime of r and not before or after. If this condition holds, we can conclude the program is safe.

The Tofte-Talpin language makes efficient use of memory provided that the lifetimes of values coincide with the lexical structure of the program. However, if

lifetimes deviate from program structure then this style of region-based memory management may force programs to use considerably more memory than necessary. Consider the following (yet to be region-annotated) program fragments.

```
% Scope 1:  The Call Site

let x = v in
:
let y = f (x) in
:
y is dead
```

```
% Scope 2:  The Function

fun f (x) =
:
  x is dead
:
  let y = v' in
:
  return y
```

The value v is an argument to the function f and must be allocated in the scope of the function call. However, suppose that when f is executed, v dies quickly. The value v' exhibits the inverse behavior. It is allocated inside f but is returned as the function result. Both v and v' have lifetimes that span two lexical scopes, but neither is live for very long in either scope. Consequently, vanilla region inference does not perform well in this setting. The best it can do is wrap the function call in a pair of `letregion` commands.

```
% Scope 1:  The Call Site

letregion r in
  let x = v at r in
  :
  letregion r' in
    let y = f (r, r', x) in
    :
    y is dead
  end (r')
end (r)
:
```

```
% Scope 2:  The Function

fun f (r, r', x) =
:
  x is dead
:
  let y = v' at r' in
:
  return y
```

Here, the regions r and r' are live much longer than they need to be due to the inflexibility of the `letregion` construct. Both regions must be allocated outside the function call. Notice also that even though v is dead when the function call returns, the outer region r cannot be deallocated until after the inner region r' has been deallocated. Lexical scoping enforces a stack-like, last-allocated/first-deallocated memory management discipline.

In this example, a much better solution to this memory management problem is to separate region allocation (`newregion`) from deallocation (`freeregion`). The following program takes this approach. In principle, since the lifetimes of regions r and r' do not overlap, the memory for these regions could be reused.

<pre> % Scope 1: The Call Site let newregion r in let x = v at r in : let r', y = f (r, x) in : y is dead let freeregion r' in : </pre>	<pre> % Scope 2: The Function fun f (r, x) = : x is dead let freeregion r in : let newregion r' in let y = v' at r' in : return (r', y) </pre>
---	--

Unfortunately, we cannot write this program in the Tofte-Talpin language because it is based on the idea of lexical scoping. Another consequence of this language design is that any program transformation that alters program structure can affect memory management. One of the most devastating transformations for the Tofte-Talpin type system is the continuation-passing style (CPS) transformation; each successive computation is placed in the scope of all previous computations, with the result that no regions can be deallocated until the entire computation has been completed. In the following example, the CPS transformation prevents the region r from being deallocated until after `code` has been executed when it could be deallocated as soon as f has completed its computation.

$$\begin{array}{l}
 \text{letregion } r \text{ in} \\
 \quad f(r, v) \\
 \text{end;} \\
 \text{code}
 \end{array}
 \Rightarrow
 \begin{array}{l}
 \text{letregion } r \text{ in} \\
 \quad f(r, v, \lambda.\text{code}) \\
 \text{end}
 \end{array}$$

The observation that the Tofte-Talpin type system will make poor use of memory in such cases has been made before. Birkedal et al. [1996] and Aiken et al. [1995] have proposed optimizations that allow regions to be freed early. However, although their optimizations are safe, there is no simple proof- or type-checker that an untrusting client can use to check the output code. Therefore, in order to construct a verifiably safe, efficient region-based language, we must reexamine the fundamental question: “*When can a program access a value v ?*”

1.2 Contributions

Our solution to the problem of provably safe yet efficient region-based memory management takes its inspiration from operating systems such as Hydra [Wulf et al. 1981]. Hydra solves the access control problem by associating an unforgeable key or *capability* with every object and requiring that the user present this capability to gain access to the object. Furthermore, when the need arises, Hydra revokes capabilities, thereby preventing future access to the protected objects.

We define a new strongly typed compiler intermediate language for region-based memory management that uses a compile-time notion of capability to ensure that

region operations are performed safely. Unlike Tofte and Talpin’s language, lexical scoping plays no part in the verification of our Capability Language (CL). Instead, the type system threads static information in the form of capabilities along the control-flow path of a program. In order to use a value in region r at a particular control-flow point, the program must present the capability for that region. As in traditional capability systems, our type system keeps track of capability copies carefully so that it can determine when a capability has truly been revoked. Unlike in traditional capability systems, which allow the operating system to revoke capabilities without the cooperation of client code, programs in our language must voluntarily give up their capabilities. However, the capabilities in our language are a purely static concept, and thus their implementation requires no run-time overhead. This mechanism provides an efficient way to check the safety of explicit, arbitrarily ordered region allocation and deallocation instructions.

We have a purely *syntactic* argument, based on Subject Reduction and Progress lemmas in the style of Felleisen and Wright [Wright and Felleisen 1994], that the type system of CL is sound. In contrast, Tofte and Talpin formulate the soundness of their system using a more complicated greatest fixed point argument [Tofte and Talpin 1997], and the soundness of Aiken et al. ’s optimizations [Aiken et al. 1995] depends upon this argument. Part of the reason for the extra complexity is that Tofte and Talpin simultaneously show that region *inference* translates lambda calculus terms into operationally equivalent region calculus terms, a stronger property than we prove. However, when system security is the main concern, soundness is the critical property.

We also have a formal translation of a variant of the Tofte-Talpin language into our calculus. Given a type-safe Tofte-Talpin program, the translation always produces a type-safe CL program. Therefore, when the translation is combined with a region inference algorithm [Tofte and Birkedal 1998], it provides a way to compile source language programs into type-safe low-level code that can be used in secure extensible systems.

The technical sections of this article are derived from earlier work presented at the Twenty-Sixth Symposium on the Principles of Programming Languages [Crory et al. 1999]. Section 2 presents the syntax and semantics of CL formally and motivates the design decisions that we made. At the end of this section, we present the type soundness theorem and discuss the most interesting parts of our proof. The complete proof may be found in Appendix A. Section 3 demonstrates that CL is at least as expressive as the Tofte-Talpin language. We define the semantics of a variant of the latter language and give a translation into CL. The translation preserves the type safety property, and Appendix B proves this fact. We further demonstrate the expressiveness of CL by sketching a couple of optimizations that are not possible in the more restrictive language. Section 4 informally explores several other applications of capabilities. This section also explains further connections with related work. Finally, Section 5 concludes.

2. THE CAPABILITY LANGUAGE

The central technical contribution of this article is CL, a statically typed intermediate language that supports the explicit allocation, deallocation, and access of memory regions.

As mentioned in the introduction, the type system for the language propagates static information (capabilities) along the control-flow path of a program. Therefore, the most elegant and natural form for the language is continuation-passing style (CPS) [Reynolds 1972]. That is, functions in CL do not return values; instead, functions finish by calling a continuation function that is typically provided as an argument. The fact that there is only one means of transferring control between functions in CPS—rather than the two means (call and return) in direct style—simplifies the tracking of capabilities. A direct style formulation is possible, but the complications involved obscure the central issues. In the remainder of this paper, we assume familiarity with CPS.

The syntax of the capability abstract machine appears in Figure 1. In the remainder of this section, we explain and motivate the main constructs and typing rules of the language one by one.

Notation. We treat terms and types that differ only in the names of bound variables as equivalent. We use the notation $\{x_1 \mapsto X_1, \dots, x_n \mapsto X_n\}$ or the notation $\{x_1:X_1, \dots, x_n:X_n\}$ to denote finite partial maps, which are equivalent up to re-ordering of their elements. In the subsections that follow, we will build a more sophisticated notion of type equivalence on top of these standard conventions. Given a finite partial map M , $Dom(M)$ denotes the domain of the map. The notation $M\{x \mapsto X\}$ or $M\{x:X\}$ denotes a new map M' that maps x to X but is otherwise identical to M whereas $M \setminus x$ denotes a new map M' that is undefined at x but is otherwise identical to M . The notation $X[Y/x]$ denotes standard capture-avoiding substitution of Y for x in X .

If Δ is a sequence of bindings of the form $\alpha_i:\kappa_i$ or $\alpha_i \leq C$ (where i ranges from 1 to n) then $Dom(\Delta)$ is the sequence of constructor variables $\alpha_1, \dots, \alpha_n$. Occasionally, we will use the notation $[c_1, \dots, c_n/\Delta]$ to refer to the simultaneous capture-avoiding substitution $[c_1, \dots, c_n/\alpha_1, \dots, \alpha_n]$ where $\alpha_1, \dots, \alpha_n = Dom(\Delta)$. We use the notation $\Delta\Delta'$ to indicate the constructor context formed by concatenating the elements of Δ' onto Δ . This notation is only defined if $Dom(\Delta)$ and $Dom(\Delta')$ have no elements in common.

2.1 Preliminaries

We specify the operational behavior of CL using a call-by-value allocation semantics [Morrisett et al. 1995; Morrisett and Harper 1997], which makes the allocation of data in memory explicit. The semantics, which is specified in Figure 2, is given by a deterministic rewriting system $P \mapsto P'$ mapping machine states to new machine states. A machine state consists of a pair (M, e) of a memory and a term being executed. A memory is a finite mapping of *region names* (ν) to *regions* where a region is a block of memory that stores a collection of individual objects. Regions are created at run time by the declaration `newrgn ρ, x` , which extends memory with a new region (ν),¹ binds ρ to the name of that region, and binds x to the *handle* (`handle(ν)`) for that region. Both ρ and x are considered bound variables for the purposes of alpha-conversion.

¹A “new” region is one that does not occur anywhere in the current memory (i.e., the region’s name does not occur in the domain of current memory, nor does it occur in any stored value) or in the expression being executed.

<i>kinds</i>	$\kappa ::= \text{Type} \mid \text{Rgn} \mid \text{Cap}$
<i>constructor vars</i>	α, ρ, ϵ
<i>constructors</i>	$c ::= \alpha \mid \tau \mid r \mid C$
<i>types</i>	$\tau ::= \alpha \mid \text{int} \mid r \text{ handle} \mid \langle \tau_1, \dots, \tau_n \rangle \text{ at } r \mid \forall[\Delta].(C, \tau_1, \dots, \tau_n) \rightarrow 0 \text{ at } r$
<i>regions</i>	$r ::= \rho \mid \nu$
<i>capabilities</i>	$C ::= \epsilon \mid \emptyset \mid \{r^\varphi\} \mid C_1 \oplus C_2 \mid \overline{C}$
<i>multiplicities</i>	$\varphi ::= 1 \mid +$
<i>con. contexts</i>	$\Delta ::= \cdot \mid \Delta, \alpha : \kappa \mid \Delta, \epsilon \leq C$
<i>value contexts</i>	$\Gamma ::= \cdot \mid \Gamma, x : \tau$
<i>region types</i>	$\Upsilon ::= \{\ell_1 : \tau_1, \dots, \ell_n : \tau_n\}$
<i>memory types</i>	$\Psi ::= \{\nu_1 : \Upsilon_1, \dots, \nu_n : \Upsilon_n\}$
<i>word values</i>	$v ::= x \mid i \mid \nu.\ell \mid \text{handle}(\nu) \mid v[c]$
<i>heap values</i>	$h ::= \text{fix } f[\Delta](C, x_1 : \tau_1, \dots, x_n : \tau_n).e \mid \langle v_1, \dots, v_n \rangle$
<i>arithmetic ops</i>	$p ::= + \mid - \mid \times$
<i>declarations</i>	$d ::= x = v \mid x = v_1 p v_2 \mid x = h \text{ at } v \mid x = \pi_i v \mid \text{newrgn } \rho, x \mid \text{freergn } v$
<i>terms</i>	$e ::= \text{let } d \text{ in } e \mid \text{if } 0 \text{ then } e_2 \text{ else } e_3 \mid v(v_1, \dots, v_n) \mid \text{halt } v$
<i>memory regions</i>	$R ::= \{\ell_1 \mapsto h_1, \dots, \ell_n \mapsto h_n\}$
<i>memories</i>	$M ::= \{\nu_1 \mapsto R_1, \dots, \nu_n \mapsto R_n\}$
<i>machine states</i>	$P ::= (M, e)$

Fig. 1. Capability abstract machine syntax.

Region names and handles are distinguished in order to maintain a phase distinction between compile-time and run-time expressions. Region names are significant at compile time: The type-checker identifies which region an object inhabits via a region name (see below). However, region names, like other type constructors, have no run-time significance and may be erased from executable code. In contrast, region handles hold the run-time data necessary to manipulate regions. In addition to accounting for a phase distinction, the separation of region names and handles also allows us to refine the contexts in which region handles are needed. Handles are needed when allocating objects within a region (to increment the region’s allocation pointer) and when freeing a region (to return region memory to the free storage list), but are not needed when reading data from a region.

Regions are freed by the declaration `freergn v`, where v is the handle for the region to be freed. Objects h large enough to require heap allocation (i.e., function closures and tuples), called *heap values*, are allocated by the declaration $x = h \text{ at } v$, where v is the handle for the region in which h is to be allocated. Data are read from a region in two ways: functions are read by a function call, and tuples are read by the declaration $x = \pi_i(v)$, which binds x to the data residing in the i th field of the object at address v . Each of these operations may be performed only when the region in question has not already been freed. Enforcing this restriction is the purpose of the capability mechanism discussed in Section 2.2.

A region maps locations (ℓ) to heap values. An address is given by a pair $\nu.\ell$ of a region name and a location. We often abbreviate address lookup $M(\nu)(\ell)$ by $M(\nu.\ell)$ and address update $M\{\nu \mapsto M(\nu)\{\ell \mapsto h\}\}$ by $M\{\nu.\ell \mapsto h\}$. In the course of execution, word-sized values (v) will be substituted for value variables and type

constructors for constructor variables, but heap values (h) are not substituted for value variables. When executing the declaration $x = h$ at v (where v is $\text{handle}(\nu)$, the handle for region ν), h is allocated in region ν (say at ℓ) and the address $\nu.\ell$ is substituted for x in the following code. Hence, programs always refer to heap-allocated values indirectly via an address.

A term in CL consists of a series of declarations ending in either a branch or a function call (or a halt). The class of declarations includes those constructs discussed above, plus two standard constructs, $x = v$ for binding variables to values and $x = v_1 p v_2$ (where p ranges over $+$, $-$ and \times) for integer arithmetic.

For example, the program below allocates a region and puts a pair of integers inside it. Next, the components of the pair are projected from the tuple, and the region is deallocated. Finally, the program sums the two integers and terminates.

```

let newrgn  $\rho, x_\rho$    in    % Allocate region  $\rho$ 
let  $y = \langle 1, 2 \rangle$  at  $x_\rho$  in % Allocate pair in  $\rho$ 
let  $t_1 = \pi_1 y$      in    % Access region  $\rho$ , no handle required
let  $t_2 = \pi_2 y$      in
let freergn ( $x_\rho$ )  in    % Deallocate region  $\rho$ 
let  $z = t_1 + t_2$     in
halt  $z$               % Terminate
    
```

Types. The types of CL include type constructor variables and *int*, a type of region handles, as well as tuple and function types. If r is a region, then r handle is the type of r 's region handle. The tuple type $\langle \tau_1, \dots, \tau_n \rangle$ at r contains the usual n field tuples, but also specifies that such tuples are allocated in region r , where r is either a region name ν or, more frequently, a region variable ρ .

The function type $(C, \tau_1, \dots, \tau_n) \rightarrow 0$ at r contains functions taking n arguments (with types τ_1 through τ_n) that may be called when capability C is satisfied (see the next subsection). The 0 return type is intended to suggest the fact that CPS functions invoke their continuations rather than returning as a direct-style function does. The suffix “at r ”, like the corresponding suffix for tuple types, indicates the region in which the function is allocated.

Functions may be made polymorphic over types, regions, or capabilities by adding a constructor context Δ to the function type as in $\forall[\Delta].(C, \tau_1, \dots, \tau_n) \rightarrow 0$ at r .² For convenience, types, regions, and capabilities are combined into a single syntactic class of “constructors” and are distinguished by kinds. Thus, a type is a constructor with kind **Type**; a region is a constructor with kind **Rgn**; and a capability is a constructor with kind **Cap**. We use the metavariable c to range over constructors, but use the metavariables τ , r , and C when those constructors are types, regions and capabilities, respectively. We also use the metavariables ρ and ϵ for constructor variables of kind **Rgn** and **Cap**, and use the metavariable α for type variables and generic constructor variables.

²Technically, all function types have the form $\forall[\Delta].(C, \tau_1, \dots, \tau_n) \rightarrow 0$ at r . However, when Δ is empty, as in the type $\forall[\cdot].(C, \tau_1, \dots, \tau_n) \rightarrow 0$ at r , we abbreviate it by $(C, \tau_1, \dots, \tau_n) \rightarrow 0$ at r .

If $e =$	then $P =$
$\text{let } x = v \text{ in } e'$	$(M, e'[v/x])$
$\text{let } x = i \text{ p } j \text{ in } e'$	$(M, e'[k/x])$ where $k = i \text{ p } j$
$\text{let } x = h \text{ at } (\text{handle}(\nu)) \text{ in } e'$ and $\nu \in \text{Dom}(M)$	$(M\{\nu.\ell \mapsto h\}, e'[\nu.\ell/x])$ where $\ell \notin \text{Dom}(M(\nu))$
$\text{let } x = \pi_i(\nu.\ell) \text{ in } e'$ and $\nu \in \text{Dom}(M)$ and $\ell \in \text{Dom}(M(\nu))$	$(M, e'[v_i/x])$ where $M(\nu.\ell) = \langle v_1, \dots, v_n \rangle$ ($1 \leq i \leq n$)
$\text{let newrgn } \rho, x \text{ in } e'$	$(M\{\nu \mapsto \{\}\}, e'[\nu, \text{handle}(\nu)/\rho, x])$ where $\nu \notin M$ and $\nu \notin e'$
$\text{let freergn } (\text{handle}(\nu)) \text{ in } e'$ and $\nu \in \text{Dom}(M)$	$(M \setminus \nu, e')$
$\text{if } 0 \text{ then } e_2 \text{ else } e_3$	(M, e_2)
$\text{if } 0 \text{ i then } e_2 \text{ else } e_3$ and $i \neq 0$	(M, e_3)
$\nu.\ell[c_1, \dots, c_m](v_1, \dots, v_n)$ and $M(\nu.\ell) = \text{fix } f[\Delta](C, x_1:\tau_1, \dots, x_n:\tau_n).e$ and $\text{Dom}(\Delta) = \alpha_1, \dots, \alpha_m$	$(M, S_2(S_1(e)))$ where $S_1 = [c_1, \dots, c_m/\alpha_1, \dots, \alpha_m]$ and $S_2 = [\nu.\ell, v_1, \dots, v_n/f, x_1, \dots, x_n]$

Fig. 2. Capability operational semantics: $(M, e) \mapsto P$.

For example, a polymorphic identity function that is allocated in region r , but whose continuation function may be in any region, may be given type

$$\forall[\alpha:\text{Type}, \rho:\text{Rgn}].(C, \alpha, (C, \alpha) \rightarrow 0 \text{ at } \rho) \rightarrow 0 \text{ at } r$$

for some appropriate C . Let f be such a function; let v be its argument with type τ ; and let g be its continuation with type $(C, \tau) \rightarrow 0 \text{ at } r$. Then f is called by $f[\tau][r](v, g)$. A more detailed explanation of functions is contained in the next subsection where we discuss the role of capabilities.

Figure 3 specifies all well-formed constructors and constructor contexts. The two main judgments $\Delta \vdash \Delta'$ and $\Delta \vdash c : \kappa$ assume that the constructor context Δ is well-formed. The first judgment states that Δ' is a well-formed constructor context, and the second judgment states c is a well-formed constructor with kind κ .

$\Delta \vdash \Delta'$

$$\frac{}{\Delta \vdash \cdot} \text{ (ctxt-empty)}$$

$$\frac{\Delta \vdash \Delta'}{\Delta \vdash \Delta', \alpha : \kappa} \text{ } (\alpha \notin \text{Dom}(\Delta\Delta')) \text{ (ctxt-var)}$$

$$\frac{\Delta \vdash \Delta' \quad \Delta\Delta' \vdash C : \text{Cap}}{\Delta \vdash \Delta', \epsilon \leq C} \text{ } (\epsilon \notin \text{Dom}(\Delta\Delta')) \text{ (ctxt-sub)}$$

 $\Delta \vdash c : \kappa$

$$\frac{}{\Delta \vdash \alpha : \kappa} \text{ } (\Delta(\alpha) = \kappa) \text{ (type-var)}$$

$$\frac{}{\Delta \vdash \epsilon : \text{Cap}} \text{ } ((\epsilon \leq C) \in \Delta) \text{ (type-sub)}$$

$$\frac{}{\Delta \vdash \text{int} : \text{Type}} \text{ (type-int)}$$

$$\frac{\Delta \vdash r : \text{Rgn}}{\Delta \vdash r \text{ handle} : \text{Type}} \text{ (type-handle)}$$

$$\frac{\Delta \vdash \tau_i : \text{Type} \text{ (for } 1 \leq i \leq n) \quad \Delta \vdash r : \text{Rgn}}{\Delta \vdash \langle \tau_1, \dots, \tau_n \rangle \text{ at } r : \text{Type}} \text{ (type-tuple)}$$

$$\frac{\Delta \vdash \Delta' \quad \Delta\Delta' \vdash \tau_i : \text{Type} \text{ (for } 1 \leq i \leq n) \quad \Delta\Delta' \vdash C : \text{Cap} \quad \Delta \vdash r : \text{Rgn}}{\Delta \vdash \forall[\Delta']. (C, \tau_1, \dots, \tau_n) \rightarrow 0 \text{ at } r : \text{Type}} \text{ (type-arrow)}$$

$$\frac{}{\Delta \vdash \nu : \text{Rgn}} \text{ (type-name)}$$

$$\frac{}{\Delta \vdash \emptyset : \text{Cap}} \text{ (type-}\emptyset\text{)}$$

$$\frac{\Delta \vdash r : \text{Rgn}}{\Delta \vdash \{r^\varphi\} : \text{Cap}} \text{ (type-single)}$$

$$\frac{\Delta \vdash C_1 : \text{Cap} \quad \Delta \vdash C_2 : \text{Cap}}{\Delta \vdash C_1 \oplus C_2 : \text{Cap}} \text{ (type-plus)}$$

$$\frac{\Delta \vdash C : \text{Cap}}{\Delta \vdash \bar{C} : \text{Cap}} \text{ (type-bar)}$$

Fig. 3. Capability static semantics: Type and context formation.

The typing rules also use region types (Υ), which assign a type to every location allocated in a region, and memory types (Ψ), which assign a region type to every region allocated in memory. However, it is not necessary to understand these constructs in the preliminary development, and therefore we will defer discussing them until we describe the static semantics of the abstract machine in formal detail (see Section 2.4).

2.2 Capabilities

The central problem is how to ensure statically that no region is used after it is freed. The typing rules enforce this with a system of capabilities that specify what

operations are permitted. The main typing judgment is

$$\Psi; \Delta; \Gamma; C \vdash e$$

which states that (when memory has type Ψ , free constructor variables have kinds given by Δ , and free value variables have types given by Γ) it is legal to execute the term e , *provided that the capability C is held*. A related typing judgment is

$$\Psi; \Delta; \Gamma; C \vdash d \Rightarrow \Delta'; \Gamma'; C'$$

which states that if the capability C is held, it is legal to execute the declaration d , which results in new constructor context Δ' , new value context Γ' , and new capability C' .

Capabilities indicate the set of regions that are presently valid to access, that is, those regions that have not been freed. Capabilities are formed by joining together a collection of singleton capabilities $\{r\}$ that provide access to only one region, and capability variables ϵ that provide access to an unspecified set of regions. Capability joins, written $C_1 \oplus C_2$, are associative and commutative, but are not always idempotent; in Section 2.3 we will define equality on capabilities formally, and we will see examples where $C \oplus C$ is not equivalent to C . The empty capability, which provides access to no regions, is denoted by \emptyset . We will often abbreviate the capability $\{r_1\} \oplus \dots \oplus \{r_n\}$ by $\{r_1, \dots, r_n\}$.

In order to read a field from a tuple in region r , it is necessary to hold the capability to access r , as in the rule:

$$\frac{\begin{array}{l} \Delta \vdash C = C' \oplus \{r\} : \text{Cap} \\ \Psi; \Delta; \Gamma \vdash v : \langle \tau_1, \dots, \tau_n \rangle \text{ at } r \end{array}}{\Psi; \Delta; \Gamma; C \vdash x = \pi_i v \Rightarrow \Delta; \Gamma\{x:\tau_i\}; C} \quad (x \notin \text{Dom}(\Gamma) \wedge 1 \leq i \leq n)$$

The first subgoal indicates that the capability held (C) is equivalent to some capability that includes $\{r\}$.

A similar rule is used to allocate an object in a region. Since the type of a heap value reflects the region in which it is allocated, the heap value typing judgment (the second subgoal below) must be provided with that region.

$$\frac{\begin{array}{l} \Delta \vdash C = C' \oplus \{r\} : \text{Cap} \\ \Psi; \Delta; \Gamma \vdash h \text{ at } r : \tau \\ \Psi; \Delta; \Gamma \vdash v : r \text{ handle} \end{array}}{\Psi; \Delta; \Gamma; C \vdash x = h \text{ at } v \Rightarrow \Delta; \Gamma\{x:\tau\}; C} \quad (x \notin \text{Dom}(\Gamma))$$

Functions. Functions are defined by the following form

$$\mathbf{fix} f[\Delta](C, x_1:\tau_1, \dots, x_n:\tau_n).e$$

where f stands for the function itself and may appear free in the body, where Δ specifies the function's constructor arguments, and where C is the function's capability precondition. When Δ is empty and f does not appear free in the function body we abbreviate the \mathbf{fix} form by $\lambda(C, x_1:\tau_1, \dots, x_n:\tau_n).e$.

In order to call a function residing in region r , it is again necessary to hold the capability to access r , as well as to hold a capability equivalent to the function's

capability precondition:

$$\frac{\Delta \vdash C = C'' \oplus \{r\} : \text{Cap} \quad \Delta \vdash C = C' : \text{Cap} \quad \Psi; \Delta; \Gamma \vdash v : (C', \tau_1, \dots, \tau_n) \rightarrow 0 \text{ at } r \quad \Psi; \Delta; \Gamma \vdash v_i : \tau_i}{\Psi; \Delta; \Gamma; C \vdash v(v_1, \dots, v_n)}$$

The body of a function may then assume the function’s capability precondition is satisfied, as indicated by the capability C in the premise of the rule:³

$$\frac{\Psi; \Delta; \Gamma \{x_1:\tau_1, \dots, x_n:\tau_n\}; C \vdash e}{\Psi; \Delta; \Gamma \vdash \lambda(C, x_1:\tau_1, \dots, x_n:\tau_n).e \text{ at } r : \tau_f} \quad (x_i \notin \text{Dom}(\Gamma))$$

As might be expected, the annotation “at r ” indicates that the closure value resides in region r . The resultant function type τ_f is $(C, \tau_1, \dots, \tau_n) \rightarrow 0 \text{ at } r$.

Often, we will extend the required capability for a function with a quantified capability variable (similar to a *row variable*). This variable may be instantiated with whatever capabilities are leftover after satisfying the required capability. Consequently, the function may be used in a variety of contexts. For example, functions with type

$$\forall[\epsilon:\text{Cap}].(\{r\} \oplus \epsilon, \dots) \rightarrow 0 \text{ at } r$$

may be called with any capability that extends $\{r\}$.

In source languages such as Standard ML [Milner et al. 1997], polymorphism is normally restricted to second-class status—polymorphic functions cannot be stored in data structures or passed as arguments to other functions—because these languages support type inference for unannotated programs. However, in a typed intermediate language such as CL, it is unnecessary to place such restrictions on polymorphic functions; explicit type annotations make type checking first-class polymorphism straightforward. The additional expressiveness of first-class polymorphism is necessary whenever a function (say f) allocates a new region ρ and returns its result in ρ along with ρ itself. In this case, f ’s continuation must be polymorphic with respect to the returned region, leading to the following type for f .⁴

$$\forall[\epsilon:\text{Cap}].(\{r\} \oplus \epsilon, \dots, \forall[\rho:\text{Rgn}].(\{\rho, r\} \oplus \epsilon, \tau_{res} \text{ at } \rho) \rightarrow 0 \text{ at } r) \rightarrow 0 \text{ at } r$$

This and similar scenarios were among the main motivations for choosing a continuation-passing style language. In a direct-style calculus, we would have to define new syntax to allow functions to return static type, region, and capability information to their calling contexts. Similarly, the context following an if statement would require special annotations to compute the join of capability information from both branches. Polymorphic continuations provide a uniform solution to both these problems.

When a function or continuation is polymorphic, its type constructor arguments may be instantiated one at a time, leading to partially applied polymorphic functions with the form $v[c]$. As mentioned in Section 2.1, type constructors c may

³This rule specializes the full rule for fix to the case where the function is neither polymorphic nor recursive.

⁴The type $\tau_{res} \text{ at } \rho$ could be a pair in region ρ containing the actual function result as well as a handle for region ρ so the continuation can deallocate ρ .

be erased before code is executed. Consequently, this partial application can be treated as a first-class value with the following typing rule:

$$\frac{\Psi; \Delta; \Gamma \vdash v : \forall[\alpha:\kappa, \Delta'].(C, \tau_1, \dots, \tau_n) \rightarrow 0 \text{ at } r \quad \Delta \vdash c : \kappa}{\Psi; \Delta; \Gamma \vdash v[c] : (\forall[\Delta'].(C, \tau_1, \dots, \tau_n) \rightarrow 0)[c/\alpha] \text{ at } r}$$

The common case is still to apply multiple type arguments at once. We often abbreviate multiple type applications $v[c_1] \cdots [c_n]$ by $v[c_1, \dots, c_n]$. As indicated in the rule for function call, a function must be fully applied before it can be called.

Allocation and Deallocation. The most delicate issue is the typing of region allocation and deallocation. Intuitively, the typing rules for the **newrgn** and **freergn** declarations should add and remove capabilities for the appropriate region. Naive typing rules could be:

$$\frac{}{\Psi; \Delta; \Gamma; C \vdash \text{newrgn } \rho, x \Rightarrow \Delta, \rho:\text{Rgn}; \Gamma\{x:\rho \text{ handle}\}; C \oplus \{\rho\}} \text{ (wrong)}$$

$$\frac{\Psi; \Delta; \Gamma \vdash v : r \text{ handle} \quad C' = C \setminus \{r\}}{\Psi; \Delta; \Gamma; C \vdash \text{freergn } v \Rightarrow \Delta; \Gamma; C'} \text{ (wrong)}$$

We will be able to use something much like the first rule for allocation, but the naive rule for freeing regions is fundamentally flawed. For example, consider the following function:

```
fix f[\rho_1:\text{Rgn}, \rho_2:\text{Rgn}](\{\rho_1, \rho_2\}, x:\rho_1 \text{ handle}, y:\langle \text{int} \rangle \text{ at } \rho_2).
  let freergn x in
  let z = \pi_1 y in \dots
```

This function is well-formed according to the naive typing rule: The function begins with the capability $\{\rho_1, \rho_2\}$, and ρ_1 is removed by the **freergn** declaration, leaving $\{\rho_2\}$. The tuple y is allocated in ρ_2 , so the projection is legal. However, this code is operationally incorrect if ρ_1 and ρ_2 are instantiated by the same region r . In that case, the first declaration frees r , and the second attempts to read from r .

This problem is a familiar one. To free a region safely it is necessary to delete all copies of the capability. However, instantiating region variables can create aliases, making it impossible to tell by inspection whether any copies exist.

2.3 Alias Control

We desire a system for alias control that can easily be enforced by the type system, without expensive and complex program analyses. One possibility is a linear type system [Girard 1987; Wadler 1990; 1993]. In a linear type system, aliasing would be trivially controlled; any use of a region name would consume that name, ensuring that it could not be used elsewhere. Thus, in a linear type system, the naive rules for allocating and deallocating regions would be sound. Unfortunately, a linear type system is too restrictive to permit many useful programs. For example, suppose f has type

$$\forall[\rho_1:\text{Rgn}, \rho_2:\text{Rgn}].(\{\rho_1, \rho_2\}, \langle \text{int} \rangle \text{ at } \rho_1, \langle \text{int} \rangle \text{ at } \rho_2, \dots) \rightarrow 0 \text{ at } r'$$

and suppose that v_1 and v_2 are integer tuples allocated in the same region r . Then f could not be called with v_1 and v_2 as arguments, because that would require

instantiating ρ_1 and ρ_2 with the same region. More generally, one could not type any function that takes two arguments that might or might not be allocated in the same region.

Approaches based on syntactic control of interference [Reynolds 1978; 1989] are more permissive than a linear type system, but are still too restrictive for our purposes; it is still impossible to instantiate multiple arguments with the same region.

Uniqueness. Our approach, instead of trying to *prevent* aliasing, is to use the type system to *track* aliasing. More precisely, we track *non*-aliasing, that is, uniqueness. We do this by tagging regions in capabilities with one of two *multiplicities*. The first form, $\{r^+\}$, is the capability to access region r as it has been understood heretofore. The second form, $\{r^1\}$, also permits accessing region r , but adds the additional information that r is unique; that is, r represents a different region from any other region appearing in a capability formed using $\{r^1\}$. For example, the capability $\{r_1^+, r_2^1\}$ not only indicates that it is permissible to access r_1 and r_2 , but also indicates that r_1 and r_2 represent distinct regions.

Since $\{r^1\}$ guarantees that r does not appear anywhere else in a capability formed using it, it is the capability, not just to access r , but also to free r . Thus we may type region deallocation with the rule

$$\frac{\Delta; \Gamma \vdash v : r \text{ handle} \quad \Delta \vdash C = C' \oplus \{r^1\} : \text{Cap}}{\Psi; \Delta; \Gamma; C \vdash \mathbf{freergn} \ v \Rightarrow \Delta; \Gamma; C'}$$

. Allocation of a region adds the new capability as unique:

$$\frac{\Psi; \Delta; \Gamma; C \vdash \mathbf{newrgn} \ \rho, x \Rightarrow \quad (\rho \notin \text{Dom}(\Delta), x \notin \text{Dom}(\Gamma))}{\Delta, \rho : \text{Rgn}; \Gamma \{x : \rho \text{ handle}\}; C \oplus \{\rho^1\}}$$

One of the main tasks of the type system is to ensure that uniqueness information is properly maintained; if a unique capability $\{r^1\}$ ever appears twice in the current capability then the **freergn** rule will not have the guarantee it requires and the system will be unsound. Immediately after initial allocation of a region, it is clear that a unique capability for that region does not appear twice: The typing rule for **newrgn** adds a single unique capability to the context. From this point forward, careful choice of the axioms for capability equality ensure subsequent instructions do not duplicate unique capabilities. In particular, although the capabilities $\{r^+\}$ and $\{r^+, r^+\}$ are considered equivalent, the capabilities $\{r^1\}$ and $\{r^1, r^1\}$ are not. More generally, we cannot prove the equation

$$\Delta \vdash C = C \oplus C : \text{Cap}$$

unless C contains no unique capabilities. Now, for example, if the current capability C contains one copy of $\{r^1\}$ before checking a **freergn** command, then the equation

$$\Delta \vdash C = C' \oplus \{r^1\} : \text{Cap}$$

that appears in the premise of the **freergn** rule cannot duplicate it. Hence, we can be certain that C' contains no privileges on the deallocated region r and other unique capabilities, say $\{r_2^1\}$, $\{r_3^1\}$, and $\{r_4^1\}$ that appear once in C , also appear exactly once in C' .

When C is equivalent to $C \oplus C$, we say that C is *duplicatable*. Note that capability variables are unduplicatable, since they can stand for any capability, including unduplicatable ones. Occasionally this prevents the typing of desired programs, so we provide a stripping operator \bar{C} that replaces all 1 multiplicities in C with + multiplicities. For example, $\{\bar{r}_1^1, \bar{r}_2^1\} = \{r_1^+, r_2^+\}$. For any capability C , the capability \bar{C} is duplicatable. When programs need an unknown but duplicatable capability, they may use a stripped variable \bar{c} . As you will see in Section 3, the stripping operator is essential in the translation of Tofte and Talpin’s region-based language into CL.

The complete rules for equivalence of capabilities and other constructors appear in Figure 4. Notice that the single rule eq-flag equates the duplicatable capability $\{r^+\}$ with the barred capability $\{\bar{r}^1\}$. Consequently, the form $\{r^+\}$ is redundant given the presence of the bar operator. However, the + notation is a pleasing foil for the 1 notation, and the two flags give us a convenient way to distinguish between regions that appear once and regions that potentially appear many times in a single capability.

Subcapabilities. The capabilities $\{r^1\}$ and $\{r^+\}$ are not the same, but the former should provide all the privileges of the latter. We therefore say that the former is a subcapability of the latter and write $\{r^1\} \leq \{r^+\}$. In the complete system, the various rules from Section 2.2 are modified to account for subcapabilities. For example, the function call rule becomes

$$\frac{\Delta \vdash C \leq C'' \oplus \{r^+\} \quad \Delta \vdash C \leq C' \quad \Psi; \Delta; \Gamma \vdash v : (C', \tau_1, \dots, \tau_n) \rightarrow 0 \text{ at } r \quad \Psi; \Delta; \Gamma \vdash v_i : \tau_i}{\Psi; \Delta; \Gamma; C \vdash v(v_1, \dots, v_n)}$$

. Suppose f has type $\forall[\rho_1:\text{Rgn}, \rho_2:\text{Rgn}].(\{\rho_1^+, \rho_2^+, \dots\} \rightarrow 0 \text{ at } r)$. If we hold capability $\{r^+\}$, we may call f by instantiating ρ_1 and ρ_2 with r , since $\{r^+\} = \{r^+, r^+\}$. Using the subcapability relation, we may also call f when we hold $\{r^1\}$, again by instantiating ρ_1 and ρ_2 with r , since $\{r^1\} \leq \{r^+\} = \{r^+, r^+\}$.

Figure 5 contains the subcapability rules. When reading these rules, remember that $\Delta \vdash \{\bar{r}^1\} = \{r^+\} : \text{Cap}$. We use this fact to derive the judgment $\Delta \vdash \{r^1\} \leq \{r^+\} : \text{Cap}$ that we discussed informally above:

$$\frac{\frac{\Delta \vdash \{\bar{r}^1\} = \{r^+\} : \text{Cap} \quad \Delta \vdash \{\bar{r}^1\} \leq \{r^+\} : \text{Cap}}{\Delta \vdash \{r^1\} \leq \{r^+\} : \text{Cap}} \text{ (sub-trans)}}{\Delta \vdash \{r^1\} \leq \{r^+\} : \text{Cap}} \text{ (sub-eq) (eq-flag)}$$

The subcapability relation accounts only for the forgetting of uniqueness information. Intuitively there could be a second source of subcapabilities, those generated by forgetting an entire capability. For example, $\{r_1^+, r_2^+\}$ seems to provide all the privileges of $\{r_1^+\}$, so it is reasonable to suppose $\{r_1^+, r_2^+\}$ to be a subcapability of $\{r_1^+\}$. Indeed, one can construct a sound CL incorporating this axiom, but we omit it because doing so allows us to specify memory management obligations and to prove a stronger property about space usage. Notice also, that by omitting this axiom, we do not give up any flexibility; one may always write a function that

$$\boxed{\Delta \vdash \Delta_1 = \Delta_2}$$

$$\frac{}{\Delta \vdash \cdot = \cdot} \text{ (ctxt-eq-empty)}$$

$$\frac{\Delta \vdash \Delta_1 = \Delta_2}{\Delta \vdash \Delta_1, \alpha : \kappa = \Delta_2, \alpha : \kappa} \text{ (}\alpha \notin \text{Dom}(\Delta \Delta_1)\text{) (ctxt-eq-kind)}$$

$$\frac{\Delta \vdash \Delta_1 = \Delta_2 \quad \Delta \Delta_1 \vdash C_1 = C_2 : \text{Cap}}{\Delta \vdash \Delta_1, \epsilon \leq C_1 = \Delta_2, \epsilon \leq C_2} \text{ (}\epsilon \notin \text{Dom}(\Delta \Delta_1)\text{) (ctxt-eq-bound)}$$

$$\boxed{\Delta \vdash c_1 = c_2 : \kappa}$$

$$\frac{\Delta \vdash c : \kappa}{\Delta \vdash c = c : \kappa} \text{ (eq-reflex)}$$

$$\frac{\Delta \vdash c_2 = c_1 : \kappa}{\Delta \vdash c_1 = c_2 : \kappa} \text{ (eq-symm)}$$

$$\frac{\Delta \vdash c_1 = c_2 : \kappa \quad \Delta \vdash c_2 = c_3 : \kappa}{\Delta \vdash c_1 = c_3 : \kappa} \text{ (eq-trans)}$$

$$\frac{\Delta \vdash C_1 = C'_1 : \text{Cap} \quad \Delta \vdash C_2 = C'_2 : \text{Cap}}{\Delta \vdash C_1 \oplus C_2 = C'_1 \oplus C'_2 : \text{Cap}} \text{ (eq-congruence-plus)}$$

$$\frac{\Delta \vdash C = C' : \text{Cap}}{\Delta \vdash \bar{C} = \bar{C}' : \text{Cap}} \text{ (eq-congruence-bar)}$$

$$\frac{\Delta \vdash C : \text{Cap}}{\Delta \vdash \emptyset \oplus C = C : \text{Cap}} \text{ (eq-}\emptyset\text{)}$$

$$\frac{\Delta \vdash C_1 : \text{Cap} \quad \Delta \vdash C_2 : \text{Cap}}{\Delta \vdash C_1 \oplus C_2 = C_2 \oplus C_1 : \text{Cap}} \text{ (eq-comm)}$$

$$\frac{\Delta \vdash C_i : \text{Cap} \quad (\text{for } 1 \leq i \leq 3)}{\Delta \vdash (C_1 \oplus C_2) \oplus C_3 = C_1 \oplus (C_2 \oplus C_3) : \text{Cap}} \text{ (eq-assoc)}$$

$$\frac{\Delta \vdash C : \text{Cap}}{\Delta \vdash \bar{C} = \bar{C} \oplus \bar{C} : \text{Cap}} \text{ (eq-dup)}$$

$$\frac{}{\Delta \vdash \bar{\emptyset} = \emptyset : \text{Cap}} \text{ (eq-bar-}\emptyset\text{)}$$

$$\frac{\Delta \vdash r : \text{Rgn}}{\Delta \vdash \{\bar{r}^1\} = \{r^+\} : \text{Cap}} \text{ (eq-flag)}$$

$$\frac{\Delta \vdash C : \text{Cap}}{\Delta \vdash \bar{\bar{C}} = \bar{C} : \text{Cap}} \text{ (eq-bar-idem)}$$

$$\frac{\Delta \vdash C_1 : \text{Cap} \quad \Delta \vdash C_2 : \text{Cap}}{\Delta \vdash \bar{C}_1 \oplus \bar{C}_2 = \overline{C_1 \oplus C_2} : \text{Cap}} \text{ (eq-distrib)}$$

Fig. 4. Capability static semantics: equality.

$\Delta \vdash C_1 \leq C_2$

$$\frac{\Delta \vdash C_1 = C_2 : \mathbf{Cap}}{\Delta \vdash C_1 \leq C_2} \text{ (sub-eq)}$$

$$\frac{\Delta \vdash C_1 \leq C_2 \quad \Delta \vdash C_2 \leq C_3}{\Delta \vdash C_1 \leq C_3} \text{ (sub-trans)}$$

$$\frac{\Delta \vdash C_1 \leq C'_1 \quad \Delta \vdash C_2 \leq C'_2}{\Delta \vdash C_1 \oplus C_2 \leq C'_1 \oplus C'_2} \text{ (sub-congruence-plus)}$$

$$\frac{\Delta \vdash C \leq C'}{\Delta \vdash \overline{C} \leq \overline{C'}} \text{ (sub-congruence-bar)}$$

$$\frac{}{\Delta \vdash \epsilon \leq C} ((\epsilon \leq C) \in \Delta) \text{ (sub-var)}$$

$$\frac{\Delta \vdash C : \mathbf{Cap}}{\Delta \vdash C \leq \overline{C}} \text{ (sub-bar)}$$

Fig. 5. Capability static semantics: Equality and subcapability relations.

can be called with extra capabilities by using a capability variable, as discussed in Section 2.2.

By omitting the axiom $C_1 \oplus C_2 \leq C_1$, our type system may formally specify who has responsibility for freeing a region. Failure to follow informal conventions is a common source of bugs in languages (such as C) that use manual memory management. Our type system rules out such bugs. For example, consider the type

$$\forall[\rho:\mathbf{Rgn}, \epsilon:\mathbf{Cap}].(\epsilon \oplus \{r^+, \rho^1\}, \rho \text{ handle}, (\epsilon \oplus \{r^+\}) \rightarrow 0 \text{ at } r) \rightarrow 0 \text{ at } r$$

. In our system $\epsilon \oplus \{r^+, \rho^1\} \not\leq \epsilon \oplus \{r^+\}$. Consequently, before any function with this type can return (i.e., call the continuation of type $(\epsilon \oplus \{r^+\}) \rightarrow 0 \text{ at } r$), it must take action to satisfy the capability $\epsilon \oplus \{r^+\}$, that is, it must free ρ .

In general, our type system prevents “region leaks”: programs must deallocate all memory regions if they terminate (Theorem 2.5). Therefore, the operating system does not have to clean up after a program halts. The typing rule for `halt` states that no capabilities may be held, and since capabilities may not be forgotten, this means that all regions must have been freed.

$$\frac{\Psi; \Delta; \Gamma \vdash v : \text{int} \quad \Delta \vdash C = \emptyset : \mathbf{Cap}}{\Psi; \Delta; \Gamma; C \vdash \text{halt } v}$$

The type system certainly does not prevent all forms of memory leaks. At times during the course of evaluation there may be dead regions that have not yet been deallocated or dead objects within live regions. However, the type system does provide a degree of static error checking and the possibility to encode certain memory management protocols.

Bounded Quantification. The system presented to this point is sound, but it is not yet sufficient for compiling real source languages. We need to be able to recover uniqueness after a region name is duplicated. To see why, suppose we hold

the capability $\{r^1\}$ and f has type

$$\forall[\rho_1:\text{Rgn}, \rho_2:\text{Rgn}].(\{\rho_1^+, \rho_2^+\}, \dots, (\{\rho_1^+, \rho_2^+\}, \dots) \rightarrow 0 \text{ at } \rho_1) \rightarrow 0 \text{ at } r$$

. We would like to be able to instantiate ρ_1 and ρ_2 with r (which we may do, since $\{r^1\} \leq \{r^+, r^+\}$), and then free r when f calls the continuation in its final argument. Unfortunately, the continuation only possesses the capability $\{r^+, r^+\} = \{r^+\}$, not the capability $\{r^1\}$ necessary to free r . It does not help to strengthen the capability of the continuation to (for example) $\{\rho_1^1\}$, because then f may not call it.

We *may* recover uniqueness information by quantifying a capability variable. Suppose we again hold capability $\{r^1\}$ and g has type

$$\forall[\rho_1:\text{Rgn}, \rho_2:\text{Rgn}, \epsilon:\text{Cap}].(\epsilon, \dots, (\epsilon, \dots) \rightarrow 0 \text{ at } \rho_1) \rightarrow 0 \text{ at } r$$

. We may instantiate ϵ with $\{r^1\}$, and then the continuation will possess that same capability, allowing it to free r . Unfortunately, the body of function g no longer has the capability to access ρ_1 and ρ_2 , since its type draws no connection between them and ϵ .

We solve this problem by using bounded quantification to relate ρ_1 , ρ_2 , and ϵ . Suppose h has type

$$\forall[\rho_1:\text{Rgn}, \rho_2:\text{Rgn}, \epsilon \leq \{\rho_1^+, \rho_2^+\}].(\epsilon, \dots, (\epsilon, \dots) \rightarrow 0 \text{ at } \rho_1) \rightarrow 0 \text{ at } r$$

. If we hold capability $\{r^1\}$, we may call h by instantiating ρ_1 and ρ_2 with r and instantiating ϵ with $\{r^1\}$. This instantiation is permissible because $\{r^1\} \leq \{r^+, r^+\}$. As with g , the continuation will possess the capability $\{r^1\}$, allowing it to free r , but the body of h (like that of f) will have the capability to access ρ_1 and ρ_2 , since $\epsilon \leq \{\rho_1^+, \rho_2^+\}$.

Bounded quantification solves the problem by revealing some information about a capability ϵ , while still requiring the function to be parametric over ϵ . Hence, when the function calls its continuation we regain the stronger capability (to free r), although that capability was temporarily hidden in order to duplicate r . More generally, bounded quantification allows us to hide some privileges when calling a function, and regain those privileges in its continuation. Thus, we support statically checkable attenuation and amplification of capabilities.

Static Semantics So Far. Together, parametric polymorphism, bounded parametric polymorphism, and notions of uniqueness and aliasing provide a flexible language for expressing the lifetimes of regions. Figures 6 and 7 formally summarize the rules for type checking instructions and values that depend upon these concepts. We have already explained the majority of these rules in previous sections, and the rules that we have not yet specified are the obvious ones (integers are given type `int`, etc.). Notice, however, that the form of the judgment for heap values h is slightly different from the judgments for instructions and small values v . The judgment $\Psi; \Delta; \Gamma \vdash h \text{ at } r : \tau$ states that when memory has type Ψ , free constructor variables have kinds (and bounds) given by Δ and free value variables have types given by Γ , the heap value h resides in region r and has type τ .

2.4 The Static Semantics of the Abstract Machine

We have described the type constructor language of CL and the typing rules for the main term-level constructs. In fact, the previous section contains all of the

$\Psi; \Delta; \Gamma \vdash h \text{ at } r : \tau$

$$\frac{\Delta \vdash \tau_f : \text{Type} \quad \Psi; \Delta \Delta'; \Gamma \{f: \tau_f, x_1: \tau_1, \dots, x_n: \tau_n\}; C \vdash e \quad \left(\begin{array}{l} \tau_f = \forall[\Delta']. (C, \tau_1, \dots, \tau_n) \rightarrow \mathbf{0} \text{ at } r \\ f, x_1, \dots, x_n \notin \text{Dom}(\Gamma) \end{array} \right)}{\Psi; \Delta; \Gamma \vdash \mathbf{fix} f[\Delta'](C, x_1: \tau_1, \dots, x_n: \tau_n). e \text{ at } r : \tau_f} \text{ (h-fix)}$$

$$\frac{\Psi; \Delta; \Gamma \vdash v_i : \tau_i \quad (\text{for } 1 \leq i \leq n) \quad \Delta \vdash r : \text{Rgn}}{\Psi; \Delta; \Gamma \vdash \langle v_1, \dots, v_n \rangle \text{ at } r : \langle \tau_1, \dots, \tau_n \rangle \text{ at } r} \text{ (h-tuple)}$$

$$\frac{\Psi; \Delta; \Gamma \vdash h \text{ at } r : \tau' \quad \Delta \vdash \tau' = \tau : \text{Type}}{\Psi; \Delta; \Gamma \vdash h \text{ at } r : \tau} \text{ (h-eq)}$$

$\Psi; \Delta; \Gamma \vdash v : \tau$

$$\frac{}{\Psi; \Delta; \Gamma \vdash x : \tau} (\Gamma(x) = \tau) \text{ (v-var)}$$

$$\frac{}{\Psi; \Delta; \Gamma \vdash i : \text{int}} \text{ (v-int)}$$

$$\frac{\Psi; \Delta; \Gamma \vdash v : \forall[\alpha: \kappa, \Delta']. (C, \tau_1, \dots, \tau_n) \rightarrow \mathbf{0} \text{ at } r \quad \Delta \vdash c : \kappa}{\Psi; \Delta; \Gamma \vdash v[c] : (\forall[\Delta']. (C, \tau_1, \dots, \tau_n) \rightarrow \mathbf{0})[c/\alpha] \text{ at } r} \text{ (v-type)}$$

$$\frac{\Psi; \Delta; \Gamma \vdash v : \forall[\epsilon \leq C'', \Delta']. (C', \tau_1, \dots, \tau_n) \rightarrow \mathbf{0} \text{ at } r \quad \Delta \vdash C \leq C''}{\Psi; \Delta; \Gamma \vdash v[C] : (\forall[\Delta']. (C', \tau_1, \dots, \tau_n) \rightarrow \mathbf{0})[C/\epsilon] \text{ at } r} \text{ (v-sub)}$$

$$\frac{\Psi; \Delta; \Gamma \vdash v : \tau' \quad \Delta \vdash \tau' = \tau : \text{Type}}{\Psi; \Delta; \Gamma \vdash v : \tau} \text{ (v-eq)}$$

Fig. 6. Capability static semantics: Heap and word values.

information programmers or compilers require to write type-safe programs in CL. However, in order to prove a type soundness result in the style of Wright and Felleisen [Wright and Felleisen 1994], we must be able to type check programs at every step during their evaluation. In this section, we give the static semantics of the run-time values that are not normally manipulated by programmers, but are nevertheless necessary to prove our soundness result.

At first, the formal definition of the semantics may appear quite complex because we use a number of different judgment forms. However, most of these forms follow naturally from the development of previous sections and other work on allocation semantics by Morrisett et al. [1995; 1997; 1998]. The extra complexity in the definition of the language will pay off when we come to prove type soundness: all of the main invariants are expressed directly in the typing rules and therefore most of the proof follows from straightforward inductions over these rules.

Figure 8 specifies the rules for typing memory, most of which are straightforward. The judgments $\vdash \Psi$ and $\vdash \Upsilon$ specify when memory types and region types are well-formed. When all elements in a sequence x_1, \dots, x_n are different, we write x_1, \dots, x_n distinct. Memory and region types (as well as regions and memory themselves) are finite partial maps so all elements in the domain of the map are distinct. The typing judgments for memory and region types have side conditions

$\Psi; \Delta; \Gamma; C \vdash d \Rightarrow \Delta'; \Gamma'; C'$

$$\frac{\Psi; \Delta; \Gamma \vdash v : \tau}{\Psi; \Delta; \Gamma; C \vdash x = v \Rightarrow \Delta; \Gamma\{x:\tau\}; C} \quad (x \notin \text{Dom}(\Gamma)) \text{ (val)}$$

$$\frac{\Psi; \Delta; \Gamma \vdash v_1 : \text{int} \quad \Psi; \Delta; \Gamma \vdash v_2 : \text{int}}{\Psi; \Delta; \Gamma; C \vdash x = v_1 \ p \ v_2 \Rightarrow \Delta; \Gamma\{x:\text{int}\}; C} \quad (x \notin \text{Dom}(\Gamma)) \text{ (prim)}$$

$$\frac{\Psi; \Delta; \Gamma \vdash v : r \ \text{handle} \quad \Psi; \Delta; \Gamma \vdash h \ \text{at} \ r : \tau \quad \Delta \vdash C \leq C' \oplus \{r^+\}}{\Psi; \Delta; \Gamma; C \vdash x = h \ \text{at} \ v \Rightarrow \Delta; \Gamma\{x:\tau\}; C} \quad (x \notin \text{Dom}(\Gamma)) \text{ (hval)}$$

$$\frac{\Psi; \Delta; \Gamma \vdash v : \langle \tau_1, \dots, \tau_n \rangle \ \text{at} \ r \quad \Delta \vdash C \leq C' \oplus \{r^+\}}{\Psi; \Delta; \Gamma; C \vdash x = \pi_i v \Rightarrow \Delta; \Gamma\{x:\tau_i\}; C} \quad (x \notin \text{Dom}(\Gamma) \wedge 1 \leq i \leq n) \text{ (proj)}$$

$$\frac{\Psi; \Delta; \Gamma; C \vdash \text{newrgn} \ \rho, x \Rightarrow \Delta\{\rho:\text{Rgn}\}; \Gamma\{x:\rho \ \text{handle}\}; C \oplus \{\rho^1\}}{\left(\begin{array}{l} \rho \notin \text{Dom}(\Delta) \\ x \notin \text{Dom}(\Gamma) \end{array} \right) \text{ (newrgn)}}$$

$$\frac{\Psi; \Delta; \Gamma \vdash v : r \ \text{handle} \quad \Delta \vdash C = C' \oplus \{r^1\} : \text{Cap}}{\Psi; \Delta; \Gamma; C \vdash \text{freergn} \ v \Rightarrow \Delta; \Gamma; C'} \text{ (freergn)}$$

$\Psi; \Delta; \Gamma; C \vdash e$

$$\frac{\Psi; \Delta; \Gamma; C \vdash d \Rightarrow \Delta'; \Gamma'; C' \quad \Psi; \Delta'; \Gamma'; C' \vdash e}{\Psi; \Delta; \Gamma; C \vdash \text{let} \ d \ \text{in} \ e} \text{ (letdec)}$$

$$\frac{\Psi; \Delta; \Gamma \vdash v : \text{int} \quad \Psi; \Delta; \Gamma; C \vdash e_2 \quad \Psi; \Delta; \Gamma; C \vdash e_3}{\Psi; \Delta; \Gamma; C \vdash \text{if} \ 0 \ \text{then} \ e_2 \ \text{else} \ e_3} \text{ (if)}$$

$$\frac{\Psi; \Delta; \Gamma \vdash v : (C', \tau_1, \dots, \tau_n) \rightarrow \mathbf{0} \ \text{at} \ r \quad \Psi; \Delta; \Gamma \vdash v_i : \tau_i \quad (\text{for } 1 \leq i \leq n) \quad \Delta \vdash C \leq C'' \oplus \{r^+\} \quad \Delta \vdash C \leq C'}{\Psi; \Delta; \Gamma; C \vdash v(v_1, \dots, v_n)} \text{ (app)}$$

$$\frac{\Psi; \Delta; \Gamma \vdash v : \text{int} \quad \Delta \vdash C = \emptyset : \text{Cap}}{\Psi; \Delta; \Gamma; C \vdash \text{halt} \ v} \text{ (halt)}$$

Fig. 7. Capability static semantics: Declarations and expressions.

to this effect.

The judgment $\vdash M : \Psi$ states that memory M is described by Ψ , and the judgment $\Psi \vdash R \ \text{at} \ \nu : \Upsilon$ states that region R with name ν is described by Υ . Informally, these judgments ensure that for addresses $\nu.\ell$, $\Psi(\nu.\ell)$ is type τ if and only if the memory M described by Ψ contains a value v at address $\nu.\ell$ that has type τ .

The next judgment, $\Psi \vdash C \ \text{sat}$, is called the satisfiability judgment, and it formalizes the connection between the static capability and the run-time state of memory. Clearly, the current capability must not contain any regions that are not in memory; this could lead to a runtime error. However, it is equally important that memory not contain regions for which we have no capability, as such regions can never be freed. Consequently, satisfiability ensures that at any time during

execution of the abstract machine, our capability is equal to $\{\nu_i^{\varphi_i}\}$ where each ν_i occurs exactly once and is present in the current memory. Furthermore, by virtue of the fact that $\cdot \vdash \{\nu^1\} \neq \{\nu^1, \nu^1\} : \text{Cap}$, no unique regions may appear more than once in C . Each of these properties are essential to ensure that regions are used safely.

Figure 9 contains rules for small values that only appear at run time (addresses and region handles). The rules for typing an address $\nu.l$ are quite unusual, but crucial to the type soundness proof. The first rule, v-addr, is used during the lifetime of the region ν : if the region ν is in memory then ν will also be in the domain of the memory type Ψ . Therefore rule v-addr applies, and $\nu.l$ will have type $\Psi(\nu.l)$. Now consider some point in the computation after the region ν has been deallocated. The region ν is no longer in the memory, but the addresses $\nu.l$ may still appear embedded in tuples or closures allocated in other regions, and, therefore, it must be given a type. If a region ν does not appear in memory type Ψ , the type system has the flexibility to give $\nu.l$ any function type (by rule v-addr-arrow) or tuple type (by rule v-addr-tuple).

At first glance, these rules would appear to lead to unsoundness: the address $\nu.l$ is a dangling pointer, and it may be given a valid type. Fortunately, though, capabilities prevent anything from going wrong. The satisfiability judgment ensures that programs only ever possess capabilities for regions that appear in memory, and, as we explained earlier, programs can only access the regions for which they have capabilities. Consequently, a dangling pointer may be given a valid tuple or function type, but capabilities prevent it from being accessed.

We now have all components necessary to define a well-formed machine state. The state (M, e) is well-formed if the memory M can be described by a well-formed heap type Ψ , there exists a capability C such that C satisfies the heap type Ψ , and finally, the expression e is well-formed with respect to Ψ and C :

$$\frac{\vdash M : \Psi \quad \Psi \vdash C \text{ sat} \quad \Psi; \cdot; \cdot; C \vdash e \text{ (program)}}{\vdash (M, e)}$$

2.5 Formal Properties

Type Soundness. The most important property of CL is *Type Soundness*. Type Soundness states that a program will never enter a *stuck state* during execution. A state (M, e) is *stuck* if there does not exist (M', e') such that $(M, e) \mapsto (M', e')$ and e is not `halt` i . For example, a state that tries to project a value from a tuple that does not appear in memory is stuck.

THEOREM (TYPE SOUNDNESS). *If $\vdash (M, e)$ and $(M, e) \mapsto^* (M', e')$ then (M', e') is not stuck.*

In the previous sections of this article, we have explained how to type memory, how to relate the memory typing to static capabilities, and finally, given a collection of capabilities, how the rules for typing expressions prevent unsafe accesses to the store. These invariants are the main elements in the formal proof of soundness. However, there are many details to fill in. The proof is in the style of Wright and Felleisen [Wright and Felleisen 1994] and uses the standard Type Preservation

$\vdash \Upsilon$	$\frac{\cdot \vdash \tau_i \quad (\text{for } 1 \leq i \leq n)}{\vdash \{\ell_1 : \tau_1, \dots, \ell_n : \tau_n\}} \quad (\ell_1, \dots, \ell_n \text{ distinct}) \text{ (region-type)}$
$\vdash \Psi$	$\frac{\vdash \Upsilon_i \quad (\text{for } i \leq i \leq n)}{\vdash \{\nu_1 : \Upsilon_1, \dots, \nu_n : \Upsilon_n\}} \quad (\nu_1, \dots, \nu_n \text{ distinct}) \text{ (memory-type)}$
$\Psi \vdash R \text{ at } \nu : \Upsilon$	$\frac{\Psi; \cdot \vdash h_i \text{ at } \nu : \tau_i \quad (\text{for } 1 \leq i \leq n)}{\Psi \vdash \{\ell_1 \mapsto h_1, \dots, \ell_n \mapsto h_n\} \text{ at } \nu : \{\ell_1 : \tau_1, \dots, \ell_n : \tau_n\}} \quad (\ell_1, \dots, \ell_n \text{ distinct}) \text{ (region)}$
$\vdash M : \Psi$	$\frac{\vdash \Psi \quad \Psi \vdash R_i \text{ at } \nu_i : \Upsilon_i \quad (\text{for } 1 \leq i \leq n)}{\vdash \{\nu_1 \mapsto R_1, \dots, \nu_n \mapsto R_n\} : \Psi} \quad \left(\begin{array}{l} \Psi = \{\nu_1 : \Upsilon_1, \dots, \nu_n : \Upsilon_n\} \\ \nu_1, \dots, \nu_n \text{ distinct} \end{array} \right) \text{ (memory)}$
$\Psi \vdash C \text{ sat}$	$\frac{\cdot \vdash C = \{\nu_1^{\varphi_1}, \dots, \nu_n^{\varphi_n}\} : \text{Cap}}{\{\nu_1 : \Upsilon_1, \dots, \nu_n : \Upsilon_n\} \vdash C \text{ sat}} \quad (\nu_1, \dots, \nu_n \text{ distinct}) \text{ (sat)}$

Fig. 8. Capability static semantics: Memory.

$\Psi; \Delta; \Gamma \vdash v : \tau$	$\frac{}{\Psi; \Delta; \Gamma \vdash \nu.\ell : \tau} \quad (\Psi(\nu.\ell) = \tau) \text{ (v-addr)}$
	$\frac{\Delta \vdash \langle \tau_1, \dots, \tau_n \rangle \text{ at } \nu : \text{Type}}{\Psi; \Delta; \Gamma \vdash \nu.\ell : \langle \tau_1, \dots, \tau_n \rangle \text{ at } \nu} \quad (\nu \notin \text{Dom}(\Psi)) \text{ (v-addr-tuple)}$
	$\frac{\Delta \vdash \forall[\Delta']. (C, \tau_1, \dots, \tau_n) \rightarrow 0 \text{ at } \nu : \text{Type}}{\Psi; \Delta; \Gamma \vdash \nu.\ell : \forall[\Delta']. (C, \tau_1, \dots, \tau_n) \rightarrow 0 \text{ at } \nu} \quad (\nu \notin \text{Dom}(\Psi)) \text{ (v-addr-arrow)}$
	$\frac{}{\Psi; \Delta; \Gamma \vdash \text{handle}(\nu) : \nu \text{ handle}} \quad \text{(v-handle)}$

Fig. 9. Capability static semantics: Run-time values.

and Progress lemmas. Progress states that well-typed states are not stuck, and Preservation states that evaluation steps preserve well-typedness.

LEMMA (TYPE PRESERVATION). *If $\vdash (M, e)$ and $(M, e) \mapsto (M', e')$ then $\vdash (M', e')$*

LEMMA (PROGRESS). *If $\vdash (M, e)$ then either:*

- (1) *There exists (M', e') such that $(M, e) \mapsto (M', e')$ or*
- (2) *$e = \text{halt } i$.*

Because of the length and tedium of the proofs of these lemmas, we have moved them, along with the proof of soundness itself, to Appendix A.

Complete Collection. The second important property of the language is that well-typed terminating programs return all of their memory resources to the system before they halt. We call this property *Complete Collection*.

THEOREM (COMPLETE COLLECTION). *If $\vdash (M, e)$ then either (M, e) diverges or $(M, e) \mapsto^* (\{\}, \mathbf{halt} \ i)$.*

By Subject Reduction and Progress, terminating programs end in well-formed machine states $(M, \mathbf{halt} \ i)$. The typing rule for the `halt` expression requires that the capability C be empty. Using this fact, we can infer that the memory M contains no regions. Appendix A also contains a formal proof of this theorem.

Decidability. A third property of interest for many typed languages is decidability. However, while we strongly believe that CL as presented is decidable, we have not studied this property in formal detail. We are less concerned with decidability than we are with type soundness or the expressiveness of the language because type checking can always be made decidable by providing additional annotations that direct reconstruction of the full typing derivation. In a user-level programming language, excessive type annotations will hinder programmer productivity, but compilers have a much higher tolerance for handling such tedious details. In fact, Necula and Lee [Necula 1997; Necula and Lee 1998] in their work on Proof-Carrying Code, suggest that compilers produce full proofs of type safety encoded in a first-order predicate logic. They argue that proof checking (as opposed to proof inference) can be engineered to be quite fast and that the trusted computing base for a proof-checker is somewhat smaller than for a sophisticated type reconstruction algorithm (i.e., theorem prover).

Of course, there must be some way for the compiler to obtain the full typing derivation if the proof-carrying code methodology is to be used. Therefore, in the next section, we describe a translation from a variant of Tofte and Talpin’s region-based language into CL. Our proof that the translation produces well-typed code is constructive (the proof builds a typing derivation for the output code), and it makes it possible to fully annotate and subsequently verify compiler output.

3. FROM REGIONS TO CAPABILITIES

The first part of this section develops a high-level type-and-effect system for regions based on the work of Tofte and Talpin [1994; 1997]. The second part of this section defines a formal translation from the Tofte-Talpin language into CL. By composing the translation with Tofte and Birkedal’s region inference technology [Tofte and Birkedal 1998], we may obtain a type-preserving compiler front-end.

3.1 A Region-Based Calculus

Preliminaries. The source language for our compiler is the region-based calculus shown in Figure 10. This language is an explicitly typed variant of the calculus first presented by Tofte and Talpin [1994; 1997]. Like CL, it contains integers, tuples, and functions. However, unlike the CL, allocation and deallocation of regions is combined in a single construct: `letregion ρ, x_ρ in e` . This construct allocates a new region ρ and places the handle for that region in the term variable x_ρ . Next, it executes the expression e , and finally, the region ρ is deallocated. As discussed in

<i>kinds</i>	$\kappa ::= \text{Type} \mid \text{Rgn} \mid \text{Eff}$
<i>constructor variables</i>	$\alpha, \beta, \rho, \epsilon$
<i>constructor contexts</i>	$\Delta ::= \cdot \mid \Delta, \alpha : \kappa_1$
<i>type schemes</i>	$\sigma ::= \tau \mid \forall[\Delta]. \tau_1 \xrightarrow{\psi} \tau_2 \text{ at } r$
<i>constructors</i>	$c, \tau, r, \psi ::= \alpha \mid \text{int} \mid r \text{ handle} \mid \langle \tau_1, \dots, \tau_n \rangle \text{ at } r \mid \tau_1 \xrightarrow{\psi} \tau_2 \text{ at } r \mid \emptyset \mid \{r\} \mid \psi_1 \cup \psi_2$
<i>term variables</i>	x, f
<i>term contexts</i>	$\Gamma ::= \cdot \mid \Gamma, x : \sigma$
<i>terms</i>	$e ::= x \mid i \mid e_1 \text{ p } e_2 \mid \text{if } 0 e_1 \text{ then } e_2 \text{ else } e_3 \mid \langle e_1, \dots, e_n \rangle \text{ at } e_{n+1} \mid \pi_i e \mid \text{letrec } f[\Delta](x) : \sigma \text{ at } e_1 = e_2 \text{ in } e_3 \mid f[c_1, \dots, c_n] \mid e_1 e_2 \mid \text{letregion } \rho, x_\rho \text{ in } e$

Fig. 10. Region syntax.

the introduction, this lexically scoped in construct is not as flexible as the separate `newrgn` and `freergn` constructs provided by CL. The main goal of this section is to show how to compile `letregion` expressions into these lower-level primitives.

As in the original Tofte-Talpin calculus, the region language has prenex predicative polymorphism. The term `letrecf` $[\Delta](x) : \sigma \text{ at } e_1 = e_2 \text{ in } e_3$ allocates a closure f of polytype σ . The closure is polymorphic over its type context Δ , which may contain ordinary type variables as well as region variables and effect variables (explained below). The closure is allocated in the region r if the expression e_1 evaluates to a region handle for r . The expression e_2 describes the body of the function.

Unlike previous work on region-based type systems, we treat all type constructors, including region constructors, as compile-time-only objects. Therefore, the term $f[c_1, \dots, c_n]$, which denotes type application, has no runtime effect. During type checking, the type scheme for the polymorphic function f is instantiated with the types c_1, \dots, c_n to obtain the resultant type for the expression, but the dynamic semantics of the program (not shown here) do not depend upon these types. Hence, the types may be erased before the program is run without affecting the computation. As in CL, the data structures that are required to allocate and deallocate regions are treated as ordinary values of type `handle`(r).

Types and Effects. The main interest of the type constructor portion of the region language is the presence of *effects*. Effects, like capabilities, are used to control a program’s access to regions and, in particular, to prevent access to regions that have been deallocated. Intuitively, the effect of a term is the set of regions that the term accesses. Formally, an effect is either the empty effect (\emptyset), an effect variable (ϵ), a singleton set ($\{r\}$), or the union of two effects ($\psi_1 \cup \psi_2$). The \cup operator is associative, commutative, and idempotent, and \emptyset is the unit for the union operator. We write $\Delta \vdash_R \psi_1 = \psi_2 : \text{Eff}$ for equality on effects, and we use the abbreviation $\Delta \vdash_R \psi_1 \subseteq \psi_2$ when $\Delta \vdash_R \psi_1 \cup \psi_3 = \psi_2 : \text{Eff}$ for some effect ψ_3 .

All functions have *latent effects* that are incurred when the function is called and its body is executed. The effect that appears on arrow types ($\forall[\Delta]. \tau \xrightarrow{\psi} \tau'$)⁵

⁵Tofte and Talpin require each arrow type be annotated with an “arrow effect,” which is con-

$\Delta \vdash_R \Delta'$

$$\frac{}{\Delta \vdash_R \cdot} \quad \frac{\Delta \vdash_R \Delta'}{\Delta \vdash_R \Delta', \alpha : \kappa} \quad (\alpha \notin \text{Dom}(\Delta))$$

$\Delta \vdash_R \sigma$

$$\frac{\frac{\Delta \vdash_R \tau : \text{Type}}{\Delta \vdash_R \tau} \quad \frac{\frac{\vdash_R \Delta \Delta' \quad \Delta \Delta' \vdash_R \tau_i : \text{Type} \quad (\text{for } 1 \leq i \leq n+1)}{\Delta \Delta' \vdash_R \psi : \text{Eff}} \quad \Delta \vdash_R r : \text{Rgn}}{\Delta \vdash_R \forall[\Delta'].(\tau_1, \dots, \tau_n) \xrightarrow{\psi} \tau_{n+1} \text{ at } r}}$$

$\Delta \vdash_R c : \kappa$

$$\frac{}{\Delta \vdash_R \alpha : \kappa} \quad (\Delta(\alpha) = \kappa) \quad \frac{}{\Delta \vdash_R \text{int} : \text{Type}} \quad \frac{\Delta \vdash_R r : \text{Rgn}}{\Delta \vdash_R r \text{ handle} : \text{Type}}$$

$$\frac{\frac{\Delta \vdash_R \tau_i : \text{Type} \quad (\text{for } 1 \leq i \leq n) \quad \Delta \vdash_R r : \text{Rgn}}{\Delta \vdash_R \langle \tau_1, \dots, \tau_n \rangle \text{ at } r : \text{Type}} \quad \frac{\frac{\Delta \vdash_R r : \text{Rgn} \quad \Delta \vdash_R \psi : \text{Eff}}{\Delta \vdash_R \tau_i : \text{Type} \quad (\text{for } 1 \leq i \leq n+1)}}{\Delta \vdash_R (\tau_1, \dots, \tau_n) \xrightarrow{\psi} \tau_{n+1} \text{ at } r : \text{Type}}}{\frac{}{\Delta \vdash_R \emptyset : \text{Eff}} \quad \frac{\Delta \vdash_R r : \text{Rgn}}{\Delta \vdash_R \{r\} : \text{Eff}} \quad \frac{\Delta \vdash_R \psi_1 : \text{Eff} \quad \Delta \vdash_R \psi_2 : \text{Eff}}{\Delta \vdash_R \psi_1 \cup \psi_2 : \text{Eff}}}$$

Fig. 11. Region type formation.

$\Delta \vdash_R c_1 = c_2 : \kappa$

$$\frac{\frac{\Delta \vdash_R c : \kappa}{\Delta \vdash_R c = c : \kappa} \quad \frac{\Delta \vdash_R c_2 = c_1 : \kappa}{\Delta \vdash_R c_1 = c_2 : \kappa} \quad \frac{\Delta \vdash_R c_1 = c_2 : \kappa \quad \Delta \vdash_R c_2 = c_3 : \kappa}{\Delta \vdash_R c_1 = c_3 : \kappa}}$$

$$\frac{\Delta \vdash_R \psi : \text{Eff}}{\Delta \vdash_R \emptyset \cup \psi = \psi : \text{Eff}} \quad \frac{\Delta \vdash_R \psi_1 : \text{Eff} \quad \Delta \vdash_R \psi_2 : \text{Eff}}{\Delta \vdash_R \psi_1 \cup \psi_2 = \psi_2 \cup \psi_1 : \text{Eff}}$$

$$\frac{\Delta \vdash_R \psi_i : \text{Eff} \quad (\text{for } 1 \leq i \leq 3)}{\Delta \vdash_R (\psi_1 \cup \psi_2) \cup \psi_3 = \psi_1 \cup (\psi_2 \cup \psi_3) : \text{Eff}} \quad \frac{\Delta \vdash_R \psi_1 : \text{Eff} \quad \Delta \vdash_R \psi_2 : \text{Eff}}{\Delta \vdash_R \psi_1 \cup \psi_2 = \psi_1 \cup (\psi_2 \cup \psi_2) : \text{Eff}}$$

$\Delta \vdash_R \psi_1 \subseteq \psi_2$

$$\frac{\Delta \vdash_R \psi_1 \cup \psi_3 = \psi_2 : \text{Eff}}{\Delta \vdash_R \psi_1 \subseteq \psi_2}$$

Fig. 12. Effect equality and subset.

specifies the set of regions that a function of that type may access when it is invoked.

The rules for type constructor formation may be found in Figure 11. Effect equality and subset relations are summarized formally in Figure 12. Equality on types is syntactic up to renaming of bound variables and modulo equality of effects.

Static Semantics. The static semantics (Figure 13) for terms use a judgment of the form $\Delta; \Gamma \vdash_R e : \tau, \psi$ to track the effects produced by each expression. This

strained to have the form $\epsilon \cup \psi$. The type variable ϵ is used to name the effect and plays a role in their inference system. Because we are interested in type checking rather than type inference, we do not need to name the effects on arrows.

$$\begin{array}{c}
 \frac{}{\Delta; \Gamma \vdash_R x : \tau, \emptyset} \quad (\Gamma(x) = \tau) \quad \frac{}{\Delta; \Gamma \vdash_R i : \text{int}, \emptyset} \\
 \frac{\Delta; \Gamma \vdash_R e_1 : \text{int}, \psi_1 \quad \Delta; \Gamma \vdash_R e_2 : \text{int}, \psi_2}{\Delta; \Gamma \vdash_R e_1 p e_2 : \text{int}, \psi_1 \cup \psi_2} \\
 \frac{\Delta; \Gamma \vdash_R e_i : \tau_i, \psi_i \quad (\text{for } 1 \leq i \leq n) \quad \Delta; \Gamma \vdash_R e_{n+1} : r \text{ handle}, \psi_{n+1}}{\Delta; \Gamma \vdash_R \langle e_1, \dots, e_n \rangle \text{ at } e_{n+1} : \tau, \psi_1 \cup \dots \cup \psi_{n+1} \cup \{r\}} \\
 \frac{\Delta; \Gamma \vdash_R e : \langle \tau_1, \dots, \tau_n \rangle \text{ at } r, \psi}{\Delta; \Gamma \vdash_R \pi_i e : \tau_i, \psi \cup \{r\}} \quad (1 \leq i < n) \\
 \frac{\Delta; \Gamma \vdash_R e_1 : \text{int}, \psi_1 \quad \Delta; \Gamma \vdash_R e_2 : \tau, \psi_2 \quad \Delta; \Gamma \vdash_R e_3 : \tau, \psi_3}{\Delta; \Gamma \vdash_R \text{if0 } e_1 \text{ then } e_2 \text{ else } e_3 : \tau, \psi_1 \cup \psi_2 \cup \psi_3} \\
 \frac{\Delta \vdash_R \sigma \quad \Delta \Delta'; \Gamma\{f:\sigma, x:\tau_1\} \vdash_R e_2 : \tau_2, \psi \quad \Delta; \Gamma \vdash_R e_1 : r \text{ handle}, \psi_1 \quad \Delta; \Gamma\{f:\sigma\} \vdash_R e_3 : \tau_3, \psi_3}{\Delta; \Gamma \vdash_R \text{letrec } f[\Delta'](x) : \sigma \text{ at } e_1 = e_2 \text{ in } e_3 : \tau_3, \psi_1 \cup \psi_3 \cup \{r\}} \quad \left(\begin{array}{l} x, f \notin \text{Dom}(\Gamma) \\ \sigma = \forall[\Delta']. \tau_1 \xrightarrow{\psi} \tau_2 \text{ at } r \end{array} \right) \\
 \frac{(\Gamma(f) = \forall[\alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n]. \tau_1 \xrightarrow{\psi} \tau_2 \text{ at } r) \quad \Delta \vdash_R c_i : \kappa_i \quad (\text{for } 1 \leq i \leq n)}{\Delta; \Gamma \vdash_R f[c_1, \dots, c_n] : (\tau_1 \xrightarrow{\psi} \tau_2)[c_1, \dots, c_n/\alpha_1, \dots, \alpha_n] \text{ at } r, \emptyset} \\
 \frac{\Delta; \Gamma \vdash_R e_1 : \tau_1 \xrightarrow{\psi} \tau_2 \text{ at } r, \psi_1 \quad \Delta; \Gamma \vdash_R e_2 : \tau_1, \psi_2}{\Delta; \Gamma \vdash_R e_1 e_2 : \tau_2, \psi_1 \cup \psi_2 \cup \psi \cup \{r\}} \\
 \frac{\Delta\{\rho:\text{Rgn}\}; \Gamma\{x:\rho \text{ handle}\} \vdash_R e : \tau, \psi}{\Delta; \Gamma \vdash_R \text{letregion } \rho, x_\rho \text{ in } e : \tau, \psi \setminus \{\rho\}} \quad \left(\begin{array}{l} \rho \notin \text{ftv}(\tau) \cup \text{Dom}(\Delta) \\ x_\rho \notin \text{Dom}(\Gamma) \end{array} \right) \\
 \frac{\Delta; \Gamma \vdash_R e : \tau, \psi \quad \Delta \vdash_R \tau = \tau' : \text{Type} \quad \Delta \vdash_R \psi \subseteq \psi'}{\Delta; \Gamma \vdash_R e : \tau', \psi'}
 \end{array}$$

Fig. 13. Region term static semantics.

judgment states that under the type context Δ and the value context Γ , a term e has type τ and produces effect ψ . For example, the rule for projection states that if e is an expression with type $\langle \tau_1, \dots, \tau_n \rangle \text{ at } r$ and produces effect ψ then a projection $\pi_i e$ produces the effect $\psi \cup \{r\}$. The projection operation reads from the region r , and the subexpression e may read from or write to any of the regions in ψ . Hence the resulting effect must be the union of the two.

The rules involving functions are more complex. First, consider a function call $e_1 e_2$. Assume that e_1 generates the effect ψ_1 and evaluates to a closure of type $(\tau) \xrightarrow{\psi} \tau' \text{ at } r$, and that e_2 produces the effect ψ_2 and has type τ . After both expressions have been evaluated, the code for the function is projected from a closure that resides in region r . Now, because the function itself produces the effect ψ , the overall effect of the call is the union of $\{r\}$ with ψ , ψ_1 , and ψ_2 . In contrast to the value application rule, the rule for type application produces no effect; remember, types are erased before an expression is executed. Finally, examine the rule for the **letrec** term. There are three components to the effect produced by this expression: ψ_1 , the effect of evaluating the handle expression; $\{r\}$,

```

letregion  $\rho_1, x_{\rho_1}$  in
  letregion  $\rho_2, x_{\rho_2}$  in
    letrec count [ $\rho$ ] ( $x_{\rho} : \rho$  handle,  $x : \langle \text{int} \rangle$  at  $\rho$ ) at  $x_{\rho_1} : \sigma_{\text{count}} =$ 
      let  $n = \pi_1(x)$  in                                     % (1)
        if0  $n$ 
          then ()
          else count [ $\rho$ ] ( $x_{\rho}, \langle n - 1 \rangle$  at  $x_{\rho}$ ) % (2)
        % end
      in
        count [ $\rho_2$ ] ( $x_{\rho_2}, \langle 10 \rangle$  at  $x_{\rho_2}$ )
      % end letrec
    % end region  $\rho_2$  scope and deallocate
  % end region  $\rho_1$  scope and deallocate

where  $\sigma_{\text{count}} = \forall[\rho]. (\rho$  handle,  $\langle \text{int} \rangle$  at  $\rho$ )  $\xrightarrow{\{\rho_1, \rho\}}$  unit

```

Fig. 14. Count in the region calculus.

the effect of writing the closure data structure in region r ; and ψ_3 , the effect of the subsequent expression e_3 . As well as checking that types match up properly, we must be sure that the effect produced by the body of the function is a subset of the declared effect. Technically, the rule for functions specifies that the effect of the body exactly equals the declared effect. However, if the body produces a lesser effect, it is always possible to use the last rule (subsumption) to increase the effect of the body so it equals the declared effect.

Finally, we examine the rule for the **letregion** construct. Here, we use the notation $\psi \setminus \{\rho\}$ to denote the effect ψ with all occurrences of $\{\rho\}$ replaced by \emptyset . This rule discharges the effect $\{r\}$ from the effect produced by the subexpression e . Intuitively, because the **letregion** construct discharges effects whereas all other constructs pass on effect information from their subexpressions to their enclosing expressions, any access to a region outside of the scope of a **letregion** will be detected and the type checker will reject the expression as a whole.

Figure 14 shows an example program, a function **count** that counts down to zero. In order to have interesting allocation behavior the integers involved in the count are allocated in a reference cell. The **count** function is stored in region ρ_1 and takes two arguments, a handle for region ρ and an integer reference x allocated in region ρ . If x is nonzero, **count** decrements it, storing the result again in ρ , and recurses. The function has two effects: a read on ρ_1 , resulting from the recursive call, and a read/write effect on ρ , resulting from line 1's read and line 2's store. Therefore, we give the function **count** the effect $\{\rho_1, \rho_2\}$. Overall, the code in Figure 14 allocates two regions (ρ_1 and ρ_2), stores the closure for **count** in ρ_1 , stores an integer reference in ρ_2 , calls **count**, and then deallocates ρ_1 and ρ_2 .

3.2 The Translation

In order to make a formal connection to region-based calculi and to corroborate our claims that we can use the region inference techniques developed by Tofte and others as a front-end for a capability-based compiler, we have defined a type-directed and type-preserving translation from the region calculus of the previous section to CL. Appendix B contains a proof that any well-formed source term is

$$\begin{aligned}
 \mathcal{K}[\text{Type}] &= \text{Type} \\
 \mathcal{K}[\text{Rgn}] &= \text{Rgn} \\
 \mathcal{K}[\text{Eff}] &= \text{Cap} \\
 \mathcal{K}[\alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n] &= \alpha_1:\mathcal{K}[\kappa_1], \dots, \alpha_n:\mathcal{K}[\kappa_n] \\
 \\
 \mathcal{T}[\alpha] &= \alpha \\
 \mathcal{T}[\text{int}] &= \text{int} \\
 \mathcal{T}[\langle \tau_1, \dots, \tau_n \rangle \text{ at } r] &= \langle \mathcal{T}[\tau_1], \dots, \mathcal{T}[\tau_n] \rangle \text{ at } \mathcal{T}[r] \\
 \mathcal{T}[\tau_1 \xrightarrow{\psi} \tau_2 \text{ at } r] &= \forall[\rho':\text{Rgn}, \epsilon:\text{Cap}, \epsilon' \leq \epsilon \oplus \mathcal{T}[\psi] \oplus \{\rho'^1\}].(\epsilon', \\
 &\quad \mathcal{T}[\tau_1], \forall[.](\epsilon', \mathcal{T}[\tau_2]) \rightarrow 0 \text{ at } \rho' \rightarrow 0 \text{ at } \mathcal{T}[r]) \\
 \mathcal{T}[r \text{ handle}] &= \mathcal{T}[r] \text{ handle} \\
 \mathcal{T}[\emptyset] &= \emptyset \\
 \mathcal{T}[\{r\}] &= \{\mathcal{T}[r]^1\} \\
 \mathcal{T}[\psi_1 \cup \psi_2] &= \mathcal{T}[\psi_1] \oplus \mathcal{T}[\psi_2] \\
 \\
 \mathcal{S}[\tau] &= \mathcal{T}[\tau] \\
 \mathcal{S}[\forall[\Delta].\tau_1 \xrightarrow{\psi} \tau_2 \text{ at } r] &= \forall[\mathcal{K}[\Delta].\mathcal{T}[\tau_1 \xrightarrow{\psi} \tau_2 \text{ at } r] \\
 \\
 \mathcal{S}[\{x_1:\sigma_1, \dots, x_n:\sigma_n\}] &= \{x_1:\mathcal{S}[\sigma_1], \dots, x_n:\mathcal{S}[\sigma_n]\}
 \end{aligned}$$

Fig. 15. Region-to-capability kind and type translation.

translated into a well-formed target term.

Kind and Type Translation. The translation is a continuation-passing style transformation in which we simultaneously transform effects into capabilities. The kind and type transformation is presented in Figure 15. The kind translation is trivial; effects become capabilities, and the other kinds are unchanged. The translation of most types is equally simple. The translation of base types is the identity, and, in general, to translate other types we recursively translate their components and recombine using the corresponding capability constructor. Thus, tuples are mapped to tuples and handles are mapped to handles, etc.

The translation of function types is more involved. Recall that in the usual CPS-translation, an arrow type $(\tau_1) \rightarrow \tau_2$ is transformed so that it accepts a translated τ_1 and a τ_2 continuation:

$$(\mathcal{T}[\tau_1], \mathcal{T}[\tau_2] \rightarrow 0) \rightarrow 0$$

The translation of region arrow types has the same structure, but there are several complications that arise as we transform effects into capabilities. The first complication involves finding a region for the continuation closure. We solve this problem by allocating a new region ρ' to hold the continuation; the translated function abstracts this region. The second complication is that an effect for a function may only mention a subset of the regions that are live at the call site. Nevertheless, the resulting CL function must thread the capability describing all the live regions from the context where the function is called through the body of the function to the continuation. We accomplish this task by abstracting an additional capability variable ϵ that makes each function context-sensitive. Using this mechanism, we can thread any capability in the calling context through the function to its continuation.

The third complication is that the type translation must ensure that equal types in the region calculus are translated to equal types in CL. For the most part, this obligation is satisfied trivially because the equality relation for most region types and their corresponding CL analogues is simple syntactic equality up to α -conversion of bound variables. However, the equality relation for effects is set equality whereas the equality relation for *arbitrary* capabilities is not set equality (\oplus is not necessarily idempotent). Fortunately, equality of duplicatable capabilities is exactly set equality. Therefore, the type translation carefully translates all arrow effects into duplicatable capabilities.

These three insights naturally lead us to translate a region function type of the form $\tau_1 \xrightarrow{\psi} \tau_2$ at r into the CL function type

$$\forall[\rho':\text{Rgn}, \epsilon:\text{Cap}, \epsilon' \leq \overline{\epsilon \oplus \mathcal{T}[\psi] \oplus \{\rho'^1\}}]. (\epsilon', \mathcal{T}[\tau_1], \tau_{cont} \text{ at } \rho') \rightarrow \mathbf{0} \text{ at } \mathcal{T}[r]$$

Notice that the capability for the translated function is ϵ' where ϵ' is a subtype of $\overline{\epsilon \oplus \mathcal{T}[\psi] \oplus \{\rho'^1\}}$. This capability gives the translated function access to all the regions it requires: the regions in $\mathcal{T}[\psi]$ are the regions accessed by the source language function; ρ' is the region containing the continuation; and ϵ contains the regions from the calling context that are threaded through the call to the continuation. Notice also that the capability that appears in this type is duplicatable, so we can prove that equal types are translated to equal types. Finally, as explained in Section 2.3, the continuation type τ_{cont} at ρ' should be

$$(\epsilon', \mathcal{T}[\tau_2]) \rightarrow \mathbf{0} \text{ at } \rho'.$$

Bounded quantification allows the continuation to recover the uniqueness information necessary to deallocate the regions used in the function.

Given these definitions, it is straightforward to prove that the essential properties of types (well-formedness, equality, and substitution) are preserved through the translation. Each of the following lemmas can be proven by a simple induction on the well-formedness or equality derivation.

LEMMA (WELL-FORMEDNESS PRESERVATION).

- (1) If $\Delta \vdash_R \Delta'$ then $\mathcal{K}[\Delta] \vdash \mathcal{K}[\Delta']$
- (2) If $\Delta \vdash_R c : \kappa$ then $\mathcal{K}[\Delta] \vdash \mathcal{T}[c] : \mathcal{K}[\kappa]$

LEMMA (EQUALITY PRESERVATION).

- (1) If $\Delta \vdash_R \psi = \psi' : \text{Eff}$ then $\mathcal{K}[\Delta] \vdash \overline{\mathcal{T}[\psi]} = \overline{\mathcal{T}[\psi']} : \text{Cap}$
- (2) If $\Delta \vdash_R c = c' : \kappa$ and κ is not **Eff** then $\mathcal{K}[\Delta] \vdash \mathcal{T}[c] = \mathcal{T}[c'] : \text{Cap}$

LEMMA (SUBSTITUTION PRESERVATION). If $\Delta, \alpha:\kappa \vdash_R \tau : \text{Type}$ and $\Delta \vdash_R c : \kappa$ then $\mathcal{K}[\Delta] \vdash \mathcal{T}[\tau[c/\alpha]] = \mathcal{T}[\tau][\mathcal{T}[c]/\alpha]$.

Term Translation. The heart of the term translation is a continuation-passing style [Fischer 1972; Plotkin 1975] transformation. There are many variations of this transformation [Danvy and Filinski 1992; Sabry and Felleisen 1993; Harper and Lillibridge 1993; Danvy et al. 1999], some of which produce more efficient code than others, and some of which lead to simpler correctness proofs. We have chosen a simple translation that is straightforward to prove type-preserving so that we may focus on the details relevant to region-based memory management.

We begin with an informal description of the basic mechanics of the CPS term translation, ignoring all of the details relevant to regions or capabilities. There are three main arguments to the translation:

- a type-checking context Φ ,
- a source-language term e , and
- a target-language continuation k .

If the source term e is well-formed under the context Φ with type τ , and k is a $\mathcal{T}[\tau]$ -continuation, then the translation $\mathcal{C}_\Phi(e)k$ should produce a well-formed target term.

Operationally, the target term computes e , producing a value v as a result, and then calls the continuation k with v as its argument. Therefore, if the source term e is already a value v , such as an integer or a variable, then the translation is simply the function call $k(v)$. On the other hand, assuming a left-to-right evaluation order, if the source term e actually represents a computation, say the computation $\langle e_1, e_2 \rangle$, the CPS translation arranges to compute e_1 producing value v_1 , compute e_2 producing value v_2 , allocate the pair $\langle v_1, v_2 \rangle$, and finally pass the resulting pointer to the continuation k . We might write such a translation as follows.

$$\begin{aligned} \mathcal{C}_\Phi(\langle e_1^{\tau_1}, e_2^{\tau_2} \rangle)k &= \mathcal{C}_\Phi(e_1)(\lambda x_1. \\ &\quad \mathcal{C}_{\Phi, x_1: \mathcal{T}[\tau_1]}(e_2)(\lambda x_2. \\ &\quad \text{let } x = \langle x_1, x_2 \rangle \text{ in } k(x))) \end{aligned}$$

The translation of each subcomponent of e requires a continuation, and that continuation contains code for all subsequent subcomponents. Finally, the primitive operation op is applied to the resulting values, and the result is passed to k . The compilation of arithmetic operations and the projections have this form.

There are a couple of further details to notice about the translation. First, we have taken the liberty of annotating expressions with their types where necessary (e.g., $e_1^{\tau_1}$). Second, when the translation introduces new variables, such as x_1 , we add those variables, with their translated types, to the context Φ . The latter decision has no influence on the behavior of the translation, but it facilitates the statement and proof of the type correctness theorem.

The translation of function application $e_1 e_2$ begins in the same way as other operations: translate e_1 , passing the resulting value to a continuation that contains the translation of e_2 . The continuation for e_2 contains the function application itself. Because user-defined CPS functions (unlike the primitive operations) do not return, the continuation k is passed directly to the translated function.

$$\begin{aligned} \mathcal{C}_\Phi(e_1^{\tau_1 \rightarrow \tau_2} e_2)k &= \mathcal{C}_\Phi(e_1)(\lambda x_1. \\ &\quad \mathcal{C}_{\Phi, x_1: \mathcal{T}[\tau_1 \rightarrow \tau_2]}(e_2)(\lambda x_2. \\ &\quad x_1(x_2, k))) \end{aligned}$$

Finally, expressions that declare functions must be translated so the result expects an extra continuation argument (x_{cont}) and calls that continuation to return.

$$\begin{aligned} \mathcal{C}_\Phi(\text{let } f : \tau_1 \rightarrow \tau_2 = \lambda x. e \text{ in } e')k &= \\ \text{let } f : \tau_f = \lambda(x, x_{cont}). \mathcal{C}_{\Phi, x: \mathcal{T}[\tau_1], x_{cont}: \tau_{cont}}(e)x_{cont} \text{ in} & \\ \mathcal{C}_{\Phi, f: \mathcal{T}[\tau_1 \rightarrow \tau_2]}(e')k & \end{aligned}$$

Here, the type of the function’s continuation, τ_{cont} , is $(\mathcal{T}[\tau_2]) \rightarrow 0$. The type τ_f of the function itself is $(\mathcal{T}[\tau_1], \tau_{cont}) \rightarrow 0$.

This simple CPS translation provides the basic structure for the translation from the region language into CL. However, as many previous researchers have observed, this translation introduces unnecessary or *administrative* redexes. For example, under the scheme we have presented so far, the translation of a simple pair $\langle 2, 3 \rangle$ with respect to a continuation k is

$$(\lambda x_1. (\lambda x_2. \mathbf{let} \ x_3 = \langle x_1, x_2 \rangle \ \mathbf{in} \ k(x_3)) (3)) (2)$$

instead of the much simpler term $\mathbf{let} \ x_3 = \langle 2, 3 \rangle \ \mathbf{in} \ k(x_3)$. While we are not concerned with the time required to execute the extra function applications, we are concerned about the space required by additional function closures. If we based our region translation directly on this naive translation, we would be forced to allocate additional regions for each of the λ -expressions above. Previous work has avoided these problems by defining the translation in terms of a two-level type system and passing the translation meta level continuations instead of target-level continuations. However, because we are only interested in the space properties of our translation, we will use a simpler solution. Rather than allocating continuation closures, we will use a **let** expression to bind the result of a computation and pass it to a continuation. This solution avoids additional allocation and does not lead to the complexities of a two-level type system. Hence, the translation of the pair $\langle 2, 3 \rangle$ with respect to continuation k will be

$$\begin{aligned} & \mathbf{let} \ x_1 = 2 \ \mathbf{in} \\ & \mathbf{let} \ x_2 = 3 \ \mathbf{in} \\ & \mathbf{let} \ x_3 = \langle x_1, x_2 \rangle \ \mathbf{in} \\ & \mathcal{A}(k, x_3) \end{aligned}$$

The notation $\mathcal{A}(k, x_3)$ denotes static (i.e., translation-time rather than run-time) application of the continuation k to the value x_3 . The continuation k is not represented as a target-language λ -expression, but, intuitively, this “application” is simply $k(x_3)$. The continuation k is actually a pair $\langle x^k; e^k \rangle$. The variable x^k is the continuation’s parameter, and e^k is its body. Given this representation, it is natural to define $\mathcal{A}(\langle x^k; e^k \rangle, v)$ to be $\mathbf{let} \ x^k = v \ \mathbf{in} \ e^k$. The translation of $\mathcal{A}(k, x)$ occurs at compile-time.

Using this notation, we can define a CPS translation for pairs in general as follows.

$$\begin{aligned} \mathcal{C}_\Phi(\langle e_1^{\tau_1}, e_2 \rangle)k &= \mathcal{C}_\Phi(e_1)\langle x_1; \\ & \mathcal{C}_{\Phi, x_1: \mathcal{T}[\tau_1]}(e_2)\langle x_2; \\ & \mathbf{let} \ x = \langle x_1, x_2 \rangle \ \mathbf{in} \ \mathcal{A}(k, x) \rangle \end{aligned}$$

From Region Expressions to Capability Language. With the basic CPS transformation in hand, we are ready to investigate the formal details of the translation of the region language expressions into CL expressions. As discussed above, the translation, \mathcal{C} , has the following form.

$$\mathcal{C}_\Phi(e)k$$

$$\begin{aligned}
 \mathcal{A}(\langle x; e \rangle, v) &= (\text{let } x = v \text{ in } e) \\
 \mathcal{C}_{\Delta; \Gamma; \Theta}(x)k &= \mathcal{A}(k, x) \\
 \mathcal{C}_{\Delta; \Gamma; \Theta}(i)k &= \mathcal{A}(k, i) \\
 \mathcal{C}_{\Delta; \Gamma; \Theta}(f[c_1, \dots, c_n])k &= \mathcal{A}(k, f[\mathcal{T}[[c_1]], \dots, \mathcal{T}[[c_n]]]) \\
 \mathcal{C}_{\Delta; \Gamma; \Theta}(e_1 \text{ p } e_2)k &= \\
 &\quad \mathcal{C}_{\Delta; \Gamma; \Theta}(e_1)\langle x_1; \\
 &\quad \mathcal{C}_{\Delta; \Gamma; \Theta, x_1: \text{int}}(e_2)\langle x_2; \\
 &\quad \text{let } x = x_1 \text{ p } x_2 \text{ in} \\
 &\quad \mathcal{A}(k, x) \rangle \rangle \\
 \mathcal{C}_{\Delta; \Gamma; \Theta}(\langle e_1^{\tau_1}, \dots, e_n^{\tau_n} \text{ at } e_{n+1}^{\tau_{n+1}} \rangle)k &= \\
 &\quad \mathcal{C}_{\Delta; \Gamma; \Theta}(e_1)\langle x_1; \\
 &\quad \vdots \\
 &\quad \mathcal{C}_{\Delta; \Gamma; \Theta, x_1: \mathcal{T}[[\tau_1]], \dots, x_{n-1}: \mathcal{T}[[\tau_{n-1}]]}(e_n)\langle x_n; \\
 &\quad \mathcal{C}_{\Delta; \Gamma; \Theta, x_1: \mathcal{T}[[\tau_1]], \dots, x_n: \mathcal{T}[[\tau_n]]}(e_{n+1})\langle x_{n+1}; \\
 &\quad \text{let } x = \langle x_1, \dots, x_n \rangle \text{ at } x_{n+1} \text{ in } \mathcal{A}(k, x) \\
 &\quad \rangle \rangle \dots \rangle \\
 \mathcal{C}_{\Delta; \Gamma; \Theta}(\pi_i e)k &= \mathcal{C}_{\Delta; \Gamma; \Theta}(e)\langle x_1; \text{let } x = \pi_i x_1 \text{ in } \mathcal{A}(k, x) \rangle \\
 \mathcal{C}_{\Delta; \Gamma; \Theta}(\text{if } 0 \text{ } e_1 \text{ then } e_2 \text{ else } e_3)k &= \\
 &\quad \mathcal{C}_{\Delta; \Gamma; \Theta}(e_1)\langle x_1; \\
 &\quad \text{if } 0 \text{ } x_1 \text{ then } \mathcal{C}_{\Delta; \Gamma; \Theta, x_1: \text{int}}(e_2)k \text{ else } \mathcal{C}_{\Delta; \Gamma; \Theta, x_1: \text{int}}(e_3)k \rangle \\
 \mathcal{C}_{\Delta; \Gamma; \Theta}(\text{letregion } \rho, x_\rho \text{ in } e)k &= \\
 &\quad \text{let newrgn } \rho, x_\rho \text{ in} \\
 &\quad \mathcal{C}_{\Delta\{\rho: \text{Rgn}\}; \Gamma\{x_\rho: \rho \text{ handle}\}; \Theta'}(e)\langle x'; \text{let freergn } x_\rho \text{ in } \mathcal{A}(k, x') \rangle \\
 \text{where} \\
 \Theta' &= \langle \Delta^\Theta; \Gamma^\Theta; C^\Theta \oplus \{\rho^1\}; B^\Theta \oplus \{\rho^1\} \rangle
 \end{aligned}$$

Fig. 16. Region-to-capability term translation.

The context Φ is actually $\Delta; \Gamma; \Theta$. The first two components are a region type context and a region value context. The third component, Θ , is a *translation environment*. This environment contains a CL type context Δ^Θ , a CL value context Γ^Θ , and a pair of capabilities C^Θ and B^Θ . The context Δ^Θ describes the kinds of the new type variables introduced by the translation and, if they are capability type variables, then possibly their bounds. The value context Γ^Θ describes the types of the new value variables introduced by the translation. Intuitively, the capability C^Θ represents the current capability at a given point in the translation; it contains the uniqueness information necessary for deallocating capabilities and will often be a single abstract capability variable ϵ . The capability B^Θ is always a bound on C^Θ . It contains all regions accessed by the current function, and it will never be a single abstract capability variable (see the translation of function declarations below for the differences between B^Θ and C^Θ). If Θ is $\langle \Delta^\Theta; \Gamma^\Theta; C^\Theta; B^\Theta \rangle$, then we use the notation $\Theta, x: \tau$ to denote the translation environment $\langle \Delta^\Theta; \Gamma^\Theta, x: \tau; C^\Theta; B^\Theta \rangle$.

$$\begin{aligned}
& \mathcal{C}_{\Delta;\Gamma;\Theta}(\mathbf{letrec} f[\Delta'](x) : (\forall[\Delta'].\tau_1 \xrightarrow{\psi_f} \tau_2 \text{ at } r) \text{ at } e_1 = e_2 \text{ in } e_3)k = \\
& \quad \mathcal{C}_{\Delta;\Gamma;\Theta}(e_1)\langle x_1; \\
& \quad \quad \mathbf{let } f = (\mathbf{fix } f[\mathcal{K}[\Delta'], \rho; \mathbf{Rgn}, \epsilon; \mathbf{Cap}, \epsilon' \leq \overline{B^{\Theta'}}](\epsilon', x: \mathcal{T}[\tau_1], x_{cont}: \tau_{cont}). \\
& \quad \quad \quad \mathcal{C}_{\Delta\Delta'; \Gamma\{f: \sigma, x: \tau_1\}; \Theta'}(e_2)\langle x_2; x_{cont}(x_2) \rangle) \text{ at } x_1 \\
& \quad \mathbf{in} \\
& \quad \mathcal{C} \quad (e_3)k \\
& \quad \Delta; \Gamma\{f: \forall[\Delta'].\tau_1 \xrightarrow{\psi_f} \tau_2 \text{ at } r\}; \Theta, x_1: \mathcal{T}[\text{r handle}]
\end{aligned}$$

where

$$\begin{aligned}
\tau_{cont} &= \forall[\cdot].(\epsilon', \mathcal{T}[\tau_2]) \rightarrow 0 \text{ at } \rho \\
\Delta^{\Theta'} &= \Delta^{\Theta}, \rho; \mathbf{Rgn}, \epsilon; \mathbf{Cap}, \epsilon' \leq \overline{B^{\Theta'}} \\
\Gamma^{\Theta'} &= \Gamma^{\Theta}, x_1: r \text{ handle}, x_{cont}: \tau_{cont} \\
C^{\Theta'} &= \epsilon' \\
B^{\Theta'} &= \epsilon \oplus \mathcal{T}[\psi_f] \oplus \{\rho^1\} \\
\Theta' &= \langle \Delta^{\Theta'}; \Gamma^{\Theta'}; C^{\Theta'}; B^{\Theta'} \rangle
\end{aligned}$$

$$\begin{aligned}
& \mathcal{C}_{\Delta;\Gamma;\Theta}(e_1^{\tau_f} e_2)k = \\
& \quad \mathcal{C}_{\Delta;\Gamma;\Theta}(e_1)\langle x_1; \\
& \quad \mathcal{C}_{\Delta;\Gamma;\Theta, x_1: \mathcal{T}[\tau_f]}(e_2)\langle x_2; \\
& \quad \quad \mathbf{let newrgn } \rho, x_\rho \text{ in} \\
& \quad \quad \mathbf{let } f_{cont} = (\mathbf{fix } f_{cont}[\cdot](C^{\Theta} \oplus \{\rho^1\}, x: \mathcal{T}[\tau_2])). \mathbf{let freergn } x_\rho \text{ in } \mathcal{A}(k, x) \text{ at } x_\rho \\
& \quad \quad \mathbf{in } x_1[\rho, B^{\Theta}, C^{\Theta} \oplus \{\rho^1\}](x_2, f_{cont}) \rangle
\end{aligned}$$

where

$$\tau_f = \tau_1 \xrightarrow{\psi_f} \tau_2 \text{ at } r$$

Fig. 17. Region-to-capability term translation (functions).

The formal translation is presented in Figures 16 and 17. In the translation, we make the assumption that all variables are fresh and that when we introduce a variable in a term or in a continuation, it is “fresh” (i.e. it does not conflict with any of the other variables in the source term, type-checking context, or continuation).

The invariant guiding the transformation has three main parts:

- (1) The region language term e is well-formed under the type and value contexts Δ and Γ . Formally, $\Delta; \Gamma \vdash_R e : \tau, \psi$.
- (2) The continuation $k = \langle x^k; e^k \rangle$ is well-formed in the current context. Formally, $\{\cdot\}; \mathcal{K}[\Delta], \Delta^{\Theta}; \mathcal{S}[\Gamma], \Gamma^{\Theta}, x^k: \mathcal{T}[\tau]; C^{\Theta} \vdash e^k$.
- (3) Finally, the current capability C^{Θ} is a subcapability of $\overline{B^{\Theta}}$, and moreover, B^{Θ} contains a superset of the regions mentioned in the effect of e . Formally,
 - $\mathcal{K}[\Delta], \Delta^{\Theta} \vdash C^{\Theta} \leq \overline{B^{\Theta}}$ and
 - $\mathcal{K}[\Delta], \Delta^{\Theta} \vdash \overline{B^{\Theta}} = \overline{B^{\Theta} \oplus \mathcal{T}[\psi]}$.

As in the case for the simplified CPS translation, the translation of source-language values is the simplest. For example, to translate a variable x or the integer i , we simply apply the continuation k to x or i respectively. The type application expression $f[c_1, \dots, c_n]$ is also a value because we take the interpretation that types are erased at run time. Therefore, we apply the continuation k directly to $f[\mathcal{T}[c_1], \dots, \mathcal{T}[c_n]]$.

The translation of tuples also follows our informal description closely. We translate each of the computations e_1, \dots, e_{n+1} that make up the tuple in sequence and bind the resulting values to x_1, \dots, x_{n+1} . Once we have translated all expressions, we allocate the tuple $\langle x_1, \dots, x_n \rangle$ using the region handle x_{n+1} .

In order for the tuple allocation operation to be safe, we must ensure the region (say, ρ) that corresponds to the region handle x_{n+1} is still live. In other words, we must be able to prove that the current capability C^\ominus contains a capability for ρ . The invariants above provide us with the means to deduce this fact using the following informal reasoning. Invariant 1 states that the expression $\langle e_1, \dots, e_n \rangle$ at e_{n+1} is well-formed and has effect ψ . Inspection of the region language typing rule for tuples reveals that ψ contains an effect on ρ . Now, the second part of invariant 3 states that the capability B^\ominus contains capabilities on all regions that appear in the effect ψ , including, of course, ρ . Finally, using the first part of invariant 3, we know that the current capability C^\ominus is a subcapability of B^\ominus and, therefore, that it too contains ρ . Consequently, the tuple allocation operation is safe.

Using similar reasoning, it is straightforward to verify informally that the translation of arithmetic operations, projections, and if statements will not fail to type check because they lack sufficient capabilities. In the translation of an if statement, we have duplicated the continuation k inside both branches. Our main concern here is a proof of type preservation, so we have opted for the simplicity of this translation. In practice, such duplication can lead to considerable code growth, and therefore, in many cases, an implementation will want to bind the continuation to a variable before entering the if statement.

The translation of the term `letregion ρ, x_ρ in e` is not much more difficult: `letregion ρ, x_ρ in e` is translated into a `newrgn ρ, x_ρ` declaration followed by the translation of the inner expression e and finally a `freergn` declaration to deallocate ρ . Once again, some simple reasoning allows us to check that the stated invariants hold throughout the transformation. In particular, the translation of the inner expression e reflects the fact that a new region ρ has just been allocated; the translation environment for that step contains capabilities $C^\ominus \oplus \{\rho^1\}$ and $B^\ominus \oplus \{\rho^1\}$. Since we know that C^\ominus is a subcapability of $\overline{B^\ominus}$, we may conclude that $C^\ominus \oplus \{\rho^1\}$ is a subcapability of $\overline{B^\ominus \oplus \{\rho^1\}}$ and therefore that invariant 3, part 1 is satisfied. Next, inspection of the typing rule for `letregion` reveals that if ψ' is the effect of e then $\psi' \setminus \{\rho\}$ is the effect of the entire `letregion` expression. Since, B^\ominus contains $\psi' \setminus \{\rho\}$, we know that $B^\ominus \oplus \{\rho^1\}$ contains capabilities for all regions in ψ' , including, of course, ρ . Hence, we also satisfy invariant 3, part 2 during the translation of e . Finally, recall that the region type system ensures that the region ρ can only be used inside of e . Therefore, in the continuation for the translation of e , we safely free the region and proceed with k . Invariant 2 specified that k only expected the capability C^\ominus , and this is exactly the capability of the context after freeing the region ρ .

The most complex part of the translation involves functions. Fortunately, the type translation, which we have already explained, specifies the main invariants; the translation of functions terms follows directly from this specification. More specifically, we extend the function's type context Δ' with ρ , a region for the continuation's closure, ϵ , a capability for hiding extra regions in the calling context,

and ϵ' , the current capability, which is bounded by $\overline{\epsilon \oplus \mathcal{T}[\psi_f] \oplus \{\rho^1\}}$. We also add a value argument to the translated function, the continuation x_{cont} . It is simple to verify that the translated function has the translated function type. The body of the function is translated under these assumptions. The continuation for this part of the translation does nothing but invoke the function's new argument x_{cont} .

Finally, we examine the translation of a function application. First, the translation allocates a new region ρ for the continuation closure. Next, the translation actually allocates the continuation in ρ . This continuation is defined to expect the capability $C^\Theta \oplus \{\rho^1\}$. This is the maximum capability at this point in the computation, and it permits the continuation to deallocate the region containing its own closure. Just as a conventional stack-based language implementation, which allocates a stack frame on a function call and deallocates that stack frame on function return, the CPS region-based implementation allocates a continuation closure on a function call and then deallocates the continuation closure on return.

After allocating the continuation, the translation continues with the translation of the actual function application. The translated function value, x_1 , will have the type

$$\forall[\rho':\text{Rgn}, \epsilon:\text{Cap}, \epsilon' \leq \overline{\epsilon \oplus \mathcal{T}[\psi_f] \oplus \{\rho^1\}}].(\dots) \rightarrow 0 \text{ at } r.In$$

Therefore, before calling the function, the code must instantiate the variables ρ' , ϵ , and ϵ' properly. The code naturally instantiates ρ' with the region ρ that was just allocated. At this point in the program, the capability $C^\Theta \oplus \{\rho^1\}$ represents the current capability, and $B^\Theta \oplus \{\rho^1\}$ is its upper bound. Therefore, the code instantiates ϵ with B^Θ and ϵ' with $C^\Theta \oplus \{\rho^1\}$, which is legal provided we can prove that

$$C^\Theta \oplus \{\rho^1\} \leq \overline{B^\Theta \oplus \psi_f \oplus \{\rho^1\}}.$$

Given invariant 3, which states that $C^\Theta \leq \overline{B^\Theta}$ and that $\overline{B^\Theta} = \overline{B^\Theta \oplus \psi_f}$, it is easy to verify this fact. Now, we need only verify that the term arguments, x_2 and the continuation f_{cont} , have the types expected by the translated function, and this too can be easily checked.

Properties of the Translation. We have proven that our translation preserves types.

THEOREM (CPS TYPE PRESERVATION). *If $\cdot; \cdot \vdash e : \text{int}, \emptyset$ then $\{\}; \cdot; \cdot; \emptyset \vdash \mathcal{C}_{\cdot; \Theta}(e) \langle x; \text{halt } x \rangle$ where x is fresh and Θ is the empty environment: $\langle \cdot; \cdot; \emptyset; \emptyset \rangle$.*

The proof proceeds by induction on the typing derivation of the source language term with invariants 1, 2, and 3 forming the induction hypothesis. Appendix B contains a formal proof of the representative cases.

We would also like to prove our translation is correct and that it preserves the space used by the program, but we have not yet tackled these problems. Recently, Minamide [1999] has proven that a standard CPS translation preserves the maximum amount of reachable space within a constant factor. He defines a space-profiling semantics for the simply typed lambda calculus that refines the work of Blleloch and Greiner [1996]. If we were to augment our semantics with this sort of

space-profiling information, we may be able to prove a similar result for our translation. An informal inspection of the translation indicates that the resultant term allocates no more data structures than the source term with the exception of the continuation closure that we require to return from a function call, and intuitively, this closure corresponds to the stack space that is required to save local variables across a function call. However, a formal investigation of the space properties of our translation is beyond the scope of this paper.

An Example. The program in Figure 18 is the translation of the `count` function from the previous section. We have made several simplifications to the output of the formal translation to improve the readability of the program. In particular, we have optimized away many of the administrative redexes and performed a tail-call optimization on the recursive call to the `count` function. Rather than writing capabilities $\{\overline{\rho_1^1}, \rho^1\}$, we use the equivalent form $\{\rho_1^+, \rho^+\}$.

The program begins by allocating regions ρ_1 and ρ_2 using the `newrgn` declaration, and puts the closure for `count` into ρ_1 . The `count` function requires a capability ϵ' at least as good as the capability $\{\rho_1^+, \rho^+, \rho_{cont}^+\}$ needed to access itself, its argument, and its continuation; and it passes on that same capability ϵ' to its continuation k . As we type check the body of the `count` function, we verify that we possess the capabilities necessary to make all data accesses legal. Comments in the code indicate where these checks occur. When calling `count`, we pass it the continuation `cont`. This continuation requires the capability $\{\rho_1^1, \rho_2^1, \rho_3^1\}$ in order to free the three regions. Hence, at the application site, `count`'s capability, ϵ' , is instantiated with the stronger capability needed by the continuation.

Another Example. In this context, the `count` function uses all of the regions that are currently allocated, and the capability variable ϵ is redundant. When the code instantiates ϵ at the call site for `count`, it does so with exactly the regions ρ_1 , ρ , and ρ_{cont} which already appear in the bound on ϵ' . However, in general, ϵ will hide some “left-over” capability. For example, if we had allocated a fourth region, ρ_4 , we would need to instantiate ϵ with the capability $\{\rho_4^1\} \oplus \{\rho_1^1, \rho_2^1, \rho_3^1\}$ and make corresponding changes to the continuation. Now, ϵ would hide the capability on the fourth region from `count` but preserve it across the call so it could be deallocated in the continuation:

```
%% count with  $\epsilon$  hiding a left-over capability
let newrgn  $\rho_1, x_{\rho_1}$  in
let newrgn  $\rho_2, x_{\rho_2}$  in
let newrgn  $\rho_3, x_{\rho_3}$  in
let newrgn  $\rho_4, x_{\rho_4}$  in
let count = ... as before ...
% capability held is  $\{\rho_1^1, \rho_2^1, \rho_3^1, \rho_4^1\}$ 
let ten =  $\langle 10 \rangle$  at  $x_{\rho_2}$  in
let cont =
  ( $\lambda (\{\rho_1^1, \rho_2^1, \rho_3^1, \rho_4^1\}) \dots$ ) at  $x_{\rho_2}$ 
in
count [ $\rho_2, \rho_3, \{\rho_4^1\} \oplus \{\rho_1^1, \rho_2^1, \rho_3^1\}, \{\rho_4^1\} \oplus \{\rho_1^1, \rho_2^1, \rho_3^1\}$ ] ( $x_{\rho_2}, \text{ten}, \text{cont}$ )
```

```

let newrgn  $\rho_1, x_{\rho_1}$  in
let newrgn  $\rho_2, x_{\rho_2}$  in
% capability held is  $\{\rho_1^1, \rho_2^1\}$ 
let count =
  (fix count
    [ $\rho$ :Rgn,  $\rho_{cont}$ :Rgn,  $\epsilon$ :Cap,  $\epsilon' \leq \bar{\epsilon} \oplus \{\rho_1^+, \rho^+, \rho_{cont}^+\}$ ]
    ( $\epsilon', x_\rho$ : $\rho$  handle,  $x$ : $\langle \text{int} \rangle$  at  $\rho, k$ :( $\epsilon'$ )  $\rightarrow 0$  at  $\rho_{cont}$ ) .
    % capability held is  $\epsilon' \leq \bar{\epsilon} \oplus \{\rho_1^+, \rho^+, \rho_{cont}^+\}$ 
    let  $n = \pi_1(x)$  in                                     %  $\rho$  ok
    if0  $n$ 
      then  $k()$                                            %  $\rho_{cont}$  ok
    else
      let  $n' = n - 1$  in
      let  $x' = \langle n' \rangle$  at  $x_\rho$  in                       %  $\rho$  ok
      count [ $\rho, \rho_{cont}, \bar{\epsilon} \oplus \{\rho_1^+, \rho^+, \rho_{cont}^+\}, \epsilon'$ ] ( $x_\rho, x', k$ ) %  $\rho_1$  ok
    ) at  $x_{\rho_1}$  in
let newrgn  $\rho_3, x_{\rho_3}$  in
% capability held is  $\{\rho_1^1, \rho_2^1, \rho_3^1\}$ 
let ten =  $\langle 10 \rangle$  at  $x_{\rho_2}$  in
let cont =
  ( $\lambda (\{\rho_1^1, \rho_2^1, \rho_3^1\})$  .
    % capability held is  $\{\rho_1^1, \rho_2^1, \rho_3^1\}$ 
    let freergn  $x_{\rho_3}$  in                                 %  $\rho_3$  unique
    let freergn  $x_{\rho_2}$  in                                 %  $\rho_2$  unique
    let freergn  $x_{\rho_1}$  in                                 %  $\rho_1$  unique
    halt 0
  ) at  $x_{\rho_3}$ 
in
count [ $\rho_2, \rho_3, \{\rho_1^1, \rho_2^1, \rho_3^1\}, \{\rho_1^1, \rho_2^1, \rho_3^1\}$ ] ( $x_{\rho_2}, \text{ten}, \text{cont}$ )

```

Fig. 18. The function count.

The power of bounded quantification comes into play when a function is called with several regions, some of which may or may not be the same. For example, the original code could be rewritten to have `ten` and `cont` share a region, without changing the function count in any way:

```

%%% count with ten and cont sharing  $\rho_2$ 
let newrgn  $\rho_1, x_{\rho_1}$  in
let newrgn  $\rho_2, x_{\rho_2}$  in
let count = ... as before ...
% capability held is  $\{\rho_1^1, \rho_2^1\}$ 
let ten =  $\langle 10 \rangle$  at  $x_{\rho_2}$  in
let cont =
  ( $\lambda (\{\rho_1^1, \rho_2^1\}) \dots$ ) at  $x_{\rho_2}$ 
in
count [ $\rho_2, \rho_2, \{\rho_1^1, \rho_2^1\}, \{\rho_1^1, \rho_2^1\}$ ] ( $x_{\rho_2}, \text{ten}, \text{cont}$ )

```

In this example, ρ_{cont} is instantiated with ρ_2 and ϵ' is instantiated with $\{\rho_1^1, \rho_2^1\}$ (which is again the capability required by `cont`). However, `count` proceeds exactly

```

let newrgn  $\rho_1, x_{\rho_1}$  in
let newrgn  $\rho_2, x_{\rho_2}$  in
% capability held is  $\{\rho_1^1, \rho_2^1, \rho_3^1\}$ 
let count =
  (fix count
    [ $\rho$ :Rgn,  $\rho_{cont}$ :Rgn,  $\epsilon \leq \{\rho_1^+, \rho_{cont}^+\}$ ]
    ( $\epsilon \oplus \{\rho^1\}, x_{\rho}:\rho$  handle,  $x:(int)$  at  $\rho$ ,
      $k:(\epsilon) \rightarrow 0$  at  $\rho_{cont}$ ) .
    % capability held is  $\epsilon \oplus \{\rho^1\}$ 
    let  $n = \pi_1(x)$  in %  $\rho$  ok
    let freergn  $x_{\rho}$  in %  $\rho$  unique
    % capability held is  $\epsilon$ 
    if0  $n$ 
    then  $k()$  %  $\rho_{cont}$  ok
    else
      let  $n' = n - 1$  in
      let newrgn  $\rho', x_{\rho'}$  in
      % capability held is  $\epsilon \oplus \{\rho'^1\}$ 
      let  $x' = \langle n' \rangle$  at  $x_{\rho'}$  in %  $\rho'$  ok
      count [ $\rho', \rho_{cont}, \epsilon$ ] ( $x_{\rho'}, x', k$ ) %  $\rho_1$  ok
  ) at  $x_{\rho_1}$  in
let ten =  $\langle 10 \rangle$  at  $x_{\rho_2}$  in
let newrgn  $\rho_3, x_{\rho_3}$  in
let cont =
  ( $\lambda (\{\rho_1^1, \rho_3^1\})$  .
   % capability held is  $\{\rho_1^1, \rho_3^1\}$ 
   let freergn  $x_{\rho_3}$  in %  $\rho_3$  unique
   let freergn  $x_{\rho_1}$  in %  $\rho_1$  unique
   halt 0
  ) at  $x_{\rho_3}$ 
in
count [ $\rho_2, \rho_3, \{\rho_1^1, \rho_3^1\}$ ] ( $x_{\rho_2}, ten, cont$ )

```

 Fig. 19. The function `count` with efficient memory usage.

as before because ϵ' is still as good as $\{\rho_1^+, \rho^+, \rho_{cont}^+\}$ since:

$$\begin{aligned}
 \{\rho_1^1, \rho_2^1\} &\leq \{\rho_1^+, \rho_2^+\} \\
 &= \{\rho_1^+, \rho_2^+\} \oplus \{\rho_1^+, \rho_2^+, \rho_2^+\} \\
 &= \{\rho_1^1, \rho_2^1\} \oplus \{\rho_1^+, \rho_2^+, \rho_2^+\} \\
 &= (\bar{\epsilon} \oplus \{\rho_1^+, \rho_2^+, \rho_2^+\})[\{\rho_1^1, \rho_2^1\}/\epsilon]
 \end{aligned}$$

An Optimization. In the examples above, even though `count` is tail-recursive, we allocate a new cell each time around the loop, and we do not deallocate any of the cells until the count is complete. However, since ρ never contains any live values other than the current argument, it is safe to reduce the program's space usage by deallocating the argument's region each time around the loop, as shown in Figure 19. Note that this optimization is not possible when region lifetimes must be lexically scoped.

In order to deallocate its argument, the revised `count` requires a unique capability for its argument's region ρ . Note that if the program were again rewritten so that `ten` and `cont` shared a region (which would lead to a run-time error, since `ten` is deallocated early), the program would no longer typecheck: In order to satisfy

the bound on `count`, ϵ must be instantiated with a capability containing `count`'s region (ρ_2 in this scenario), but then `count`'s required capability will contain two occurrences of ρ_2 , one from ϵ and one from the instantiation of ρ by ρ_2 . Since the latter of these two capabilities is required to be unique, the union of two cannot be satisfied by the current capability. For instance, if ϵ is instantiated with $\{\rho_1^1, \rho_2^1\}$ then the function call does not typecheck since the current capability $\{\rho_1^1, \rho_2^1\} \not\leq \{\rho_1^1, \rho_2^1, \rho_2^1\}$.

On the other hand, the program rewritten so that `count` and `cont` share a region does not fail at run time, and does typecheck. In this case, it is legal to instantiate ϵ with $\{\rho_1^1\}$ and it is possible to prove that the capability held at the point of the function call is a subcapability of the expected capability: $\{\rho_1^1, \rho_2^1\} \leq \{\rho_1^1, \rho_2^1\}$.

4. DISCUSSION

We believe the general framework of our capability system is quite robust. There are several ways to extend the language and a number of directions for future research.

4.1 Language Extensions

In this article, we have concentrated on using CL to implement safe region-based memory management, but with a few changes, we believe our capability apparatus may be used in a variety of other settings as well. One potential application involves reducing the overhead of communication across the user-kernel address space boundary in traditional operating systems. Typically, in such systems, when data in user space are presented to the kernel, the kernel must copy that data to ensure its integrity is preserved. However, if a user process hands off a unique capability for a region to the kernel, the kernel does not have to copy that region's data; without the capability, the user can no longer read or modify the contents of that region.

By handing off a user's capability to the kernel, we ensure that the kernel has exclusive access to the data governed by the capability. We can generalize this idea and use capabilities to ensure mutually exclusive access to shared mutable data in a multithreaded environment, by viewing locks as analogous to regions. If we associate each piece of sensitive data with a lock, we can statically check that every client of the data obtains the corresponding lock and its associated capability before attempting access. When the code releases the lock, the type system would revoke the capability on the data, just as it revokes a capability after a region is freed. Flanagan and Abadi [1999] have investigated this idea in the context of a high-level lexically scoped language. Just as we compiled Tofte and Talpin's high-level region language into CL, we conjecture we could compile Flanagan and Abadi's locking language into a variant of CL with locking primitives instead of allocation primitives.

A third application of capabilities is to control and reason about aliasing on a per-object basis rather than a per-region basis. Smith, Walker, and Morrisett [Smith et al. 2000; Walker and Morrisett 2000] have investigated the idea of associating a different capability with each individual object and including the type of the object within the capability itself. When code possesses the unique capability for an object, it may deallocate the object, or, if it chooses, it may explicitly reuse the space for that object to store a value of a different type. This new design

may be viewed as an extension to conventional linear type systems [Girard 1987; Lafont 1988; Wadler 1990; Abramsky 1993] as it admits aliasing and yet allows safe deallocation of objects. Recently, these techniques have been used to extend the Typed Assembly Language implementation [Morrisett et al. 2000] with operations for explicit, but verifiably safe memory management.

In general, whenever a system wishes to restrict access to some data statically, and/or to ensure a certain sequence of operations are performed, it may consider using capabilities. In fact, Walker [2000] has shown that the combination of capabilities and a simple logic are sufficiently powerful to encode any *safety property* [Alpern and Schneider 1987; Schneider 2000].

4.2 Related Work

There are many formalisms for reasoning about computational effects in programming languages including type-and-effects systems [Gifford and Lucassen 1986; Lucassen 1987; Jouvelot and Gifford 1991; Tofte and Talpin 1994], monads [Moggi 1991; Peyton Jones and Wadler 1993; Launchbury and Peyton Jones 1995; Filinski 1996], linear types [Girard 1987; Lafont 1988; Wadler 1990; Abramsky 1993], and now capabilities. Many researchers are actively investigating the relationships between these different areas, but the overall picture is not yet fully understood. We are eager to continue this line of research and explore the formal links between our system and the others.

Our translation of Tofte and Talpin’s region calculus into CL reveals that the relationship between effects and capabilities is quite close. A necessary prerequisite for the use of either system is type inference, performed by a programmer or compiler, and much of the research into effects systems has concentrated on this difficult task. However, because of the focus on inference, effect systems are usually formulated as a bottom-up synthesis of effects. Our work may be viewed as producing verifiable evidence of the correctness of an inference. Hence, while effect systems typically work bottom-up, specifying the effects that might occur, we take a top-down approach, specifying by capabilities the effects that are *permitted to occur*. Moreover, unlike Tofte and Talpin’s effect system, our capabilities are sensitive to control-flow. Rather than constructing the overall effect of an expression by taking the union of the effects of the subexpressions, and thereby losing information about the order of evaluation, we verify that programs are safe by checking one instruction after another and using the capability produced by previous instructions to verify that the instructions that follow are safe.

A connection can also be drawn between capabilities and monadic type systems. Work relating effects to monads has viewed effectful functions as pure functions that return state transformers. This might be called an *ex post* view: the effect takes place after the function’s execution. In contrast, we take an *ex ante* view in which the capability to perform the relevant effect must be satisfied *before* the function’s execution. Nevertheless, there is considerable similarity between the views; just as monads can be used to ensure that the store is single-threaded through a computation, our typing rules thread a capability (which summarizes aspects of the store) along the execution path of a program.

Perhaps the closest relationship occurs between linear type systems and capabilities. An inspection of the axioms of capability equality reveals that they are

very similar to the structural rules of linear type systems. In particular, linear assumptions, like unique capabilities, do not allow contraction and weakening rules whereas nonlinear assumptions, like duplicatable capabilities, do allow contraction and weakening rules.⁶ One essential difference between the two formalisms is that the capability to access an object (say, $\{\rho^1\}$) is separated from the type of the object itself (say, $\langle \text{int} \rangle$ at ρ). This level of indirection makes it possible to allow aliasing and yet verify that deallocation is still safe.

There has also been a significant amount of prior research on the more specific topic of the theory and implementation of region-based memory management. With respect to implementation, Birkedal et al. [1996] describe several optimizations to the basic region-allocation scheme that are used in the ML Kit with Regions to improve space efficiency. One of their observations is that functions can be used in two different contexts: one context in which no live object remains in a region after a function call and a second context in which there may be live objects remaining in a region after a call. In order to avoid code duplication and yet ensure efficient space usage, the call site passes information to the called function at run time. Using this information, the function may make dynamic decisions about region deallocation. The type system we present here is not powerful enough to encode these *storage-mode polymorphic* functions. However, we believe these dynamic tests may be viewed as a form of intensional type analysis [Harper and Morrisett 1995; Crary et al. 1998], and, therefore, if we augment CL with a variant of Harper and Morrisett’s typecase mechanism, we may be able to verify the results of storage-mode optimizations as well.

Aiken et al. [1995] have also studied how to optimize the original Tofte-Talpin region framework. As in CL, they separate region allocation from region deallocation. However, they have not presented a technique for verifying that the results of their optimizations are safe. We conjecture, based on the soundness proof for Aiken et al.’s analyses, that the analysis could be used to produce typing annotations and that verification could take place using CL.

Gay and Aiken [1998] have developed extensions to C that gives programmers complete control over region allocation and deallocation. They use reference counting to prevent programmers from accidentally accessing deallocated regions. Hawblitzel and von Eicken [Hawblitzel and von Eicken 1999] have also used the notion of a region in their language Passport to support sharing and revocation between multiple protection domains. Both of these groups use run-time checking to ensure safety, and it would be interesting to investigate hybrid systems that combine features of our static type system with more dynamic systems.

Tofte and Talpin [1997] have studied the soundness of region-based type systems at length. They use a greatest fixed-point construction and a coinductive argument to prove the correctness of their region-inference scheme. In contrast, our formulation of CL allows us to use the syntactic proof techniques popularized by Wright and Felleisen [1994]. However, despite the high-level differences between the proof

⁶Many formulations of linear logic admit a weakening rule that allows an assumption to be completely forgotten. As explained earlier, we do not allow complete forgetting of capabilities because it leads to space leaks. Instead, we admit a more restrictive weakening rule that allows all but the last capability to be forgotten.

techniques, there are illuminating similarities in some of the details. Most notably, Tofte and Talpin’s proof involves a notion of *consistency* that relates source and target values in the region inference translation. Consistency is defined with respect to the effect (ψ) of the rest of the computation. Informally, one of the consistency conditions states that a source value is consistent with a target value in region ρ , with respect to effect ψ , if ρ does not appear in ψ . Hence, if ρ is not in the effect, or *capability*, of the rest of the computation, then we can deallocate that region because the rest of the computation cannot distinguish a dangling pointer into ρ from a value in the source language. Therefore, within the Tofte-Talpin proof, the effect of the rest of the computation plays a role very similar to a capability. We are able to give a syntactic proof of soundness for our language because continuations and their capabilities are explicit in our framework whereas Tofte and Talpin introduce this idea as a metalevel construction in their proof.

5. CONCLUSIONS

We have presented a new strongly typed language that admits operations for explicit allocation and deallocation of data structures. Furthermore, this language is expressive enough to serve as a target for region inference and admits a relatively straightforward proof of soundness. We believe that the notion of capabilities that support statically checkable attenuation, amplification, and revocation is an effective new tool for language designers.

APPENDIX

A. SOUNDNESS OF THE CAPABILITY LANGUAGE

A.1 Notation

The capability $\{r^+\}$ is a derived form that we used for expository purposes in the article. It is equivalent to $\{\overline{r^1}\}$. For the sake of simplicity, the proof operates on a new language that does not include capabilities of the form $\{r^+\}$. The syntax of capabilities is

$$C ::= \epsilon \mid \emptyset \mid \{r\} \mid C_1 \oplus C_2 \mid \overline{C}.$$

The form $\{r\}$ is the new syntax for unique capabilities. The only way to form a duplicatable capability is to use the bar operator as in $\{\overline{r}\}$; hence multiplicity annotations are unnecessary. The rule eq-flag is a derived rule. If the abbreviation $\{r^+\}$ is replaced by its definition, it is clear the rule is simply a special case of reflexivity, and therefore, we do not need it in our system. Where convenient, we continue to use $\{r^+\}$ as an abbreviation for $\{\overline{r}\}$ and $\{r^1\}$ to emphasize the fact that a capability is unique. We also continue to use the metavariable φ to range over $+$ or 1 .

We also use the following notational conventions:

- We abbreviate $(\dots((\emptyset \oplus \{r_1^{\varphi_1}\}) \oplus \{r_2^{\varphi_2}\})\dots) \oplus \{r_n^{\varphi_n}\}$ by $\{r_1^{\varphi_1}, \dots, r_n^{\varphi_n}\}$.
- We abbreviate $(\dots((\emptyset \oplus C_1) \oplus C_2)\dots) \oplus C_n$ by $C_1 \oplus \dots \oplus C_n$.
- We use the notation $C \in C'$ to denote the fact that C is a subcomponent of C' .

A.2 Overview

The proof is broken down into a series of lemmas, most of which are proven by induction on the typing derivations or by induction on the syntax of the language. The proof culminates in a proof Type Soundness and Complete Collection. The supporting lemmas are grouped as follows:

- Lemmas 9 to 11 describe when extensions to type contexts or exchanges of elements within a type context are permissible.
- Lemmas 12 to 14 state that constructors involved in equality and subtyping judgements are well-formed and that all free variables of well-formed constructors are bound by the type context.
- Definitions and Lemmas 15 to 22 describe which capabilities are equal to one another and which capabilities are subtypes of one another. They provide a higher level of abstraction than the rules for equality and subtyping and are used frequently in the rest of the proof.
- Lemmas 23 and 24 are substitution lemmas for types and values respectively.
- Lemma 25 states that well-formed small values, heap values, and declarations have well-formed types.
- Lemmas 26 to 28 are Canonical Forms lemmas. Given a type, these lemmas describe the shape of memory or of values.
- Lemmas 29 to 31 describe the conditions under which you can add labels or regions to the memory type and preserve typing.
- Lemma 33 states that satisfiability is preserved across equality and subtyping (under the empty context).
- Lemma 34 states that satisfiability is preserved when a region and the corresponding unique capability are removed both from memory and the current capability simultaneously.
- Lemmas 35 and 36 are the Preservation and Progress lemmas respectively. They are used directly in the proof of Type Soundness.

LEMMA 9. *If $\Delta \vdash \Delta'$ then $Dom(\Delta) \cap Dom(\Delta') = \emptyset$.*

PROOF. By induction on the derivation. \square

LEMMA (TYPE CONTEXT EXCHANGE). *If $Dom(\Delta_1) \cap Dom(\Delta_2) = \emptyset$ then*

- (1) *If $\Delta_0\Delta_1\Delta_2\Delta_3 \vdash \Delta$ then $\Delta_0\Delta_2\Delta_1\Delta_3 \vdash \Delta$*
- (2) *If $\Delta_0\Delta_1\Delta_2\Delta_3 \vdash c : \kappa$ then $\Delta_0\Delta_2\Delta_1\Delta_3 \vdash c : \kappa$.*

PROOF. By induction on the derivations. In the rule type-var

$$\frac{}{\Delta_0\Delta_1\Delta_2\Delta_3 \vdash \alpha : \kappa} (\Delta_0\Delta_1\Delta_2\Delta_3(\alpha) = \kappa)$$

we know $\Delta_0\Delta_1\Delta_2\Delta_3(\alpha) = \Delta_0\Delta_2\Delta_1\Delta_3(\alpha)$ because the domains of Δ_1 and Δ_2 are disjoint. Consequently, $\Delta_0\Delta_2\Delta_1\Delta_3 \vdash \alpha : \kappa$. \square

LEMMA (TYPE CONTEXT EXTENSION). *If $\Delta \vdash \Delta'$ then*

- (1) *If $\Delta \vdash \Delta''$ and $Dom(\Delta'') \cap Dom(\Delta') = \emptyset$ then $\Delta\Delta' \vdash \Delta''$*
- (2) *If $\Delta \vdash c : \kappa$ then $\Delta\Delta' \vdash c : \kappa$*

- (3) If $\Delta \vdash c_1 = c_2 : \kappa$ then $\Delta\Delta' \vdash c_1 = c_2 : \kappa$
 (4) If $\Delta \vdash c_1 \leq c_2 : \kappa$ then $\Delta\Delta' \vdash c_1 \leq c_2 : \kappa$.

PROOF. By induction on the derivation. Almost all cases follow directly from the inductive hypothesis. Rules `ctxt-sub` and `type-arrow` require Type Context Exchange where Δ_3 is \cdot . \square

LEMMA 12. If $\Delta \vdash c : \kappa$ then $\text{ftv}(c) \subseteq \text{Dom}(\Delta)$.

PROOF. By induction on the derivation. \square

LEMMA (EQUALITY REGULARITY). If $\Delta \vdash C = C' : \kappa$ then $\Delta \vdash C : \kappa$ and $\Delta \vdash C' : \kappa$.

PROOF. By induction on the derivation. \square

LEMMA (SUBTYPING REGULARITY). If $\cdot \vdash \Delta$ and $\Delta \vdash C \leq C' : \kappa$ then $\Delta \vdash C : \kappa$ and $\Delta \vdash C' : \kappa$.

PROOF. By induction on the derivation. In the rule `sub-var`, we show by induction on the derivation $\cdot \vdash \Delta$ that if $(\epsilon \leq C) \in \Delta$ then $\Delta \vdash \Delta(\epsilon) : \text{Cap}$. \square

Definition 15. An atomic capability, a , is a type variable ϵ of kind `Cap`, a singleton capability $\{r\}$, or a barred capability $\bar{\epsilon}$ or $\{\bar{r}\}$. The meta-variable a ranges over atomic capabilities.

Definition 16. $E(C)$ is the set of elements ϵ or $\{r\}$ that appear in C (where $\{x_1, \dots, x_n\}$ is notation for the set of elements x_1, \dots, x_n):

$$\begin{aligned} E(\emptyset) &= \emptyset \\ E(\{r\}) &= \{\{r\}\} \\ E(\epsilon) &= \{\{\epsilon\}\} \\ E(C_1 \oplus C_2) &= E(C_1) \cup E(C_2) \\ E(\bar{C}) &= E(C) \end{aligned}$$

LEMMA (EQUALITY). If $\Delta \vdash C : \text{Cap}$ then

- (1) $\Delta \vdash C = a_1 \oplus \dots \oplus a_n : \text{Cap}$ for some atomic capabilities a_1, \dots, a_n .
- (2) $\Delta \vdash a_1 \oplus \dots \oplus a_{i-1} \oplus a_i \oplus a_{i+1} \oplus \dots \oplus a_n = a_1 \oplus \dots \oplus a_{i-1} \oplus a_{i+1} \oplus \dots \oplus a_n \oplus a_i : \text{Cap}$.
- (3) $\Delta \vdash a_1 \oplus \dots \oplus a_n = a'_1 \oplus \dots \oplus a'_n : \text{Cap}$ where a'_1, \dots, a'_n is any permutation of a_1, \dots, a_n .
- (4) $\Delta \vdash a_1 \oplus \dots \oplus a_n = a'_1 \oplus \dots \oplus a'_m : \text{Cap}$ where $\bar{a}'_1, \dots, \bar{a}'_m$ is a subsequence of a_1, \dots, a_n with all duplicate barred elements removed.
- (5) If $\Delta \vdash C = C' : \text{Cap}$ then the sets $E(C)$ and $E(C')$ are equal.
- (6) If $E(C) = E(C')$ and $\Delta \vdash C' : \text{Cap}$ then $\Delta \vdash \bar{C} = \bar{C}' : \text{Cap}$.
- (7) If $\Delta \vdash C \oplus \{r\} = C' \oplus \{r\} : \text{Cap}$ then $\Delta \vdash C = C' : \text{Cap}$.

PROOF. Part 1 follows by induction on the derivation $\Delta \vdash C : \text{Cap}$. Case `type- \emptyset` is immediate. Case `type-single` follows from application of the equality rules `eq-symm` and `eq- \emptyset` . Case `type-plus` is more intricate. The inductive hypothesis gives us

$$\begin{aligned} \Delta \vdash C_1 &= a_1 \oplus \dots \oplus a_n : \text{Cap} \\ \Delta \vdash C_2 &= a'_1 \oplus \dots \oplus a'_m : \text{Cap}. \end{aligned}$$

By induction on m and using the rules eq- \emptyset , eq-assoc, and eq-trans

$$\Delta \vdash a'_1 \oplus \cdots \oplus a'_m = a'_1 \oplus (a'_2 \oplus \cdots \oplus (a'_{m-1} \oplus a'_m) \cdots) : \mathbf{Cap}.$$

By equality congruence and eq-trans,

$$\Delta \vdash C_1 \oplus C_2 = (a_1 \oplus \cdots \oplus a_n) \oplus a'_1 \oplus (a'_2 \oplus \cdots \oplus (a'_{m-1} \oplus a'_m) \cdots) : \mathbf{Cap}.$$

By induction on m again and using eq-assoc, eq-symm, and eq-trans,

$$\Delta \vdash C_1 \oplus C_2 = a_1 \oplus \cdots \oplus a_n \oplus a'_1 \oplus \cdots \oplus a'_m : \mathbf{Cap}.$$

For the case \overline{C} , we have $\Delta \vdash C = a_1 \oplus \cdots \oplus a_n : \mathbf{Cap}$ by IH. By congruence, $\Delta \vdash \overline{C} = \overline{a_1 \oplus \cdots \oplus a_n} : \mathbf{Cap}$. By induction on n , $\Delta \vdash \overline{C} = \overline{a_1} \oplus \cdots \oplus \overline{a_n} : \mathbf{Cap}$. For each a_i , either $\overline{a_i}$ is an atomic element or a_i is already barred, and we use the eq-bar-idem rule to show that $\Delta \vdash \overline{a_i} = a_i : \mathbf{Cap}$. In either case, by induction on n again and use of the congruence rules, we are done.

Part 2 follows by induction on $m - i$ using eq-assoc, eq-comm, as well as the transitivity and symmetry of equality. Part 3 is a corollary of part 2. Part 4 follows by induction on the number of barred duplicates and uses part 3, transitivity, symmetry, and eq-dup rules. Part 5 follows by induction on the equality judgment.

Part 6 may be proven as follows:

$\Delta \vdash C = a_1 \oplus \cdots \oplus a_n : \mathbf{Cap}$ where $E(C) = E(a_1 \oplus \cdots \oplus a_n)$ by parts 1 and 5.

$\Delta \vdash C' = a'_1 \oplus \cdots \oplus a'_m : \mathbf{Cap}$ where $E(C') = E(a'_1 \oplus \cdots \oplus a'_m)$ by parts 1 and 5.

By parts 3 and 4 and congruence of equality: $\Delta \vdash \overline{C} = \overline{a_1 \oplus \cdots \oplus a_n} = a_{j_1} \oplus a_{j_n} : \mathbf{Cap}$

$\Delta \vdash \overline{C'} = \overline{a'_1 \oplus \cdots \oplus a'_m} = a'_{j_1} \oplus a'_{j_m} : \mathbf{Cap}$

where the a_{j_i} and a'_{j_i} contain no duplicates and are ordered according to some canonical ordering. If $E(C) = E(C')$ then the a_{j_i} and the a'_{j_i} are the same and are in the same order. Hence, the constructors are syntactically equal and thus definitionally equal.

Part 7 follows by induction on the typing derivation. \square

Definition 18. forA capability C is *unique* in C' if there does not exist C'' such that $\overline{C''} \in C'$ and $C \in C''$. A capability C is *duplicatable* in C' if $\overline{C''} \in C'$ and $C \in \overline{C''}$.

LEMMA 19. *If $\Delta \vdash C' : \mathbf{Cap}$ and C is duplicatable in C' then $\Delta \vdash C' = C'' \oplus \overline{C} : \mathbf{Cap}$.*

PROOF. By induction on the typing derivation. \square

LEMMA 20. *If $\Delta \vdash C' : \mathbf{Cap}$ and C is unique in C' then $\Delta \vdash C' = C'' \oplus C : \mathbf{Cap}$.*

PROOF. By induction on the typing derivation. \square

LEMMA (CAPABILITY EQUALITY CARDINALITY PRESERVATION(CECP)). *If $\Delta \vdash C_1 = C_2 : \mathbf{Cap}$ and $\Delta \vdash a : \mathbf{Cap}$ and $a = \epsilon$ or $\{r\}$ then*

- (1) *a is unique (duplicatable) in C_1 iff a is unique (duplicatable) in C_2 .*
- (2) *The number of unique occurrences of a is the same in C_1 and C_2 .*

PROOF. By induction on the derivation. \square

LEMMA (CAPABILITY SUBTYPING CARDINALITY PRESERVATION(CSCP)). *If $\cdot \vdash C_1 \leq C_2 : \text{Cap}$ then*

- (1) *For all region names ν , $\{\nu\} \in C_1$ iff $\{\nu\} \in C_2$.*
- (2) *For all region names ν , if $\{\nu\}$ is not unique in C_1 then $\{\nu\}$ is not unique in C_2 .*

PROOF. By induction on the derivation and Capability Equality Cardinality Preservation. Note that by Subtyping Regularity and Lemma 12, no type variables ϵ appear in C_1 , and consequently, the rule `sub-var` never appears in the derivation. \square

LEMMA (TYPE SUBSTITUTION). *Let Δ_0 be $\Delta[c_1, \dots, c_n/\Delta']$ Γ_0 be $\Gamma[c_1, \dots, c_n/\Delta']$ C_0 be $C[c_1, \dots, c_n/\Delta']$ r_0 be $r[c_1, \dots, c_n/\Delta']$ τ_0 be $\tau[c_1, \dots, c_n/\Delta']$. If Δ' is b_1, \dots, b_n where for $1 \leq i \leq n$:*

- A1. *if b_i is $\alpha_i : \kappa_i$ then $\cdot \vdash c_i : \kappa_i$*
- A2. *if b_i is $\epsilon_i \leq C_i$ then $\cdot \vdash c_i \leq C_i : \text{Cap}$*

then

- (1) *If $\Delta', \Delta \vdash \Delta''$ then $\Delta_0 \vdash \Delta''[c_1, \dots, c_n/\Delta']$*
- (2) *If $\Delta', \Delta \vdash c : \kappa$ then $\Delta_0 \vdash c[c_1, \dots, c_n/\Delta'] : \kappa$*
- (3) *If $\Delta', \Delta \vdash c = c' : \kappa$ then $\Delta_0 \vdash c[c_1, \dots, c_n/\Delta'] = c'[c_1, \dots, c_n/\Delta'] : \kappa$*
- (4) *If $\Delta', \Delta \vdash c \leq c' : \kappa$ then $\Delta_0 \vdash c[c_1, \dots, c_n/\Delta'] \leq c'[c_1, \dots, c_n/\Delta'] : \kappa$*
- (5) *If $\Delta', \Delta \vdash \Delta_1 = \Delta_2$ then $\Delta_0 \vdash \Delta_1[c_1, \dots, c_n/\Delta'] = \Delta_2[c_1, \dots, c_n/\Delta']$*
- (6) *If $\Psi; \Delta', \Delta; \Gamma; r \vdash h : \tau$ then $\Psi; \Delta_0; \Gamma_0; r_0 \vdash h[c_1, \dots, c_n/\Delta'] : \tau_0$*
- (7) *If $\Psi; \Delta', \Delta; \Gamma \vdash v : \tau$ then $\Psi; \Delta_0; \Gamma_0 \vdash v[c_1, \dots, c_n/\Delta'] : \tau_0$*
- (8) *If $\Psi; \Delta', \Delta; \Gamma; C \vdash d \Rightarrow \Delta', \Delta''; \Gamma''; C''$ then $\Psi; \Delta_0; \Gamma_0; C_0 \vdash d[c_1, \dots, c_n/\Delta'] \Rightarrow (\Delta''; \Gamma''; C'')[c_1, \dots, c_n/\Delta']$*
- (9) *If $\Psi; \Delta', \Delta; \Gamma; C \vdash e$ then $\Psi; \Delta_0; \Gamma_0; C_0 \vdash e[c_1, \dots, c_n/\Delta']$.*

PROOF. By induction on the derivations. Almost all cases follow directly from the IH. In part 2, we must prove our lemma for the following rules.

$$\frac{}{\Delta', \Delta \vdash \alpha : \kappa} (\Delta', \Delta(\alpha) = \kappa) \quad \frac{}{\Delta', \Delta \vdash \epsilon : \text{Cap}} ((\epsilon \leq C) \in \Delta', \Delta)$$

In the first case, we have our result by A1 and Type Context Extension. In the second case, assume ϵ is ϵ_i . By A2, we have $\cdot \vdash c_i \leq C_i : \text{Cap}$. Because $\cdot \vdash \cdot$, Subtyping Regularity tells us that $\cdot \vdash c_i : \text{Cap}$. By Type Context Extension, we know that $\Delta_0 \vdash c_i : \text{Cap}$. In part 4, the lemma follows for the rule

$$\frac{}{\Delta', \Delta \vdash \epsilon \leq C : \text{Cap}} ((\epsilon \leq C) \in \Delta', \Delta)$$

by A2 and Type Context Extension. In part 9, the case for `let`, we can apply the induction hypothesis because inspection of the rules for declarations show that $\Psi; \Delta; \Gamma; C \vdash d \Rightarrow \Delta, \Delta''; \Gamma''; C'$ instead of the more general $\Psi; \Delta; \Gamma; C \vdash d \Rightarrow \Delta''; \Gamma''; C'$. \square

LEMMA (VALUE SUBSTITUTION). *If Γ is $\{x_1:\tau_1, \dots, x_n:\tau_n\}$, $\cdot \vdash \Gamma$ and for $1 \leq i \leq n$, $\Psi; \cdot; \cdot \vdash v_i : \tau_i$ then*

- (1) *If $\Psi; \Delta; \Gamma, \Gamma' \vdash h$ at $r : \tau$ then $\Psi; \Delta; \Gamma' \vdash h[v_1, \dots, v_n/x_1, \dots, x_n]$ at $r : \tau$*
- (2) *If $\Psi; \Delta; \Gamma, \Gamma' \vdash v : \tau$ then $\Psi; \Delta; \Gamma' \vdash v[v_1, \dots, v_n/x_1, \dots, x_n] : \tau$*
- (3) *If $\Psi; \Delta; \Gamma, \Gamma'; C \vdash d \Rightarrow \Delta'; \Gamma, \Gamma''; C'$
then $\Psi; \Delta; \Gamma'; C \vdash d[v_1, \dots, v_n/x_1, \dots, x_n] \Rightarrow \Delta'; \Gamma''; C'$*
- (4) *If $\Psi; \Delta; \Gamma, \Gamma'; C \vdash e$ then $\Psi; \Delta; \Gamma'; C \vdash e[v_1, \dots, v_n/x_1, \dots, x_n]$.*

PROOF. By induction on the typing derivations. In part 4, the case for `let`, we can use the induction hypothesis because inspection of the typing rules for declarations reveals that $\Psi; \Delta; \Gamma; C \vdash d \Rightarrow \Delta'; \Gamma, \Gamma'; C'$ instead of the more general $\Psi; \Delta; \Gamma; C \vdash d \Rightarrow \Delta'; \Gamma'; C'$. \square

LEMMA (TERM JUDGMENT REGULARITY). *If*

- A1 $\vdash \Psi$
- A2 $\cdot \vdash C : \text{Cap}$
- A3 $\cdot \vdash r : \text{Rgn}$

then

- (1) *If $\Psi; \cdot; \cdot \vdash v : \tau$ then $\cdot \vdash \tau : \text{Type}$*
- (2) *If $\Psi; \cdot; \cdot \vdash h$ at $r : \tau$ then $\cdot \vdash \tau : \text{Type}$*
- (3) *If $\Psi; \cdot; \cdot; C \vdash d \Rightarrow \Delta'; \Gamma'; C'$ then $\cdot \vdash \Delta'$ and $\Delta' \vdash \Gamma'$ and $\Delta' \vdash C' : \text{Cap}$.*

PROOF. By induction on the typing derivations. Almost all cases follow directly from the induction hypothesis and Equality Regularity or Subtyping Regularity. In part 1, consider the case for type application:

$$\frac{\Psi; \cdot; \cdot \vdash v : \forall[\alpha:\kappa, \Delta].(C', \tau_1, \dots, \tau_n) \rightarrow 0 \text{ at } r \quad \cdot \vdash c : \kappa}{\Psi; \cdot; \cdot \vdash v[c] : (\forall[\Delta].(C', \tau_1, \dots, \tau_n) \rightarrow 0)[c/\alpha] \text{ at } r}$$

By the induction hypothesis, and then inspection of the typing rules for arrow types, we can deduce a judgment of the following form:

$$\frac{\frac{\alpha:\kappa \vdash \Delta}{\cdot \vdash \alpha:\kappa, \Delta} \quad \alpha:\kappa, \Delta \vdash C' : \text{Cap} \quad \alpha:\kappa, \Delta \vdash \tau_i : \text{Type} \quad (\text{for } 1 \leq i \leq n)}{\cdot \vdash \forall[\alpha:\kappa, \Delta].(C', \tau_1, \dots, \tau_n) \rightarrow 0 \text{ at } r : \text{Type}}$$

By Type Substitution, we may deduce that $\cdot \vdash (\forall[\Delta'].(C', \tau_1, \dots, \tau_n) \rightarrow 0)[c/\alpha] \text{ at } r$. The second type application rule follows similarly. \square

LEMMA (CANONICAL MEMORY FORMS). *If*

$\vdash \{\nu_1 \mapsto R_1, \dots, \nu_n \mapsto R_n\} : \{\nu_1:\Upsilon_1, \dots, \nu_n:\Upsilon_n\}$ then for all $1 \leq i \leq n$ and for all $\ell \in \text{Dom}(\Upsilon_i)$, either

- (1) $\Upsilon_i(\ell)$ is $\langle \tau_1, \dots, \tau_m \rangle$ at ν_i and $R_i(\ell) = \langle v_1, \dots, v_m \rangle$ and for $1 \leq j \leq m$, $\Psi; \cdot; \cdot \vdash v_j : \tau_j$ or
- (2) $\Upsilon_i(\ell)$ is $\forall[\Delta].(C, \tau_1, \dots, \tau_n) \rightarrow 0$ at ν_i
and $R_i(\ell) = \text{fix } f[\Delta](C, x_1:\tau_1, \dots, x_n:\tau_n).e$
and $\Psi; \Delta; \{f:\forall[\Delta].(C, \tau_1, \dots, \tau_n) \rightarrow 0 \text{ at } \nu_i, x_1:\tau_1, \dots, x_n:\tau_n\}; C \vdash e$.

PROOF. By inspection of the typing judgments for memory, regions, and heap values. \square

LEMMA (CANONICAL MEMORY FORMS II). (1) If $\vdash M : \Psi$ and $\nu \notin M$ then $\nu \notin \Psi$

- (2) If $\Psi \vdash R$ at $\nu' : \Upsilon$ and $\nu \notin R$ and ν' is not ν then $\nu \notin \Upsilon$
- (3) If $\Psi; \cdot; \cdot \vdash h$ at $\nu' : \tau$ and ν' is not ν and $\nu \notin h$ then $\nu \notin \tau$
- (4) If $\Psi; \cdot; \cdot \vdash v : \tau$ and $\nu \notin v$ then $\nu \notin \tau$.

PROOF. By induction on the typing derivations. \square

LEMMA (CANONICAL FORMS). If $\vdash M : \Psi$ and $\Psi; \cdot; \cdot \vdash v : \tau$ then

- (1) If τ is `int` then $v = i$
- (2) If τ is ν `handle` then $v = \text{handle}(\nu)$
- (3) If τ is $\forall[\Delta].(C, \tau_1, \dots, \tau_n) \rightarrow 0$ at ν then
 - (a) $v = \nu.\ell[c_1, \dots, c_m]$
 - (b) $M(\nu.\ell) = \text{fix } f[\Delta', \Delta''](C', x_1:\tau'_1, \dots, x_n:\tau'_n).e$
 - (c) Δ' is b_1, \dots, b_m and for $0 \leq i \leq m$, either b_i is $\alpha_i:\kappa_i$ and $\cdot \vdash c_i:\kappa_i$, or b_i is $\epsilon_i \leq C_i$ and $\cdot \vdash c_i \leq C_i : \text{Cap}$.
 - (d) $\cdot \vdash \Delta = \Delta''[c_1, \dots, c_m/\Delta']$, and $\Delta \vdash C = C'[c_1, \dots, c_m/\Delta']$, and for $1 \leq i \leq n$, $\Delta \vdash \tau_i = \tau'_i[c_1, \dots, c_m/\Delta'] : \text{Type}$
 - (e) $\Psi; \Delta', \Delta''; \{f:\forall[\Delta', \Delta''].(C', \tau'_1, \dots, \tau'_n) \rightarrow 0 \text{ at } \nu, x_1:\tau'_1, \dots, x_n:\tau'_n\}; C' \vdash e$ or instead of (b),(c),(d), and (e), $\nu \notin \Psi$
- (4) If τ is $\langle \tau_1, \dots, \tau_n \rangle$ at ν then
 - (a) $v = \nu.\ell$
 - (b) $M(\nu.\ell) = \langle v_1, \dots, v_n \rangle$
 - (c) $\Psi; \cdot \vdash v_i : \tau_i$
 - or instead of (b),(c), $\nu \notin \Psi$.

PROOF. Part 1, 2 follow by inspection of the typing rules for word values. Part 3 follows by induction on the derivation, $\Psi; \cdot; \cdot \vdash v : \tau$. By Canonical Memory Forms and inspection of the typing rules for word values, either $\nu.\ell$ or one of the type application rules are last.

case $\nu.\ell$ where $\nu \notin \Psi$:

- (a) Trivial.

case $\nu.\ell$ where $\nu \in \Psi$:

- (a) Trivial.
- (b) By Canonical Memory Forms where Δ' is \cdot , Δ'' is Δ , C' is C , and for $1 \leq i \leq n$, τ'_i is τ_i .
- (c) Trivial.
- (d) Trivial.
- (e) By inspection of judgment.

case $v[C]$

$$\frac{\Psi; \cdot; \cdot \vdash v : \forall[\epsilon \leq C_a, \Delta].(C_b, \tau_1, \dots, \tau_n) \rightarrow 0 \text{ at } r \quad \cdot \vdash C \leq C_a : \text{Cap}}{\Psi; \cdot; \cdot \vdash v[C] : (\forall[\Delta].(C_b, \tau_1, \dots, \tau_n) \rightarrow 0)[C/\epsilon] \text{ at } r}$$

By Term Judgment Regularity and Lemma 12, r is ν . The inductive hypothesis is as follows:

- (a) $v = \nu.\ell[c_1, \dots, c_n]$
 - (b) $M(\nu.\ell) = \mathbf{fix} f[\Delta', \epsilon \leq C'_a, \Delta''](C'_b, x_1:\tau'_1, \dots, x_n:\tau'_n).e$
 - (c) Δ' is b_1, \dots, b_m and for $0 \leq i \leq m$, either b_i is $\alpha_i:\kappa_i$ and $\cdot \vdash c_i:\kappa_i$, or b_i is $\epsilon_i \leq C_i$ and $\cdot \vdash c_i \leq C_i : \mathbf{Cap}$.
 - (d) $\cdot \vdash \epsilon \leq C_a, \Delta = (\epsilon \leq C'_a, \Delta'')[c_1, \dots, c_m/\Delta']$
and $\epsilon \leq C_a, \Delta \vdash C_b = C'_b[c_1, \dots, c_m/\Delta']$
and for $1 \leq i \leq n$, $\epsilon \leq C_a, \Delta \vdash \tau_i = \tau'_i[c_1, \dots, c_m] : \mathbf{Type}$
 - (e) $\Psi; \Delta', \epsilon \leq C_a, \Delta'', \Gamma; C'_b \vdash e$
where $\Gamma = \{f:\forall[\Delta', \epsilon \leq C_a, \Delta''].(C'_b, \tau'_1, \dots, \tau'_n) \rightarrow \mathbf{0} \text{ at } \nu, x_1:\tau'_1, \dots, x_n:\tau'_n\}$
- or instead of (b),(c),(d), and (e), $\nu \notin \Psi$. Thus,
- (a) $v[C] = \nu.\ell[c_1, \dots, c_n, C]$ from IH.
 - If $\nu \notin \Psi$ then the result follows trivially. Thus assume $\nu \in \Psi$.
 - (b) By IH.
 - (c) By IH and the typing judgement which states $\cdot \vdash C \leq C_a : \mathbf{Cap}$.
 - (d) By IH and Type Substitution.
 - (e) By IH.

case $v[c]$ Similar.

Part 4 follows by inspection of the typing rules for word values. Notice only the $\nu.\ell$ rule when $\nu \in \Psi$, or the rule for tuples when $\nu \notin \Psi$ can apply. Assuming the former (the latter is trivial), then (a) is immediate and (b), (c) follow by Canonical Memory Forms. \square

LEMMA (MEMORY TYPE GC). *If $\vdash \Psi$ and Ψ' is $\Psi \setminus \nu$ then*

- (1) *If $\Psi; \Delta; \Gamma \vdash h \text{ at } r : \tau$ then $\Psi'; \Delta; \Gamma \vdash h \text{ at } r : \tau$*
- (2) *If $\Psi; \Delta; \Gamma \vdash v : \tau$ then $\Psi'; \Delta; \Gamma \vdash v : \tau$*
- (3) *If $\Psi; \Delta; \Gamma; C \vdash d \Rightarrow \Delta'; \Gamma'; C'$ then $\Psi'; \Delta; \Gamma; C \vdash d \Rightarrow \Delta'; \Gamma'; C'$*
- (4) *If $\Psi; \Delta; \Gamma; C \vdash e$ then $\Psi'; \Delta; \Gamma; C \vdash e$*
- (5) *If $\Psi \vdash R \text{ at } \nu : \Upsilon$ then $\Psi' \vdash R \text{ at } \nu : \Upsilon$.*

PROOF. By induction on the typing derivation. All cases follow directly from IH except the rule

$$\frac{}{\Psi; \Delta; \Gamma \vdash \nu'.\ell : \tau} (\Psi(\nu'.\ell) = \tau).$$

When ν is not ν' this case is trivial so assume ν is ν' . By Canonical Memory Forms, τ is either $\forall[\Delta'].(\epsilon, \tau_1, \dots, \tau_n) \rightarrow \mathbf{0} \text{ at } \nu$ or $\langle \tau_1, \dots, \tau_n \rangle \text{ at } \nu$. Because $\vdash \Psi$, we have $\cdot \vdash \tau : \mathbf{Type}$. By Type Context Extension, $\Delta \vdash \tau : \mathbf{Type}$. Thus, in either of the above cases, $\Psi; \Delta; \Gamma \vdash \nu'.\ell : \tau$ via one of the two rules for $\nu \notin \text{Dom}(\Psi)$. \square

LEMMA (MEMORY TYPE EXTENSION). *If ν does not appear in $\Psi, \Delta, \Gamma, r, h, v, d, e$, or R , and Ψ' is $\Psi\{\nu:\{\}\}$ then*

- (1) *If $\Psi; \Delta; \Gamma \vdash h \text{ at } r : \tau$ then $\Psi'; \Delta; \Gamma \vdash h \text{ at } r : \tau$*
- (2) *If $\Psi; \Delta; \Gamma \vdash v : \tau$ then $\Psi'; \Delta; \Gamma \vdash v : \tau$*
- (3) *If $\Psi; \Delta; \Gamma; C \vdash d \Rightarrow \Delta'; \Gamma'; C'$ then $\Psi'; \Delta; \Gamma; C \vdash d \Rightarrow \Delta'; \Gamma'; C'$*

- (4) If $\Psi; \Delta; \Gamma; C \vdash e$ then $\Psi'; \Delta; \Gamma; C \vdash e$
 (5) If $\Psi \vdash R$ at $\nu' : \Upsilon$ then $\Psi' \vdash R$ at $\nu' : \Upsilon$.

PROOF. By induction on the typing derivation. In part 2, for the rule

$$\frac{\Delta \vdash \langle \tau_1, \dots, \tau_n \rangle \text{ at } \nu' : \text{Type}}{\Psi; \Delta; \Gamma \vdash \nu'.\ell : \langle \tau_1, \dots, \tau_n \rangle \text{ at } \nu'} \quad (\nu' \notin \Psi)$$

ν' is not ν by assumption and thus the result holds, and similarly for the analogous rule for arrow types. \square

LEMMA (REGION TYPE EXTENSION). *If $\nu \in \text{Dom}(\Psi)$, $\ell \notin \text{Dom}(\Psi(\nu))$, and Ψ' is $\Psi\{\nu.\ell:\tau\}$ then*

- (1) If $\Psi; \Delta; \Gamma \vdash h$ at $r : \tau$ then $\Psi'; \Delta; \Gamma \vdash h$ at $r : \tau$
 (2) If $\Psi; \Delta; \Gamma \vdash v : \tau$ then $\Psi'; \Delta; \Gamma \vdash v : \tau$
 (3) If $\Psi; \Delta; \Gamma; C \vdash d \Rightarrow \Delta'; \Gamma'; C'$ then $\Psi'; \Delta; \Gamma; C \vdash d \Rightarrow \Delta'; \Gamma'; C'$
 (4) If $\Psi; \Delta; \Gamma; C \vdash e$ then $\Psi'; \Delta; \Gamma; C \vdash e$
 (5) If $\Psi \vdash R$ at $\nu' : \Upsilon$ then $\Psi' \vdash R$ at $\nu' : \Upsilon$
 (6) If $\Psi \vdash C$ sat then $\Psi' \vdash C$ sat.

PROOF. By induction on the typing derivation. \square

LEMMA 32. *If $\Delta \vdash C_1 \oplus C_2 : \text{Cap}$ and $\Delta \vdash C_1 \oplus C_2 = a_1 \oplus \dots \oplus a_n : \text{Cap}$ then $\Delta \vdash C_1 = a'_1 \oplus \dots \oplus a'_m : \text{Cap}$ and a'_1, \dots, a'_m is a subset of a_1, \dots, a_n .*

PROOF. By Lemma 17 (1), $\Delta \vdash C_1 = a'_1 \oplus \dots \oplus a'_m : \text{Cap}$. By Lemma 17 (5), $E(a'_1 \oplus \dots \oplus a'_m) = E(C_1) \subseteq E(C_1 \oplus C_2) = E(a_1 \oplus \dots \oplus a_n)$. By CECP, a_i is unique (duplicatable) in $a'_1 \oplus \dots \oplus a'_m$ if and only if a_i is unique (duplicatable) in $a_1 \oplus \dots \oplus a_n$. Therefore, $a \in a'_1, \dots, a'_m$ implies $a \in a_1, \dots, a_n$.

\square

LEMMA (CAPABILITY SATISFIABILITY PRESERVATION).

- (1) If $\Psi \vdash C$ sat and $\cdot \vdash C = C' : \text{Cap}$ then $\Psi \vdash C'$ sat.
 (2) If $\Psi \vdash C$ sat and $\cdot \vdash C \leq C' : \text{Cap}$ then $\Psi \vdash C'$ sat.

PROOF. (1) By symmetry and transitivity of equality and inspection of the sat derivation.

- (2) By induction on the subtyping derivation.

case equality: From Part 1.

case transitivity: By IH.

case $\cdot \vdash \epsilon \leq C$: Does not apply because the context (\cdot) is empty and therefore, by Lemma 12, the two capabilities must not contain any free variables.

case $\cdot \vdash C \leq \bar{C}$:

Assume:

$$\frac{\cdot \vdash C = \{\nu_1^{\varphi_1}, \dots, \nu_n^{\varphi_n}\} : \text{Cap} \quad (2)}{\{\nu_1:\Upsilon_1, \dots, \nu_n:\Upsilon_n\} \vdash C \text{ sat} \quad (1)} \quad (\nu_1, \dots, \nu_n \text{ distinct} \quad (3))$$

4. From 2, using rule eq-congruence-bar, $\cdot \vdash \overline{C} = \overline{\{\nu_1^{\varphi_1}, \dots, \nu_n^{\varphi_n}\}} : \text{Cap}$
 5. From 4, by constructor equality rules, $\cdot \vdash \overline{C} = \overline{\{\nu_1^+, \dots, \nu_n^+\}} : \text{Cap}$
 6. Hence, from 5,3 we have $\{\nu_1:\Upsilon_1, \dots, \nu_n:\Upsilon_n\} \vdash \overline{C} \text{ sat}$
- case

$$\frac{\cdot \vdash C_1 \leq C'_1 \text{ (1)} \quad \cdot \vdash C_2 \leq C'_2 \text{ (2)}}{\cdot \vdash C_1 \oplus C_2 \leq C'_1 \oplus C'_2 \text{ (0)}}$$

Assume:

$$\frac{\cdot \vdash C_1 \oplus C_2 = \{\nu_1^{\varphi_1}, \dots, \nu_n^{\varphi_n}\} : \text{Cap (4)}}{\{\nu_1:\Upsilon_1, \dots, \nu_n:\Upsilon_n\} \vdash C_1 \oplus C_2 \text{ sat (3)}} \quad (\nu_1, \dots, \nu_n \text{ distinct (5)})$$

6. By 4 and Equality Regularity, C_1, C_2, C'_1, C'_2 are all well-formed under the empty context \cdot .
7. By 6 and Lemma 17(1), $\cdot \vdash C_1 = \{\nu_{1,1}^{\varphi_{1,1}}, \dots, \nu_{1,n_1}^{\varphi_{1,n_1}}\} : \text{Cap}$
8. From 7 and CECP, for $1 \leq i \leq n_1$, $\nu_{1,i}^{\varphi_{1,i}} \in C_1$
9. Hence, we have $\nu_{1,i}^{\varphi_{1,i}} \in C_1 \oplus C_2$
10. and by 4 and CECP again, we can conclude $\nu_{1,i}^{\varphi_{1,i}} \in \{\nu_1^{\varphi_1}, \dots, \nu_n^{\varphi_n}\}$
11. For $1 \leq i \leq n_1$ and $i \neq j$, if $\nu_{1,i} = \nu_{1,j}$ and $\varphi_{1,i} = \varphi_{1,j} = +$ then by Lemma 17(4), we can eliminate the duplicates and assume without loss of generality that $\{\nu_{1,1}^{\varphi_{1,1}}, \dots, \nu_{1,n_1}^{\varphi_{1,n_1}}\}$ does not contain duplicates of this form.
12. For some $1 \leq i \leq n_1$ and $i \neq j$, assume (anticipating a contradiction) that $\nu_{1,i} = \nu_{1,j}$ and one of $\varphi_{1,i}$ or $\varphi_{1,j}$ is 1 then by CECP, $\{\nu_{1,i}^{\varphi_{1,i}}\}$ and $\{\nu_{1,j}^{\varphi_{1,j}}\}$ both appear in C_1
13. From 12, we have $\{\nu_{1,i}^{\varphi_{1,i}}\}$ and $\{\nu_{1,j}^{\varphi_{1,j}}\}$ both appear in $C_1 \oplus C_2$
14. and from 13, and CECP, $\{\nu_{1,i}^{\varphi_{1,i}}\}$ and $\{\nu_{1,j}^{\varphi_{1,j}}\}$ both appear in the capability $\{\nu_1^{\varphi_1}, \dots, \nu_n^{\varphi_n}\}$
15. However, 14 and 12 ($\nu_{1,i} = \nu_{1,j}$) contradict 5, indicating our assumption 12 was false
16. By 15 and 11, we may assume without loss of generality that $\nu_{1,1}, \dots, \nu_{1,n_1}$ distinct
17. By 7, 11, and 16, we can conclude $\Psi_1 \vdash C_1 \text{ sat}$ where Ψ_1 is Ψ with domain restricted to $\{\nu_{1,1}, \dots, \nu_{1,n_1}\}$
18. Analogous reasoning and definitions for Ψ_2 yields $\Psi_2 \vdash C_2 \text{ sat}$
19. From 17 and 1, using the inductive hypothesis, we have $\Psi_1 \vdash C'_1 \text{ sat}$
20. From 18 and 2, using the inductive hypothesis, we have $\Psi_2 \vdash C'_2 \text{ sat}$
21. From 19 (20) and inspection of the sat judgment, we know $\cdot \vdash C'_1 = \{\nu_{1',1}^{\varphi_{1',1}}, \dots, \nu_{1',n'_1}^{\varphi_{1',n'_1}}\} : \text{Cap}$ for some $\nu_{1',1}^{\varphi_{1',1}}, \dots, \nu_{1',n'_1}^{\varphi_{1',n'_1}}$ (and we know $\cdot \vdash C'_2 = \{\nu_{2',1}^{\varphi_{2',1}}, \dots, \nu_{2',n'_1}^{\varphi_{2',n'_1}}\} : \text{Cap}$ for some $\nu_{2',1}^{\varphi_{2',1}}, \dots, \nu_{2',n'_1}^{\varphi_{2',n'_1}}$)
22. We have

$$\begin{aligned} \cdot \vdash \{\nu_1^{\varphi_1}, \dots, \nu_n^{\varphi_n}\} &= C_1 \oplus C_2 && \text{(By 3)} \\ &\leq C'_1 \oplus C'_2 && \text{(By 0)} \\ &= \{\nu_{1',1}^{\varphi_{1',1}}, \dots, \nu_{1',n'_1}^{\varphi_{1',n'_1}}\} \oplus && \text{(By 21, eq-congruence)} \\ &\quad \{\nu_{2',1}^{\varphi_{2',1}}, \dots, \nu_{2',n'_1}^{\varphi_{2',n'_1}}\} \end{aligned}$$

23. From 22 and CSCP, we know $\nu \in \nu_1, \dots, \nu_n$ if and only if $\nu \in \nu_{1',1}, \dots, \nu_{1',n'_1}, \nu_{2',1}, \dots, \nu_{2',n'_1}$
24. From 22 and 5 and by reasoning analogously to steps 12 through 16, we can deduce that $\{\nu_{1',1}^{\varphi_{1',1}}, \dots, \nu_{1',n'_1}^{\varphi_{1',n'_1}}, \nu_{2',1}^{\varphi_{2',1}}, \dots, \nu_{2',n'_1}^{\varphi_{2',n'_1}}\}$ contains no duplicate region names aside from those which both have multiplicity flag +.
25. By Lemma 17 parts 3 and 4, we may eliminate duplicate region names with flag + and reorder them as we choose. Hence, by 23 and for some set of multiplicities $\varphi'_1, \dots, \varphi'_n$, we know:
 $\cdot \vdash \{\nu_{1',1}^{\varphi_{1',1}}, \dots, \nu_{1',n'_1}^{\varphi_{1',n'_1}}, \nu_{2',1}^{\varphi_{2',1}}, \dots, \nu_{2',n'_1}^{\varphi_{2',n'_1}}\} = \{\nu_1^{\varphi'_1}, \dots, \nu_n^{\varphi'_n}\} : \text{Cap}$
 (where ν_1, \dots, ν_n distinct)
26. From 22 and 25, $\cdot \vdash C'_1 \oplus C'_2 = \{\nu_1^{\varphi'_1}, \dots, \nu_n^{\varphi'_n}\} : \text{Cap}$
 (where ν_1, \dots, ν_n distinct)
27. From 26, $\{\nu_1 : \Upsilon_1, \dots, \nu_n : \Upsilon_n\} \vdash C'_1 \oplus C'_2 \text{ sat}$

□

LEMMA 34. *If $\Psi \vdash C \oplus \{\nu\}$ sat then $\Psi \setminus \nu \vdash C$ sat.*

PROOF. 1. Assume $\Psi \vdash C \oplus \{\nu\}$ sat

2. and $\Psi = \{\nu_1 : \Upsilon_1, \dots, \nu_n : \Upsilon_n\}$
3. From 1, we know $\cdot \vdash C \oplus \{\nu\} = \{\nu_1^{\varphi_1}, \dots, \nu_n^{\varphi_n}\} : \text{Cap}$
4. and ν_1, \dots, ν_n distinct
5. From 3, 4, and CSCP, ν appears once in $\{\nu_1^{\varphi_1}, \dots, \nu_n^{\varphi_n}\}$ and once in $C \oplus \{\nu\}$.
6. From 5, and Equality Lemma, part 3,
 $\cdot \vdash \{\nu_1^{\varphi_1}, \dots, \nu_{i-1}^{\varphi_{i-1}}, \nu^1, \nu_{i+1}^{\varphi_{i+1}}, \dots, \nu_n^{\varphi_n}\} = \{\nu_1^{\varphi_1}, \dots, \nu_{i-1}^{\varphi_{i-1}}, \nu_{i+1}^{\varphi_{i+1}}, \dots, \nu_n^{\varphi_n}\} \oplus \{\nu\} : \text{Cap}$
7. From 3, 6, transitivity of equality, and Equality Lemma, part 7,
 $\cdot \vdash C = \{\nu_1^{\varphi_1}, \dots, \nu_{i-1}^{\varphi_{i-1}}, \nu_{i+1}^{\varphi_{i+1}}, \dots, \nu_n^{\varphi_n}\} : \text{Cap}$
8. Hence, from 2,4,7 we have $\Psi \setminus \nu \vdash C$ sat □

LEMMA (PRESERVATION). *If $\vdash (M, e)$ and $(M, e) \mapsto (M', e')$ then $\vdash (M', e')$*

PROOF. The proof proceeds by cases on the structure of e . In each case, we show the form of the typing judgment that can be inferred by inspection of the typing rules and refer to it throughout the case as “the typing judgment”. Then we give the transition specified by the operational semantics. Using these two facts, we derive the result $\vdash (M', e')$.

—let v

$$\frac{\frac{\Psi; \cdot; \cdot \vdash v : \tau}{\Psi; \cdot; \cdot; C \vdash x = v \Rightarrow \cdot; \{x : \tau\}; C} \quad \Psi; \cdot; \{x : \tau\}; C \vdash e}{\vdash M : \Psi \quad \Psi \vdash C \text{ sat} \quad \Psi; \cdot; \cdot; C \vdash \text{let } x = v \text{ in } e} \vdash (M, \text{let } x = v \text{ in } e)$$

and $(M, \text{let } x = v \text{ in } e) \mapsto (M, e[v/x])$. By the typing judgment and Value Substitution, $\vdash (M, e[v/x])$.

—let h

$$\frac{\vdash M : \Psi \quad \Psi \vdash C \text{ sat} \quad (A)}{\vdash (M, \text{let } x = h \text{ at } v \text{ in } e)}$$

$$\frac{\frac{\Psi; \cdot; \cdot \vdash v : \nu \text{ handle}}{\Psi; \cdot; \cdot \vdash h \text{ at } \nu : \tau} \quad \cdot \vdash C \leq C' \oplus \{\nu\} \quad \vdots}{\frac{\Psi; \cdot; \cdot; C \vdash x = h \text{ at } v \Rightarrow \cdot; \{x:\tau\}; C \quad \Psi; \cdot; \{x:\tau\}; C \vdash e}{\Psi; \cdot; \cdot; C \vdash \text{let } x = h \text{ at } v \text{ in } e}} \quad (A)$$

where $v = \text{handle}(\nu)$

and $\nu \in \text{Dom}(M)$ and $\ell \notin \text{Dom}(M(\nu))$

and $M' = M\{\nu.\ell \mapsto h\}$

and $(M, \text{let } x = h \text{ at } v \text{ in } e) \mapsto (M', e[\nu.\ell/x])$

and let $\Psi' = \Psi\{\nu.\ell:\tau\}$

(1) (a) τ is $\langle \tau_1, \dots, \tau_n \rangle$ at ν or $\forall[].(C'', \tau_1, \dots, \tau_n) \rightarrow 0$ at ν by inspection of the heap value typing rules and the typing judgment.

(b) $\vdash M' : \Psi'$ by the typing judgment and inspection of the memory typing rule.

(2) $\Psi' \vdash C$ sat by Region Type Extension.

(3) (a) $\Psi'; \cdot; \cdot \vdash \nu.\ell : \tau$ by the typing rules for word values.

(b) $\Psi'; \cdot; \cdot; C \vdash e[\nu.\ell/x]$ by (a) and Value Substitution.

By 1(b), 2, and 3(b), we have $\vdash (M', e[\nu.\ell/x])$.

— $\pi_i v$

$$\frac{\frac{\Psi; \cdot; \cdot \vdash v : \langle \tau_1, \dots, \tau_n \rangle \text{ at } \nu}{\cdot \vdash C \leq C' \oplus \{\nu\}}}{\frac{\Psi; \cdot; \cdot; C \vdash x = \pi_i v \Rightarrow \cdot; \{x:\tau_i\}; C \quad \Psi; \cdot; \{x:\tau_i\}; C \vdash e}{\Psi; \cdot; \cdot; C \vdash \text{let } x = \pi_i v \text{ in } e}} \quad \frac{\vdash M : \Psi \quad \Psi \vdash C \text{ sat}}{\vdash (M, \text{let } x = \pi_i v \text{ in } e)}$$

where $v = \nu.\ell$

and $M(\nu.\ell) = \langle v_1, \dots, v_n \rangle$

and $(M, \text{let } x = \pi_i v \text{ in } e) \mapsto (M, e[v_i/x])$

(1) $\vdash M : \Psi$ by the typing judgment.

(2) $\Psi \vdash C$ sat by the typing judgment.

(3) (a) $\Psi; \cdot; \cdot \vdash v_i : \tau_i$ by Canonical Forms and the typing judgment.

(b) $\Psi; \cdot; \cdot; C \vdash e[v_i/x]$ by Value Substitution, (a), and the typing judgment.

By 1, 2, and 3(b), $\vdash (M, e[v_i/x])$.

—**freergn**

$$\frac{\frac{\Psi; \cdot; \cdot \vdash v : \nu \text{ handle} \quad \cdot \vdash C \leq C' \oplus \{\nu\}}{\Psi; \cdot; \cdot; C \vdash \text{freergn } v \Rightarrow \cdot; \cdot; C'} \quad \Psi; \cdot; \cdot; C' \vdash e}{\frac{\vdash M : \Psi \quad \Psi \vdash C \text{ sat}}{\Psi; \cdot; \cdot; C \vdash \text{freergn } v \text{ in } e}} \quad \frac{\vdash M : \Psi \quad \Psi \vdash C \text{ sat}}{\vdash (M, \text{freergn } v \text{ in } e)}$$

where v is $\text{handle}(\nu)$ and $(M, \text{freergn } v \text{ in } e) \mapsto (M \setminus \nu, e)$. Let Ψ' be $\Psi \setminus \nu$.

(1) $\vdash M' : \Psi'$ by Memory Type GC and the typing judgment.

(2) (a) $\Psi \vdash C$ sat and $\cdot \vdash C \leq C' \oplus \{\nu\}$ by the typing judgment.

(b) $\Psi \vdash C' \oplus \{\nu\}$ sat by Capability Satisfiability Preservation and (a)

(c) $\Psi' \vdash C'$ sat by Lemma 34 and (b)

(3) $\Psi'; \cdot; \cdot; C' \vdash e$ by the typing judgment and Memory Type GC.
 By 1, 2(e), and 3, $\vdash (M \setminus \nu, e)$.

—**newrgn**

$$\frac{\cdot \vdash C = C' : \mathbf{Cap} \quad (\dots) \quad \frac{(A) \quad (B)}{\Psi; \cdot; \cdot; C \vdash \mathbf{newrgn} \rho, x_\rho \text{ in } e}}{\vdash M : \Psi \quad \Psi \vdash C \text{ sat}}}{\vdash (M, \mathbf{newrgn} \rho, x_\rho \text{ in } e)}$$

$$\frac{}{\Psi; \cdot; \cdot; C \vdash \mathbf{newrgn} \rho, x_\rho \Rightarrow \rho : \mathbf{Rgn}; \{x_\rho : \rho \text{ handle}\}; C \oplus \{\rho\}} \quad (A)$$

$$\frac{\vdots}{\Psi; \rho : \mathbf{Rgn}; \{x_\rho : \rho \text{ handle}\}; C \oplus \{\rho\} \vdash e} \quad (B)$$

The operational rule is

$$(M, \mathbf{newrgn} \rho, x_\rho \text{ in } e) \mapsto (M', e[\nu, \mathbf{handle}(\nu)/\rho, x_\rho])$$

where $M' = M\{\nu \mapsto \{\}\}$ and $\nu \notin M$ and $\nu \notin e$.

In what follows, let $\Psi' = \Psi\{\nu : \{\}\}$.

- (1) $\vdash \Psi'$ by Memory Type Extension and the typing judgment.
- (2) Since $\nu \notin \Psi$ by assumption in the operational semantics, we can satisfy the side condition on the sat judgment. We can also prove $\cdot \vdash C \oplus \{\nu\} = C' \oplus \{\nu\} : \mathbf{Cap}$ by the congruence rule for equality. Consequently, $\Psi' \vdash C \oplus \{\nu\} \text{ sat}$.
- (3) $\Psi'; \cdot; \cdot; C \oplus \{\nu\} \vdash e[\nu, \mathbf{handle}(\nu)/\rho, x_\rho]$ from the typing judgment and application of Type and Value Substitution and then Memory Type Extension Lemmas.

By 1, 2(e), and 3, $\vdash (M', e[\nu, \mathbf{handle}(\nu)/\rho, x_\rho])$.

—**if0**

$$\frac{\vdash M : \Psi \quad \Psi \vdash C \text{ sat} \quad \frac{\Psi; \cdot; \cdot \vdash i : \mathbf{int} \quad \Psi; \cdot; \cdot; C \vdash e_2 \quad \Psi; \cdot; \cdot; C \vdash e_3}{\Psi; \cdot; \cdot; C \vdash \mathbf{if0} i \text{ then } e_2 \text{ else } e_3}}{\vdash (M, \mathbf{if0} i \text{ then } e_2 \text{ else } e_3)}$$

$(M, e) \mapsto (M, e_2)$ if $i = 0$ and $(M, e) \mapsto (M, e_3)$ otherwise. By the typing judgment, $\vdash (M, e_2)$, or $\vdash (M, e_3)$.

— $v_0(v_1, \dots, v_n)$

$$\frac{\Psi; \cdot; \cdot \vdash v_i : \tau_i \quad (\text{for } 0 \leq i \leq n) \quad \cdot \vdash C \leq C'' \oplus \{\nu\} : \mathbf{Cap} \quad \cdot \vdash \tau_0 = \forall[\cdot]. (C', \tau_1, \dots, \tau_n) \rightarrow 0 \text{ at } \nu \quad \cdot \vdash C \leq C' : \mathbf{Cap}}{\vdash M : \Psi \quad \Psi \vdash C \text{ sat} \quad \Psi; \cdot; \cdot; C \vdash v_0(v_1, \dots, v_n)}{\vdash (M, v_0(v_1, \dots, v_n))}$$

$(M, v_0(v_1, \dots, v_n)) \mapsto (M, S(e))$

where $v_0 = \nu.\ell[c_1, \dots, c_m]$

and $M(\nu.\ell) = \mathbf{fix} f[b_1, \dots, b_m](C'', x_1 : \tau_1, \dots, x_n : \tau_n).e$

and for $1 \leq i \leq m$, $b_i = \alpha_i : \kappa_i$ or $b_i = \alpha_i \leq C_i$

and $S = [c_1, \dots, c_m, \nu.\ell, v_1, \dots, v_n/\alpha_1, \dots, \alpha_m, f, x_1, \dots, x_n]$.

- (1) $\vdash M : \Psi$ by the typing judgment.
 - (2) $\Psi \vdash C'$ sat by Capability Satisfiability Preservation.
 - (3) (a) $\Psi \vdash C'' \oplus \{\nu\}$ sat by Capability Satisfiability Preservation and the typing judgment.
 - (b) $\nu \in \text{Dom}(\Psi)$ by CECP and (a).
 - (c) $\Psi; \cdot; \cdot \vdash v_0 : \forall[].(C''', \tau_1''', \dots, \tau_n''') \rightarrow 0$ and $\cdot \vdash \forall[].(C''', \tau_1''', \dots, \tau_n''') \rightarrow 0 = \forall[].(C', \tau_1, \dots, \tau_n) \rightarrow 0$ by the typing judgment.
 - (d) $\Psi; b_1, \dots, b_m; \{x_1:\tau_1, \dots, x_n:\tau_n\}; C'' \vdash e$ by Canonical Forms 3(e), (b), and (c).
 - (e) $\cdot \vdash C' = C''[c_1, \dots, c_m/\alpha_1, \dots, \alpha_m] : \text{Cap}$ by the transitivity of equality, Canonical Forms 3(d), (b), and (c).
 - (f) $\Psi; \cdot; \cdot; C' \vdash S(e)$ by Type and Value Substitution, and (e).
- By 1, 2, and 3(f), $\vdash (M, S(e))$

□

LEMMA (PROGRESS). *If $\vdash (M, e)$ then either:*

- (1) *There exists (M', e') such that $(M, e) \mapsto (M', e')$ or*
- (2) *$e = \text{halt } v$ and $\Psi; \cdot; \cdot \vdash v : \text{int}$.*

PROOF. The proof proceeds by cases on the structure of e and makes heavy use of the Canonical Forms lemma.

—**let** v Trivial.

—**let** h

$$\frac{\Psi; \cdot; \cdot \vdash v : r \text{ handle} \quad \Psi; \cdot; \cdot \vdash h \text{ at } r : \tau \quad \cdot \vdash C \leq C' \oplus \overline{\{r\}}}{\Psi; \cdot; \cdot; C \vdash x = h \text{ at } v \Rightarrow \cdot; \{x:\tau\}; C \quad \dots} \quad \frac{\vdash M : \Psi \quad \Psi \vdash C \text{ sat} \quad \Psi; \cdot; \cdot; C \vdash \text{let } x = h \text{ at } v \text{ in } e}{\vdash (M, \text{let } x = h \text{ at } v \text{ in } e)}$$

$\Psi; \cdot; \cdot \vdash v : r \text{ handle}$ directly from the typing judgment. By Term Judgment Regularity and Lemma 12, r is ν , and by Canonical Forms, v is $\text{handle}(\nu)$. By Capability Satisfiability Preservation, $\Psi \vdash C' \oplus \overline{\{\nu\}}$ sat and thus $\nu \in \text{Dom}(\Psi)$. By inspection of the memory typing rules, $\nu \in \text{Dom}(M)$. Thus $(M, e) \mapsto (M\{\nu.\ell \mapsto h\}, e[\nu.\ell/x])$.

— $\pi_i v$

$$\frac{\Psi; \cdot; \cdot \vdash v : \langle \tau_1, \dots, \tau_n \rangle \text{ at } r \quad \cdot \vdash C \leq C' \oplus \overline{\{r\}}}{\Psi; \cdot; \cdot; C \vdash x = \pi_i v \Rightarrow \cdot; \{x:\tau_i\}; C \quad \dots} \quad \frac{\vdash M : \Psi \quad \Psi \vdash C \text{ sat} \quad \Psi; \cdot; \cdot; C \vdash \text{let } x = \pi_i v \text{ in } e}{\vdash (M, \text{let } x = \pi_i v \text{ in } e)}$$

By Capability Satisfiability Preservation, $\Psi \vdash C' \oplus \overline{\{\nu\}}$ sat. By CECP, $\nu \in \text{Dom}(\Psi)$ and by Canonical Forms, $M(\nu.\ell) = \langle v_1, \dots, v_n \rangle$. By the typing judgment, $1 \leq i \leq n$. Thus $(M, \text{let } x = \pi_i v \text{ in } e) \mapsto (M, e[v_i/x])$.

- newrgn Trivial.
- freergn

$$\frac{\frac{\Psi; \cdot; \cdot \vdash v : r \text{ handle } \cdot \vdash C \leq C' \oplus \{r\} : \text{Cap}}{\Psi; \cdot; \cdot; C \vdash \text{freergn } v \Rightarrow \cdot; \cdot; C'} \dots}{\frac{\vdash M : \Psi \quad \Psi \vdash C \text{ sat}}{\Psi; \cdot; \cdot; C \vdash \text{freergn } v \text{ in } e}}{\vdash (M, \text{freergn } v \text{ in } e)}$$

By Term Judgment Regularity and Lemma 12, r is ν , and by Canonical Forms, v is $\text{handle}(\nu)$. By Capability Satisfiability Preservation, $\Psi \vdash C \oplus \{\nu\}$ sat. Thus, by CECF, $\nu \in \text{Dom}(\Psi)$, and by inspection of the memory typing rules, $\nu \in \text{Dom}(M)$. Consequently, $(M, \text{freergn } v \text{ in } e) \mapsto (M \setminus \nu, e)$.

- if0

$$\frac{\Psi; \Delta; \Gamma \vdash v : \text{int} \quad \Psi; \Delta; \Gamma; C \vdash e_2 \quad \Psi; \Delta; \Gamma; C \vdash e_3}{\Psi; \Delta; \Gamma; C \vdash \text{if0 } v \text{ then } e_2 \text{ else } e_3}$$

By Canonical Forms, v must be integer. Therefore, one of the two operational rules for **if0** applies.

- $v_0(v_1, \dots, v_n)$

$$\frac{\frac{\Psi; \cdot; \cdot \vdash v_i : \tau_i \quad (\text{for } 0 \leq i \leq n) \quad \cdot \vdash C \leq C'' \oplus \{\tau\} : \text{Cap}}{\cdot \vdash \tau_0 = \forall[\cdot](C'', \tau_1, \dots, \tau_n) \rightarrow 0 \text{ at } r} \quad \cdot \vdash C \leq C' : \text{Cap}}{\frac{\vdash M : \Psi \quad \Psi \vdash C \text{ sat}}{\Psi; \cdot; \cdot; C \vdash v_0(v_1, \dots, v_n)}}{\vdash (M, v_0(v_1, \dots, v_n))}$$

By Subtyping Regularity and Lemma 12, r is ν . By Capability Satisfiability Preservation, $\Psi \vdash C'' \oplus \{\tau\}$ sat, and by CECF, $\nu \in \text{Dom}(\Psi)$. Thus, by Canonical Forms, $v_0 = \nu.\ell[c_1, \dots, c_m]$, $M(\nu.\ell) = \text{fix } f[b_1, \dots, b_m](C, x_1:\tau_1, \dots, x_n:\tau_n).e$ and for $1 \leq i \leq n$, $b_i = \alpha_i:\kappa_i$ or $b_i = \alpha_i \leq C_i$. If we let S_1 be $[c_1, \dots, c_m/\alpha_1, \dots, \alpha_m]$ and S_2 be $[\nu.\ell, v_1, \dots, v_n/f, x_1, \dots, x_n]$ then

$$(M, v_0(v_1, \dots, v_n)) \mapsto (M, S_2(S_1(e))).$$

- halt

$$\frac{\frac{\Psi; \cdot; \cdot \vdash v : \text{int} \quad \cdot \vdash C = \emptyset : \text{Cap}}{\Psi; \cdot; \cdot; C \vdash \text{halt } v}}{\frac{\vdash M : \Psi \quad \Psi \vdash C \text{ sat}}{\vdash (M, \text{halt } v)}}$$

Part 2 holds by inspection of the typing judgement.

□

Definition 37. An abstract machine state (M, e) is stuck if e is not **halt** v and there does not exist (M', e') such that $(M, e) \mapsto (M', e')$.

THEOREM (TYPE SOUNDNESS). If $\vdash (M, e)$ and $(M, e) \mapsto^* (M', e')$ then (M', e') is not stuck.

PROOF. By induction on the number of steps taken in the operational semantics and Preservation, if $\vdash (M, e)$ and $(M, e) \mapsto^* (M', e')$ then $\vdash (M', e')$. By Progress, no well-typed state (M', e') is stuck: either e' is `halt v` or $(M', e') \mapsto (M'', e'')$. \square

THEOREM (COMPLETE COLLECTION). *If $\vdash (M, e)$ then either $(M, e) \uparrow$ or $(M, e) \mapsto^* (M', \text{halt } v)$ and $M' = \{\}$.*

PROOF. Assume $\vdash (M, e)$ and $(M, e) \mapsto^* (M', e')$, and assume there is no (M'', e'') such that $(M', e') \mapsto (M'', e'')$. By Preservation and Progress, $e' = \text{halt } v$ and

$$\frac{\vdash M' : \Psi \quad \Psi \vdash C \text{ sat} \quad \frac{\Psi; \cdot; \vdash v : \text{int} \quad \cdot \vdash C = \emptyset : \text{Cap}}{\Psi; \cdot; \cdot; C \vdash \text{halt } v}}{\vdash (M', \text{halt } v)}$$

By CECP and the sat judgment, $\nu \in \text{Dom}(\Psi)$ if and only if $\nu \in \emptyset$. Consequently, $\Psi = \{\}$. By inspection of the judgment for memory types, $M' = \{\}$. \square

B. REGION TRANSLATION TYPE PRESERVATION

In this section, we prove our translation of the region calculus into the CL is type preserving. In other words, given a well-formed source language term, the result of the translation is a well-formed CL term. Section 3 describes the syntax and static semantics of the term constructs.

B.1 Notation

We use the notation $\Delta \vdash_R \Gamma$ to indicate that all the types in Γ are well-formed under type context Δ . Formally:

$$\frac{}{\Delta \vdash_R \cdot} \quad \frac{\Delta \vdash_R \Gamma \quad \Delta \vdash_R \tau : \text{Type}}{\Delta \vdash_R \Gamma, x:\tau} \quad (x \notin \text{Dom}(\Gamma))$$

B.2 Overview

The theorem uses several lemmas from the proof of soundness including lemmas for manipulating capabilities (Lemma 17) and asserting well-formedness (Lemma 13). We also require a number of supplementary lemmas:

- Lemmas 40 and 41 describe additional well-formedness constraints on the types and effects that appear in region calculus judgments.
- Lemmas 42, 43, and 44 state that well-formedness of constructors, constructor equality, and substitution is preserved across the translation.
- Lemma 47 describes the way the subset relation is preserved during the translation of effects (sets) to capabilities (not sets). This lemma makes use of Lemmas 45 and 46 in its proof.
- Lemma 48 is a miscellaneous lemma required in the proof of the letregion construct.
- Lemma 49 gives the conditions under which the static application of a continuation closure to its arguments (*i.e.* $\mathcal{A}(k, v)$) is well-formed.

LEMMA 40.

- (1) If $\Delta \vdash_R c_1 = c_2 : \kappa$ then $\Delta \vdash_R c_1 : \kappa$ and $\Delta \vdash_R c_2 : \kappa$.
 (2) If $\Delta \vdash_R \psi_1 \subseteq \psi_2$ then $\Delta \vdash_R \psi_1 : \text{Eff}$ and $\Delta \vdash_R \psi_2 : \text{Eff}$

PROOF. By induction on the derivations. \square

LEMMA 41. If $\Delta \vdash_R \Gamma$ and $\Delta; \Gamma \vdash_R e : \tau, \psi$ then $\Delta \vdash_R \tau : \text{Type}$ and $\Delta \vdash_R \psi : \text{Eff}$

PROOF. By induction on the derivation $\Delta; \Gamma \vdash_R e : \tau, \psi$. \square

LEMMA (WELL-FORMEDNESS PRESERVATION).

- (1) If $\Delta \vdash_R \Delta'$ then $\mathcal{K}[\Delta] \vdash \mathcal{K}[\Delta']$
 (2) If $\Delta \vdash_R c : \kappa$ then $\mathcal{K}[\Delta] \vdash \mathcal{T}[c] : \mathcal{K}[\kappa]$

PROOF. By induction on the derivations. \square

LEMMA (EQUALITY PRESERVATION).

- (1) If $\Delta \vdash_R \psi = \psi' : \text{Eff}$ then $\mathcal{K}[\Delta] \vdash \overline{\mathcal{T}[\psi]} = \overline{\mathcal{T}[\psi']} : \text{Cap}$
 (2) If $\Delta \vdash_R c = c' : \kappa$ and $\kappa \neq \text{Eff}$ then $\mathcal{K}[\Delta] \vdash \mathcal{T}[c] = \mathcal{T}[c'] : \text{Cap}$

PROOF. By induction on the equality derivations. \square

LEMMA (SUBSTITUTION PRESERVATION). If $\Delta, \alpha : \kappa \vdash_R \tau : \text{Type}$ and $\Delta \vdash_R c : \kappa$ then

$\mathcal{K}[\Delta] \vdash \mathcal{T}[\tau[c/\alpha]] = \mathcal{T}[\tau][\mathcal{T}[c]/\alpha] : \text{Type}$.

PROOF. By induction on the typing derivation. \square

LEMMA 45. If $\Delta \vdash \psi = \psi' : \text{Eff}$ then $E(\mathcal{T}[\psi]) = E(\mathcal{T}[\psi'])$.

PROOF. By induction on the derivation. \square

LEMMA 46. If $\Delta \vdash \psi \subseteq \psi'$ then $E(\mathcal{T}[\psi]) \subseteq E(\mathcal{T}[\psi'])$.

PROOF. By inspection of the sub-effecting rule, $\Delta \vdash \psi \cup \psi'' = \psi' : \text{Eff}$ for some effect ψ'' .

By Lemma 45, $E(\mathcal{T}[\psi \cup \psi'']) = E(\mathcal{T}[\psi'])$.

By definition of E and the type translation, we have $E(\mathcal{T}[\psi]) \cup E(\mathcal{T}[\psi'']) = E(\mathcal{T}[\psi'])$.

Hence, $E(\mathcal{T}[\psi]) \subseteq E(\mathcal{T}[\psi'])$. \square

LEMMA 47. If $\Delta \vdash \psi \subseteq \psi'$ and $\Delta \vdash \overline{C} = \overline{C \oplus \mathcal{T}[\psi']} : \text{Cap}$ then $\Delta \vdash \overline{C} = \overline{C \oplus \mathcal{T}[\psi]} : \text{Cap}$.

PROOF. 1. By $\Delta \vdash \overline{C} = \overline{C \oplus \mathcal{T}[\psi']} : \text{Cap}$ and Equality (5), $E(C) = E(C \oplus \mathcal{T}[\psi'])$.

2. Thus, by definition of E , $E(\mathcal{T}[\psi']) \subseteq E(C)$.

3. By $\Delta \vdash \psi \subseteq \psi'$ and Lemma 46, $E(\mathcal{T}[\psi]) \subseteq E(\mathcal{T}[\psi'])$.

4. By 2 and 3, $E(\mathcal{T}[\psi]) \subseteq E(C)$.

5. Hence, by definition of E , $E(C) = E(C \oplus \mathcal{T}[\psi])$.

6. By 5 and Equality (6), $\Delta \vdash \overline{C} = \overline{C \oplus \mathcal{T}[\psi]} : \text{Cap}$. \square

LEMMA 48. If $\Delta \vdash_R \psi$ then $\mathcal{K}[\Delta] \vdash \overline{\mathcal{T}[\psi \setminus \{\rho\}] \oplus \{\rho^1\}} = \overline{\mathcal{T}[\psi] \oplus \{\rho^1\}} : \text{Cap}$.

PROOF. The proof is by induction on the structure of ψ .

— $\psi = \emptyset$. By definition, $\mathcal{T}[\emptyset \setminus \{\rho\}] = \emptyset$. The result follow immediately from the reflexivity of equality.

— $\psi = \epsilon$ or $\psi = \{\rho'\}$ and $\rho \neq \rho'$. Similar.

— $\psi = \{\rho\}$. The following reasoning provides the result.

$$\begin{aligned} \mathcal{K}[\Delta] \vdash \overline{\mathcal{T}[\{\rho\} \setminus \{\rho\}] \oplus \{\rho^1\}} &= \overline{\emptyset \oplus \{\rho^1\}} && \text{(By definition)} \\ &= \overline{\{\rho^1\}} && \text{(By rule eq-}\emptyset\text{)} \\ &= \{\rho^1\} \oplus \{\rho^1\} && \text{(By rule eq-dup)} \end{aligned}$$

— $\psi = \psi_1 \cup \psi_2$. By induction, we know that (1) $\mathcal{K}[\Delta] \vdash \overline{\mathcal{T}[\psi_i \setminus \{\rho\}] \oplus \{\rho^1\}} = \overline{\mathcal{T}[\psi_i] \oplus \{\rho^1\}}$ for $i = 1, 2$. Now, the following reasoning provides the result.

$$\begin{aligned} \mathcal{K}[\Delta] \vdash \overline{\mathcal{T}[(\psi_1 \cup \psi_2) \setminus \{\rho\}] \oplus \{\rho^1\}} & \\ &= \overline{\mathcal{T}[\psi_1 \setminus \{\rho\}] \oplus \mathcal{T}[(\psi_2 \setminus \{\rho\})] \oplus \{\rho^1\}} && \text{(By def.)} \\ &= \overline{\mathcal{T}[\psi_1 \setminus \{\rho\}] \oplus \mathcal{T}[(\psi_2 \setminus \{\rho\})] \oplus \{\rho^1\} \oplus \{\rho^1\}} && \text{(By eq-dup)} \\ &= \overline{\mathcal{T}[\psi_1 \setminus \{\rho\}] \oplus \{\rho^1\} \oplus \mathcal{T}[(\psi_2 \setminus \{\rho\})] \oplus \{\rho^1\}} && \text{(By eq-comm)} \\ &= \overline{\mathcal{T}[\psi_1] \oplus \{\rho^1\} \oplus \mathcal{T}[\psi_2] \oplus \{\rho^1\}} && \text{(By 1)} \\ &= \overline{\mathcal{T}[\psi_1] \oplus \mathcal{T}[\psi_2] \oplus \{\rho^1\} \oplus \{\rho^1\}} && \text{(By eq-comm)} \\ &= \overline{\mathcal{T}[\psi_1] \oplus \mathcal{T}[\psi_2] \oplus \{\rho^1\}} && \text{(By eq-dup)} \end{aligned}$$

□

LEMMA 49. Let $k = \langle x^k; e^k \rangle$. If $\Psi; \Delta; \Gamma, x^k; \tau; C \vdash e^k$ and $\Psi; \Delta; \Gamma \vdash v : \tau$ then $\Psi; \Delta; \Gamma; C \vdash \mathcal{A}(k, v)$.

PROOF. The term $\mathcal{A}(k, v)$ is defined to be $\mathbf{let} \ x^k = v \ \mathbf{in} \ e^k$. The following derivation is a proof the lemma.

$$\frac{\Psi; \Delta; \Gamma \vdash v : \tau \quad \Psi; \Delta; \Gamma, x^k; \tau; C \vdash e^k}{\Psi; \Delta; \Gamma; C \vdash \mathbf{let} \ x^k = v \ \mathbf{in} \ e^k}$$

□

LEMMA 50. If $\{\}; \Delta; \Gamma, x; \tau; C \vdash e$ and $\Delta \vdash \tau = \tau' : \mathbf{Type}$ then $\{\}; \Delta; \Gamma, x; \tau'; C \vdash e$.

PROOF. The proof is by induction on the typing derivation for expressions. □

THEOREM (CPS TYPE PRESERVATION). If $\cdot; \cdot \vdash e : \mathbf{int}, \emptyset$ then $\{\}; \cdot; \cdot; \emptyset \vdash \mathcal{C}_{\cdot; \cdot; \emptyset}(e) \langle x; \mathbf{halt} \ x \rangle$ where x is fresh and Θ is $\langle \cdot; \cdot; \emptyset; \emptyset \rangle$.

PROOF. The proof is by induction on the typing derivation of the expression using the following inductive hypothesis.

Given $\Delta, \Gamma, \Theta, e$, and k where $\Theta = \langle \Delta^\Theta; \Gamma^\Theta; C^\Theta; B^\Theta \rangle$ and $k = \langle x^k; e^k \rangle$. If

- A. $\Delta; \Gamma \vdash_R e : \tau, \psi$
- B. $\{\}; \mathcal{K}[\Delta], \Delta^\Theta; \mathcal{S}[\Gamma], \Gamma^\Theta, x^k; \mathcal{T}[\tau]; C^\Theta \vdash e^k$
- C. $\mathcal{K}[\Delta], \Delta^\Theta \vdash C^\Theta \leq \overline{B^\Theta}$
- D. $\mathcal{K}[\Delta], \Delta^\Theta \vdash \overline{B^\Theta} = \overline{B^\Theta \oplus \mathcal{T}[\psi]}$

Then:

E. $\{\}; \mathcal{K}[\Delta], \Delta^\Theta; \mathcal{S}[\Gamma], \Gamma^\Theta; C^\Theta \vdash \mathcal{C}_{\Delta; \Gamma; \Theta}(e)k$

We use several abbreviations and conventions in the proof. First, we will often use the meta-variable Φ to range over type-checking contexts of the form $\Psi; \Delta; \Gamma$ or $\Psi; \Delta; \Gamma; C$.

We also abbreviate derivations involving `let`. A derivation

$$\frac{\mathcal{D}_1 \quad \dots \quad \mathcal{D}_n}{\Psi; \Delta; \Gamma; C \vdash d \Rightarrow \Delta'; \Gamma'; C'} \quad \Psi; \Delta'; \Gamma'; C' \vdash e}{\Psi; \Delta; \Gamma; C \vdash \text{let } d \text{ in } e}$$

is abbreviated by

$$\frac{\mathcal{D}_1 \quad \dots \quad \mathcal{D}_n \quad \Psi; \Delta'; \Gamma'; C' \vdash e}{\Psi; \Delta; \Gamma; C \vdash \text{let } d \text{ in } e}.$$

Many of the typing rules contain the side condition that a CL variable ρ , ϵ , or x not be contained in the context Δ or Γ for that judgment. We have assumed that all the variables in the translated term have been generated fresh so that this will be the case. For the sake of brevity, we do not mention this side condition each time it occurs in the proof. Many of the rules also contain well-formedness constraints on types or capabilities (ie: $\Delta \vdash \tau$ or $\Delta \vdash C$). These well-formedness constraints always follow directly from the source typing judgement and Lemma 42. However, in order to concentrate on the more important aspects of the proof, we do not mention these conditions each time they appear in a derivation.

In the following, we prove the result for the more difficult cases: `letrec`, type application, value application, `letregion`, and equality. The other cases follow a similar, but simpler pattern.

—The case for `letrec`. The translation is

$$\begin{aligned} \mathcal{C}_{\Delta; \Gamma; \Theta}(\text{letrec } f[\Delta'](x) : \sigma \text{ at } e_1 = e_2 \text{ in } e_3)k = & \\ \mathcal{C}_{\Delta; \Gamma; \Theta}(e_1)\langle x_1; & \\ \text{let } f = (\text{fix } f[\mathcal{K}[\Delta'], \Delta''](\epsilon', x: \mathcal{T}[\tau_1], x_{cont}: \tau_{cont}). & \\ \mathcal{C}_{\Delta \Delta'; \Gamma\{f: \sigma, x: \tau_1\}; \Theta'}(e_2)\langle x_2; x_{cont}(x_2) \rangle) \text{ at } x_1 & \\ \text{in } \mathcal{C}_{\Delta; \Gamma\{f: \sigma\}; \Theta}(e_3)k \end{aligned}$$

where

$$\begin{aligned} \sigma &= \forall[\Delta']. \tau_1 \xrightarrow{\psi_f} \tau_2 \text{ at } r \\ \tau_{cont} &= \forall[.](\epsilon', \mathcal{T}[\tau_2]) \rightarrow 0 \text{ at } \rho \\ B^{\Theta'} &= \epsilon \oplus \mathcal{T}[\psi_f] \oplus \{\rho^1\} \\ \Delta'' &= \rho: \text{Rgn}, \epsilon: \text{Cap}, \epsilon' \leq \overline{B^{\Theta'}} \\ \Delta^{\Theta'} &= \Delta^\Theta, \Delta'' \\ \Gamma^{\Theta'} &= \Gamma^\Theta, x_1: r \text{ handle}, x_{cont}: \tau_{cont} \\ C^{\Theta'} &= \epsilon' \\ \Theta' &= \langle \Delta^{\Theta'}; \Gamma^{\Theta'}; C^{\Theta'}; B^{\Theta'} \rangle \end{aligned}$$

The source typing derivation A:

$$\frac{\begin{array}{l} (A_1) \Delta; \Gamma \vdash_R e_1 : r \text{ handle}, \psi_1 \\ (A_2) \Delta \Delta'; \Gamma \{f:\sigma, x:\tau_1\} \vdash_R e_2 : \tau_2, \psi_f \\ (A_3) \Delta; \Gamma \{f:\sigma\} \vdash_R e_3 : \tau_3, \psi_3 \end{array}}{\Delta; \Gamma \vdash_R \text{letrec } f[\Delta'](x) : \sigma \text{ at } e_1 = e_2 \text{ in } e_3 : \tau_3, \psi_1 \cup \psi_3 \cup \{r\}} \quad (x, f \notin \text{Dom}(\Gamma))$$

In order to make the derivations for this case more manageable, we will use the following abbreviations.

$$\begin{aligned} \Phi_0 &= \{ \}; \mathcal{K}[\Delta], \Delta^\Theta; \mathcal{S}[\Gamma], \Gamma^\Theta, x_1:r \text{ handle} \\ \Phi_1 &= \{ \}; \mathcal{K}[\Delta], \Delta^\Theta; \mathcal{S}[\Gamma], \Gamma^\Theta, x_1:r \text{ handle}; C^\Theta \\ \Phi_2 &= \{ \}; \mathcal{K}[\Delta, \Delta'], \Delta^{\Theta'}; \mathcal{S}[\Gamma, f:\sigma, x:\tau_1], \Gamma^{\Theta'}, x_1:\mathcal{T}[r \text{ handle}], x_2:\mathcal{T}[\tau_1]; C^{\Theta'} \\ \Phi_3 &= \{ \}; \mathcal{K}[\Delta, \Delta'], \Delta^{\Theta'}; \mathcal{S}[\Gamma, f:\sigma, x:\tau_1], \Gamma^{\Theta'}, x_1:\mathcal{T}[r \text{ handle}]; C^{\Theta'} \\ \Phi_4 &= \{ \}; \mathcal{K}[\Delta], \Delta^\Theta; \mathcal{S}[\Gamma, f:\sigma], \Gamma^\Theta; C^\Theta \\ \Phi_5 &= \{ \}; \mathcal{K}[\Delta], \Delta^\Theta; \mathcal{S}[\Gamma, f:\sigma], \Gamma^\Theta, x_1:\mathcal{T}[r \text{ handle}]; C^\Theta \end{aligned}$$

We begin by showing the continuation used in the translation of e_2 is well-formed under the appropriate context (call this fact B₂).

$$\frac{\frac{\text{(By v-var)}}{\Phi_2 \vdash x_{cont} : \mathcal{T}[\tau_{cont}]} \quad \frac{\text{(By v-var)}}{\Phi_2 \vdash x_2 : \mathcal{T}[\tau_1]} \quad \frac{\text{(By eq-reflex)}}{\mathcal{K}[\Delta, \Delta'], \Delta^{\Theta'} \vdash \epsilon' \leq \epsilon'} \quad \mathcal{D}_{\epsilon'}}{\Phi_2 \vdash x_{cont}(x_2)}$$

The derivation $\mathcal{D}_{\epsilon'}$ follows (recall $\Delta^{\Theta'} = \Delta^\Theta, \rho:\text{Rgn}, \epsilon:\text{Cap}, \epsilon' \leq \overline{\epsilon \oplus \mathcal{T}[\psi] \oplus \{\rho^1\}}$).

$$\begin{aligned} \mathcal{K}[\Delta, \Delta'], \Delta^{\Theta'} \vdash \epsilon' \leq \overline{\epsilon \oplus \mathcal{T}[\psi] \oplus \{\rho^1\}} & \quad \text{(By sub-var)} \\ &= \overline{\epsilon \oplus \mathcal{T}[\psi] \oplus \{\rho^+\}} \quad \text{(By eq-distrib)} \end{aligned}$$

We have now shown B₂. By use of rule sub-var, we can conclude that (C₂)

$$\mathcal{K}[\Delta, \Delta'], \Delta^{\Theta'} \vdash C^{\Theta'} \leq \overline{B^{\Theta'}}$$

We also have (D₂)

$$\begin{aligned} \mathcal{K}[\Delta, \Delta'], \Delta^{\Theta'} \vdash \overline{B^{\Theta'}} &= \overline{\epsilon \oplus \mathcal{T}[\psi_f] \oplus \{\rho^1\}} \quad \text{(expand abbreviation)} \\ &= \overline{\epsilon \oplus \{\rho^1\} \oplus \mathcal{T}[\psi_f]} \quad \text{(By eq-comm)} \\ &= \overline{\epsilon \oplus \{\rho^1\} \oplus \mathcal{T}[\psi_f] \oplus \mathcal{T}[\psi_f]} \quad \text{(By eq-dup)}. \end{aligned}$$

Using A₂, B₂, C₂, and D₂, we can apply the induction hypothesis and obtain E₂:

$$\Phi_3 \vdash \mathcal{C}_{\Delta \Delta'; \Gamma \{f:\sigma, x:\tau_1\}; \Theta'}(e_2) \langle x_2; x_{cont}(x_2) \rangle$$

Now, from A₃, B, C, and D, we can apply the induction hypothesis and obtain E₃:

$$\Phi_4 \vdash \mathcal{C}_{\Delta; \Gamma \{f:\sigma\}; \Theta, x_1:\mathcal{T}[r \text{ handle}]}(e_3)k$$

Using E₂ and E₃, we can build the typing derivation for the code in the continuation for translating e_1 (call this fact B₁):

$$\frac{\frac{\text{E}_2}{\Phi_0 \vdash \text{fix } f\dots. \mathcal{C}_{\Delta \Delta'; \Gamma \{f:\sigma, x:\tau_1\}; \Theta'}(e_2) \langle x_2; x_{cont}(x_2) \rangle} \text{at } r \quad \mathcal{D}_r \quad \text{E}_3}{\Phi_1 \vdash \text{let } f = (\dots) \text{ at } x_1 \text{ in } \mathcal{C}_{\Delta; \Gamma \{f:\sigma\}; \Theta, x_1:r \text{ handle}}(e_3)k}$$

where the derivation \mathcal{D}_{x_1} is

$$\frac{(\text{By v-var})}{\Phi_1 \vdash x_1 : r \text{ handle.}}$$

and the derivation \mathcal{D}_r is

$$\begin{aligned} \mathcal{K}[\Delta]\Delta^\Theta \vdash C^\Theta &\leq \overline{B^\Theta} && (\text{By assumption C}) \\ &= \overline{B^\Theta \oplus \{r^1\}} && (\text{By assumption D, Lemma 47}) \\ &= \overline{B^\Theta \oplus \{r^+\}} && (\text{By eq-distrib.}) \end{aligned}$$

Finally, using A₁, B₁, C, and D, we can apply the induction hypothesis to the translation of e_1 , giving us the final result:

$$\{ \}; \mathcal{K}[\Delta], \Delta^\Theta; \mathcal{S}[\Gamma], \Gamma^\Theta; C^\Theta \vdash \mathcal{C}_{\Delta; \Gamma; \Theta}(\text{letrec } f[\Delta'](x) : \sigma \text{ at } e_1 = e_2 \text{ in } e_3)k$$

—The case for type application. The translation is

$$\mathcal{C}_{\Delta; \Gamma; \Theta}(f[c_1, \dots, c_n])k = \mathcal{A}(k, f[\mathcal{T}[c_1], \dots, \mathcal{T}[c_n]]).$$

The source typing derivation A, where σ is $\forall[\alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n]. \tau_1 \xrightarrow{\psi} \tau_2 \text{ at } r :$

$$\frac{\Delta \vdash_R \sigma \quad \Delta \vdash_R c_i : \kappa_i}{\Delta; \Gamma \vdash_R f[c_1, \dots, c_n] : (\tau_1 \xrightarrow{\psi} \tau_2)[c_1, \dots, c_n/\alpha_1, \dots, \alpha_n] \text{ at } r, \emptyset} \quad (\Gamma(f) = \sigma)$$

First, we must show that the value $f[\mathcal{T}[c_1], \dots, \mathcal{T}[c_n]]$ is well-formed under the context

$$\Phi_1 = \{ \}; \mathcal{K}[\Delta], \Delta^\Theta; \mathcal{S}[\Gamma], \Gamma^\Theta$$

where $\Theta = \langle \Delta^\Theta; \Gamma^\Theta; C^\Theta; B^\Theta \rangle$

First, we can conclude that

$$\frac{(\text{By A and rule v-var})}{\Phi_1 \vdash f : \mathcal{T}[\forall[\alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n]. \tau_1 \xrightarrow{\psi} \tau_2 \text{ at } r]}$$

Now, by induction on the number of constructors c_i applied to f , and use of the rule v-type, we can conclude that

$$\Phi_1 \vdash f[\mathcal{T}[c_1], \dots, \mathcal{T}[c_n]] : \mathcal{T}[\tau_1 \xrightarrow{\psi} \tau_2 \text{ at } r][\mathcal{T}[c_1], \dots, \mathcal{T}[c_n]/\alpha_1, \dots, \alpha_n].$$

By the Substitution Lemma, Lemma 44, we have

$$\Phi_1 \vdash f[\mathcal{T}[c_1], \dots, \mathcal{T}[c_n]] : \mathcal{T}[(\tau_1 \xrightarrow{\psi} \tau_2 \text{ at } r)[c_1, \dots, c_n/\alpha_1, \dots, \alpha_n]].$$

Using this fact, assumption B, and Lemma 49, we obtain our final result:

$$\Phi_1; C^\Theta \vdash \mathcal{A}(k, f[\mathcal{T}[c_1], \dots, \mathcal{T}[c_n]])$$

—The case for application. The translation is

$$\begin{aligned}
\mathcal{C}_{\Delta; \Gamma; \Theta}(e_1^{\tau_f} e_2)k = & \\
& \mathcal{C}_{\Delta; \Gamma; \Theta}(e_1)\langle x_1; \\
& \mathcal{C}_{\Delta; \Gamma; \Theta, x_1: \mathcal{T}[\tau_f]}(e_2)\langle x_2; \\
& \text{let newrgn } \rho, x_\rho \text{ in} \\
& \text{let } f_{cont} = (\text{fix } f_{cont}[])(C^\Theta \oplus \{\rho^1\}, x: \mathcal{T}[\tau_2]). \\
& \quad \text{let freergn } x_\rho \text{ in } \mathcal{A}(k, x)) \text{ at } x_\rho \text{ in} \\
& x_1[\rho, B^\Theta, C^\Theta \oplus \{\rho^1\}](x_2, f_{cont})\rangle
\end{aligned}$$

where

$$\tau_f = \tau_1 \xrightarrow{\psi_f} \tau_2 \text{ at } r.$$

The source typing derivation for the term is (A):

$$\frac{(A_1) \Delta; \Gamma \vdash_R e_1 : \tau_1 \xrightarrow{\psi_3} \tau_2 \text{ at } r, \psi_1 \quad (A_2) \Delta; \Gamma \vdash_R e_2 : \tau_1, \psi_2}{\Delta; \Gamma \vdash_R e_1 e_2 : \tau_2, \psi_1 \cup \psi_2 \cup \psi_3 \cup \{r\}}$$

We begin showing that the result of the translation is type-correct by showing that the body of the innermost continuation (`let newrgn ρ, x_ρ in \dots`) is well-formed under the appropriate context. In order to make the derivation more manageable, we will use the following abbreviations:

$$\begin{aligned}
\Phi_1 &= \{ \}; \mathcal{K}[\Delta], \Delta^\Theta; \mathcal{S}[\Gamma], \Gamma^\Theta, x_1: \mathcal{T}[\tau_f], x_2: \mathcal{T}[\tau_2]; C^\Theta \\
\Phi_2 &= \{ \}; \mathcal{K}[\Delta], \Delta^\Theta, \rho: \text{Rgn}; \mathcal{S}[\Gamma], \Gamma^\Theta, x_1: \mathcal{T}[\tau_f], x_2: \mathcal{T}[\tau_2], x_\rho: \rho \text{ handle} \\
\Phi_3 &= \Phi_2; C^\Theta \oplus \{\rho^1\} \\
\Phi_4 &= \Phi_2, f_{cont}: \tau_{cont} \\
\Phi_5 &= \Phi_2, f_{cont}: \tau_{cont}, x: \mathcal{T}[\tau_2] \\
e'_1 &= \text{let } f_{cont} = h_{cont} \text{ at } x_\rho \text{ in } e'_2 \\
e'_2 &= x_1[\rho, B^\Theta, C^\Theta \oplus \{\rho^1\}](x_2, f_{cont}) \\
h_{cont} &= \text{fix } f_{cont}[](C^\Theta \oplus \{\rho^1\}, x: \mathcal{T}[\tau_2]). \text{let freergn } x_\rho \text{ in } \mathcal{A}(k, x) \\
\tau_{app} &= (C^\Theta \oplus \{\rho^1\}, \mathcal{T}[\tau_1], \tau_{cont}) \rightarrow 0 \text{ at } r \\
\tau_{cont} &= (C^\Theta \oplus \{\rho^1\}, \mathcal{T}[\tau_2]) \rightarrow 0 \text{ at } \rho
\end{aligned}$$

The derivation is as follows:

$$\frac{\frac{\frac{\mathcal{D}_{app}}{\Phi_4 \vdash x_1[\rho, B^\Theta, C^\Theta \oplus \{\rho^1\}] : \tau_{app} \quad \mathcal{D}_{x_2} \quad \mathcal{D}_{f_{cont}} \quad \mathcal{D}_r \quad \mathcal{D}_{C^\Theta}}{\Phi_3, f_{cont}: \tau_{cont} \vdash x_1[\rho, B^\Theta, C^\Theta \oplus \{\rho^1\}](x_2, f_{cont})} \quad \mathcal{D}_{x_\rho} \quad \mathcal{D}_h \quad \mathcal{D}_\rho}{\Phi_3 \vdash \text{let } f_{cont} = h_{cont} \text{ at } x_\rho \text{ in } e'_2}}{\Phi_1 \vdash \text{let newrgn } \rho, x_\rho \text{ in } e'_1}$$

The derivation \mathcal{D}_{app} can be proven as follows. First, by rule v-var, we can deduce

$$\begin{aligned}
\Phi_4 \vdash x_1 : \mathcal{T}[\tau_f] = & \\
& \forall[\rho': \text{Rgn}, \epsilon: \text{Cap}, \epsilon' \leq \overline{\epsilon \oplus \mathcal{T}[\psi_3] \oplus \{\rho^1\}}]. (\epsilon', \mathcal{T}[\tau_1], \tau_{cont}) \rightarrow 0 \text{ at } r.
\end{aligned}$$

From this judgment, two applications of the rule v-type give us

$$\Phi_4 \vdash x_1[\rho, B^\Theta] : \forall[\epsilon' \leq \overline{B^\Theta \oplus \mathcal{T}[\psi_3] \oplus \{\rho^1\}}]. (\epsilon', \mathcal{T}[\tau_1], \tau_{cont}) \rightarrow 0 \text{ at } r.$$

Finally, by rule v-sub, we can conclude

$$\Phi_4 \vdash x_1[\rho, B^\ominus, C^\ominus \oplus \{\rho^1\}] : \forall[].(C^\ominus \oplus \{\rho^1\}, \mathcal{T}[\tau_1], \tau_{cont}) \rightarrow 0 \text{ at } r$$

because the required subcapability relation holds:

$$\begin{aligned} \mathcal{K}[\Delta], \Delta^\ominus, \rho : \text{Rgn} \vdash C^\ominus \oplus \{\rho^1\} & \\ & \leq \overline{B^\ominus \oplus \{\rho^1\}} \quad (\text{By assumption C}) \\ & = \overline{B^\ominus \oplus \mathcal{T}[\psi_3] \oplus \{\rho^1\}} \quad (\text{By assumption D, Lemma 47}) \\ & \leq \overline{B^\ominus \oplus \mathcal{T}[\psi_3] \oplus \{\rho^1\}} \quad (\text{By rule sub-bar}) \end{aligned}$$

The derivation \mathcal{D}_{x_2} is

$$\frac{(\text{by v-var})}{\Phi_4 \vdash x_2 : \mathcal{T}[\tau_1]},$$

and $\mathcal{D}_{f_{cont}}$ is

$$\frac{(\text{by v-var})}{\Phi_4 \vdash f_{cont} : \tau_{cont}}.$$

Next, we consider the derivation \mathcal{D}_r . Here we must show that

$$\mathcal{K}[\Delta], \Delta^\ominus, \rho : \text{Rgn} \vdash C^\ominus \oplus \{\rho^1\} \leq C'' \oplus \{r^+\}$$

for some capability C'' . The reasoning is straightforward:

$$\begin{aligned} \mathcal{K}[\Delta], \Delta^\ominus, \rho : \text{Rgn} \vdash C^\ominus \oplus \{\rho^1\} & \\ & = \{\rho^1\} \oplus C^\ominus \quad (\text{By rule eq-comm}) \\ & \leq \{\rho^1\} \oplus \overline{B^\ominus} \quad (\text{By assumption C}) \\ & = \{\rho^1\} \oplus \overline{B^\ominus \oplus \mathcal{T}[\{r\}]} \quad (\text{By assumption D, Lemma 47}) \\ & = \{\rho^1\} \oplus \overline{B^\ominus \oplus \{r^1\}} \quad (\text{By def. of translation}) \\ & = \{\rho^1\} \oplus \overline{B^\ominus \oplus \{r^+\}} \quad (\text{By rule eq-distrib}) \end{aligned}$$

Next, we consider \mathcal{D}_{C^\ominus} . The judgment we must prove is

$$\mathcal{K}[\Delta], \Delta^\ominus \vdash C^\ominus \oplus \{\rho^1\} \leq C^\ominus \oplus \{\rho^1\}.$$

This follows by rules sub-eq and eq-reflex.

Next, we consider \mathcal{D}_{x_ρ} . The judgment we must prove is $\Phi_2 \vdash x_\rho : \rho$ handle. The judgment follows by rule v-var.

Next, we must prove \mathcal{D}_h :

$$\frac{\frac{(\text{by v-var})}{\Phi_5 \vdash x_\rho : \rho \text{ handle}} \quad \frac{(\text{by eq-reflex})}{\mathcal{K}[\Delta], \Delta^\ominus, \rho : \text{Rgn} \vdash C^\ominus \oplus \{\rho^1\} = C^\ominus \oplus \{\rho^1\} : \text{Cap}}{\Phi_5; C^\ominus \oplus \{\rho^1\} \vdash \text{let freern } x_\rho \text{ in } \mathcal{A}(k, x)}}{\Phi_2 \vdash h_{cont} \text{ at } \rho : \mathcal{T}[\tau_f]} \mathcal{D}_{\mathcal{A}(k, x)}$$

The judgment we are trying to prove in the derivation $\mathcal{D}_{\mathcal{A}(k, x)}$ is $\Phi_5; C^\ominus \vdash \mathcal{A}(k, x)$. Using rule v-var, we can conclude that $\Phi_5 \vdash x : \mathcal{T}[\tau_2]$. Assumption B tells us that $\Phi_5; C^\ominus \vdash e^k$. Hence, by Lemma 49, we have the result.

Finally, we show \mathcal{D}_ρ . We must prove $\mathcal{K}[\Delta], \Delta^\ominus, \rho : \text{Rgn} \vdash C^\ominus \oplus \{\rho^1\} \leq C'' \oplus \{\rho^+\}$ for some capability C'' . This fact follows using rule sub-dup (C'' is C^\ominus).

We have now satisfied all of the requirements necessary to show that the body of the innermost continuation is well-formed:

$$(B_2) \Phi_1 \vdash \mathbf{letregion} \rho, x_\rho \mathbf{in} e'_1$$

We have A_2 from the source typing derivation. If we let C_2 be C and D_2 be D , then we fulfill all of the requirements for the induction hypothesis, and we may conclude E_2 :

$$\{ \}; \mathcal{K}[\Delta], \Delta^\ominus; \mathcal{S}[\Gamma], \Gamma^\ominus, x_1: \mathcal{T}[\tau_f]; C^\ominus \vdash \\ \mathcal{C}_{\Delta; \Gamma; \Theta, x_1: \mathcal{T}[\tau_f]}(e_2) \langle x_2; \mathbf{letregion} \rho, x_\rho \mathbf{in} e'_1 \rangle$$

Now, using this fact, A_1 from the source typing derivation, C , and D , we can apply the induction hypothesis and conclude E_1 :

$$\{ \}; \mathcal{K}[\Delta], \Delta^\ominus; \mathcal{S}[\Gamma], \Gamma^\ominus; C^\ominus \vdash \mathcal{C}_{\Delta; \Gamma; \Theta}(e_1) \langle x_1; \mathcal{C}_{\Delta; \Gamma; \Theta, x_1: \mathcal{T}[\tau_f]}(e_2) \langle \dots \rangle \rangle$$

which is equivalent to the result E that we were trying to prove.

—The case for $\mathbf{letregion}$. The translation is

$$\mathcal{C}_{\Delta; \Gamma; \Theta}(\mathbf{letregion} \rho, x_\rho \mathbf{in} e)k = \\ \mathbf{letnewrgn} \rho, x_\rho \mathbf{in} \\ \mathcal{C}_{\Delta\{\rho: \mathbf{Rgn}\}; \Gamma\{x_\rho: \rho \mathbf{handle}\}; \Theta'}(e) \langle x'; \mathbf{freergn} x_\rho \mathbf{in} \mathcal{A}(k, x') \rangle$$

where

$$\Theta' = \langle \Delta^\ominus; \Gamma^\ominus; C^\ominus \oplus \{\rho^1\}; B^\ominus \oplus \{\rho^1\} \rangle.$$

The source typing derivation A

$$(A_1) \frac{\Delta\{\rho: \mathbf{Rgn}\}; \Gamma\{x_\rho: \rho \mathbf{handle}\} \vdash_R e : \tau, \psi \quad \left(\begin{array}{l} \rho \notin \mathit{ftv}(\tau) \cup \mathit{Dom}(\Delta) \\ x_\rho \notin \mathit{Dom}(\Gamma) \end{array} \right)}{\Delta; \Gamma \vdash_R \mathbf{letregion} \rho, x_\rho \mathbf{in} e : \tau, \psi \setminus \{\rho\}}$$

First, we show that the continuation for the translation of e is well-formed under an appropriate context:

$$(B_1) \frac{\mathcal{D}_{\mathcal{A}(k, x')} \quad \mathcal{D}_{C^\ominus} \quad \mathcal{D}_{x_\rho}}{\mathcal{K}[\Delta, \rho: \mathbf{Rgn}], \Delta^\ominus; \mathcal{S}[\Gamma, x_\rho: \rho \mathbf{handle}], \Gamma^\ominus, x': \mathcal{T}[\tau]; C^\ominus \oplus \{\rho^1\} \vdash \\ \mathbf{let freergn} x_\rho \mathbf{in} \mathcal{A}(k, x')}$$

The judgment we must prove in $\mathcal{D}_{\mathcal{A}(k, x')}$ is

$$(J) \mathcal{K}[\Delta, \rho: \mathbf{Rgn}], \Delta^\ominus; \mathcal{S}[\Gamma, x_\rho: \rho \mathbf{handle}], \Gamma^\ominus, x': \mathcal{T}[\tau]; C^\ominus \vdash \mathcal{A}(k, x').$$

By rule v-var, we can conclude that

$$\mathcal{K}[\Delta, \rho: \mathbf{Rgn}], \Delta^\ominus; \mathcal{S}[\Gamma, x_\rho: \rho \mathbf{handle}], \Gamma^\ominus, x': \mathcal{T}[\tau] \vdash x' : \mathcal{T}[\tau].$$

Using this fact, assumption B and Lemma 49, we can conclude (J).

The judgment we must prove in \mathcal{D}_{C^\ominus} is

$$\mathcal{K}[\Delta, \rho: \mathbf{Rgn}], \Delta^\ominus \vdash C^\ominus \oplus \{\rho^1\} = C^\ominus \oplus \{\rho^1\}$$

which follows by the rule eq-reflex.

The derivation \mathcal{D}_{x_ρ} is

$$\frac{\text{(by v-var)}}{\mathcal{K}[\Delta, \rho:\text{Rgn}], \Delta^\ominus; \mathcal{S}[\Gamma, x_\rho:\rho \text{ handle}], \Gamma^\ominus, \Gamma^k, x':\mathcal{T}[\tau] \vdash x_\rho : \rho \text{ handle}} .$$

Now, we have fulfilled all the requirements necessary to show that the body of the innermost continuation is well-formed (call this fact B_1). We have A_1 from the typing derivation. In order to apply the induction hypothesis, we must show C_1 :

$$\begin{aligned} \mathcal{K}[\Delta, \rho:\text{Rgn}], \Delta^\ominus \vdash C^\ominus \oplus \{\rho^1\} &\leq \overline{B^\ominus \oplus \{\rho^1\}} \text{ (By assumption C)} \\ &\leq \overline{B^\ominus \oplus \{\rho^1\}} \text{ (By rule sub-bar)} \\ &= \overline{B^\ominus \oplus \{\rho^1\}} \text{ (By rule eq-distrib)} \end{aligned}$$

and D_1 :

$$\begin{aligned} \mathcal{K}[\Delta, \rho:\text{Rgn}], \Delta^\ominus \vdash \overline{B^\ominus \oplus \{\rho^1\}} &= \overline{B^\ominus \oplus \mathcal{T}[\psi \setminus \{\rho\}] \oplus \{\rho^1\}} \text{ (By assumption D)} \\ &= \overline{B^\ominus \oplus \mathcal{T}[\psi] \oplus \{\rho^1\}} \text{ (By Lemma 48)} \\ &= \overline{B^\ominus \oplus \{\rho^1\} \oplus \mathcal{T}[\psi]} \text{ (By rule eq-comm)} \end{aligned}$$

Together A_1 , B_1 , C_1 , and D_1 satisfy the preconditions for applying the induction hypothesis. The result is E_1 :

$$\mathcal{K}[\Delta, \rho:\text{Rgn}], \Delta^\ominus; \mathcal{S}[\Gamma, x_\rho:\rho \text{ handle}], \Gamma^\ominus; C^\ominus \oplus \{\rho^1\} \vdash \mathcal{C}_{\mathcal{K}[\Delta, \rho:\text{Rgn}], \Delta^\ominus; \mathcal{S}[\Gamma, x_\rho:\rho \text{ handle}], \Theta'(e)}(e)\langle \dots \rangle$$

Now, we can show the result of the translation type-checks:

$$\frac{E_1}{\mathcal{K}[\Delta], \Delta^\ominus; \mathcal{S}[\Gamma], \Gamma^\ominus; C^\ominus \vdash \text{let newrgn } \rho, x_\rho \text{ in } \mathcal{C}_{\mathcal{K}[\Delta, \rho:\text{Rgn}], \Delta^\ominus; \mathcal{S}[\Gamma, x_\rho:\rho \text{ handle}], \Theta'(e)}(e)\langle \dots \rangle}$$

—The equality rule. The source typing derivation A is

$$\frac{\Delta; \Gamma \vdash_R e : \tau, \psi \quad \Delta \vdash_R \tau = \tau' : \text{Type} \quad \Delta \vdash_R \psi \subseteq \psi'}{(A_1)\Delta; \Gamma \vdash_R e : \tau', \psi'}$$

and the continuation $k = \langle x; e^k \rangle$ is well-formed under the appropriate context (B):

$$\{ \}; \mathcal{K}[\Delta], \Delta^\ominus; \mathcal{S}[\Gamma], \Gamma^\ominus, x:\mathcal{T}[\tau]; C^\ominus \vdash e^k$$

By Lemma 43 and the equality judgment $\Delta \vdash \tau = \tau' : \text{Type}$ we can deduce that

$$\mathcal{K}[\Delta] \vdash \mathcal{T}[\tau] = \mathcal{T}[\tau'] : \text{Type}.$$

Therefore, by Lemma 50, we can deduce (B₁).

$$\{ \}; \mathcal{K}[\Delta], \Delta^\ominus; \mathcal{S}[\Gamma], \Gamma^\ominus, x:\mathcal{T}[\tau']; C^\ominus \vdash e^k$$

Now, recall Assumption D states that

$$\mathcal{K}[\Delta], \Delta \vdash \overline{B^\ominus} = \overline{B^\ominus \oplus \mathcal{T}[\psi']} : \text{Cap}.$$

Using this fact, and the source typing derivation (A), which states that $\Delta \vdash \psi \subseteq \psi'$, and Lemma 47, we can deduce that (D₁)

$$\mathcal{K}[\Delta], \Delta \vdash \overline{B^\Theta} = \overline{B^\Theta \oplus \mathcal{T}[\psi]} : \text{Cap.}$$

Using A₁, B₁, C, and D₁, we can apply the induction hypothesis and obtain

$$\{ \}; \mathcal{K}[\Delta], \Delta^\Theta; \mathcal{S}[\Gamma], \Gamma^\Theta; C^\Theta \vdash \mathcal{C}_{\Delta; \Gamma; \Theta}(e)k,$$

and we are done.

We have completed the proof that the induction hypothesis is preserved by the translation. In order to obtain the proof of the CPS translation theorem, we simply instantiate the induction hypothesis. We are given part A: $\cdot; \cdot \vdash e : \text{int}, \emptyset$.

The continuation k is $\langle x; \text{halt } x \rangle$. Using the halt rule, we have part B,

$$\frac{\{ \}; \cdot; x : \text{int} \vdash x : \text{int} \quad \cdot \vdash \emptyset = \emptyset : \text{Cap}}{\{ \}; \cdot; x : \text{int}; \emptyset \vdash \text{halt } x}$$

Part C is trivial: $\cdot \vdash \emptyset \leq \emptyset$.

Part D is also straightforward,

$$\begin{aligned} \cdot \vdash \overline{\emptyset} &\leq \overline{\emptyset} && \text{(By rule sub-eq)} \\ &= \overline{\emptyset \oplus \emptyset} && \text{(By rule eq-dup)} \\ &= \overline{\emptyset \oplus \mathcal{T}[\emptyset]} && \text{(By definition of } \mathcal{T}[\cdot] \text{)} \end{aligned}$$

Therefore, we can conclude E, $\{ \}; \cdot; \cdot; \emptyset \vdash \mathcal{C}_\Theta(e)k$ where Θ is the empty translation environment $\langle \cdot; \cdot; \emptyset; \emptyset \rangle$ and k is the trivial continuation $\langle x; \text{halt } x \rangle$.

□

ACKNOWLEDGMENTS

We would like to thank Lars Birkedal, Martin Elsmann, Neal Glew, Dan Grossman, Chris Hawblitzel, Fred Smith, Mads Tofte, Stephanie Weirich, and Steve Zdancewic for their comments and suggestions on earlier drafts of this article. We are also grateful both to the anonymous reviewers of this TOPLAS article and to the reviewers of its precursor, a paper that appeared in the twenty-sixth Symposium on Principles of Programming Languages.

REFERENCES

- ABRAMSKY, S. 1993. Computational interpretations of linear logic. *Theoretical Computer Science* 111, 3–57.
- AIKEN, A., FÄHNDRICH, M., AND LEVIEN, R. 1995. Better static memory management: Improving region-based analysis of higher-order languages. In *ACM Conference on Programming Language Design and Implementation*. ACM Press, La Jolla, California, 174–185.
- ALPERN, B. AND SCHNEIDER, F. 1987. Recognizing safety and liveness. *Distributed Computing* 2, 117–126.
- BAKER, H. G. 1978. List processing in real-time on a serial computer. *Communications of the ACM* 21, 4, 280–294.
- BERSHAD, B., SAVAGE, S., PARDYAK, P., SIRER, E., FIUCZYNSKI, M., BECKER, D., CHAMBERS, C., AND EGGERS, S. 1995. Extensibility, safety and performance in the SPIN operating system. In *Fifteenth ACM Symposium on Operating Systems Principles*. ACM Press, Copper Mountain, 267–284.

- BIRKEDAL, L., ROTHWELL, N., TOFTE, M., AND TURNER, D. N. 1993. The ML Kit (version 1). Tech. Rep. 93/14, Department of Computer Science, University of Copenhagen.
- BIRKEDAL, L., TOFTE, M., AND VEJLSTRUP, M. 1996. From region inference to von Neumann machines via region representation inference. In *Twenty-Third ACM Symposium on Principles of Programming Languages*. ACM Press, St. Petersburg, 171–183.
- BELLELOCH, G. AND GREINER, J. 1996. A provably time and space efficient implementation of NESL. In *ACM International Conference on Functional Programming*. ACM Press, Philadelphia, 213–225.
- CRARY, K., WALKER, D., AND MORRISETT, G. 1999. Typed memory management in a calculus of capabilities. In *Twenty-Sixth ACM Symposium on Principles of Programming Languages*. ACM Press, San Antonio, 262–275.
- CRARY, K., WEIRICH, S., AND MORRISETT, G. 1998. Intensional polymorphism in type-erasure semantics. In *ACM International Conference on Functional Programming*. ACM Press, Baltimore, 301–312.
- DANVY, O., DZAFIC, B., AND PFENNING, F. 1999. On proving syntactic properties of CPS programs. In *Third International Workshop on Higher-Order Operational Techniques in Semantics*, A. Gordon and A. Pitts, Eds. Electronic Notes in Computer Science, vol. 26. Elsevier, Paris, 19–31.
- DANVY, O. AND FILINSKI, A. 1992. Representing control: a study of the CPS transformation. *Mathematical Structures in Computer Science* 2, 4 (Dec.), 361–391.
- FILINSKI, A. 1996. Controlling effects. Ph.D. thesis, Carnegie Mellon University, School of Computer Science, Pittsburgh, Pennsylvania.
- FISCHER, M. J. 1972. Lambda calculus schemata. In *Proceedings of the ACM Conference on Proving Assertions about Programs*. 104–109.
- FLANAGAN, C. AND ABADI, M. 1999. Types for safe locking. In *Lecture Notes in Computer Science*, S. Swierstra, Ed. Vol. 1576. Springer-Verlag, Amsterdam, 91–108. Appeared in the Eighth European Symposium on Programming.
- GAY, D. AND AIKEN, A. 1998. Memory management with explicit regions. In *ACM Conference on Programming Language Design and Implementation*. ACM Press, Montreal, 313 – 323.
- GIFFORD, D. K. AND LUCASSEN, J. M. 1986. Integrating functional and imperative programming. In *ACM Conference on Lisp and Functional Programming*. ACM Press, Cambridge, Massachusetts, 28–38.
- GIRARD, J.-Y. 1987. Linear logic. *Theoretical Computer Science* 50, 1–102.
- HARPER, R. AND LILLIBRIDGE, M. 1993. Explicit polymorphism and CPS conversion. In *Twentieth ACM Symposium on Principles of Programming Languages*. ACM Press, Charleston, 206–219.
- HARPER, R. AND MORRISETT, G. 1995. Compiling polymorphism using intensional type analysis. In *Twenty-Second ACM Symposium on Principles of Programming Languages*. ACM Press, San Francisco, 130–141.
- HAWBLITZEL, C., CHANG, C.-C., CZAJKOWSKI, G., HU, D., AND VON EICKEN, T. 1998. Implementing multiple protection domains in Java. In *1998 USENIX Annual Technical Conference*. USENIX, New Orleans.
- HAWBLITZEL, C. AND VON EICKEN, T. 1999. Type system support for dynamic revocation. In *ACM SIGPLAN workshop on Compiler Support for System Software*. ACM Press, Atlanta.
- JOUVELOT, P. AND GIFFORD, D. K. 1991. Algebraic reconstruction of types and effects. In *Eighteenth ACM Symposium on Principles of Programming Languages*. ACM Press, Orlando, 303–310.
- KOZEN, D. 1998. Efficient code certification. Tech. Rep. TR98-1661, Cornell University. Jan.
- LAFONT, Y. 1988. The linear abstract machine. *Theoretical Computer Science* 59, 157–180.
- LAUNCHBURY, J. AND PEYTON JONES, S. L. 1995. State in Haskell. *LISP and Symbolic Computation* 8, 4 (Dec.), 293–341.
- LINDHOLM, T. AND YELLIN, F. 1996. *The Java Virtual Machine Specification*. Addison-Wesley, Menlo Park, California.
- LUCASSEN, J. M. 1987. Types and effects—towards the integration of functional and imperative programming. Ph.D. thesis, MIT Laboratory for Computer Science.

- MILNER, R., TOFTE, M., HARPER, R., AND MACQUEEN, D. 1997. *The Definition of Standard ML (Revised)*. MIT Press, Boston.
- MINAMIDE, Y. 1999. Space-profiling semantics of the call-by-value lambda calculus and the CPS transformation. In *Third International Workshop on Higher-Order Operational Techniques in Semantics*, A. D. Gordon and A. Pitts, Eds. Electronic Notes in Computer Science, vol. 26. Elsevier, Paris, 103–118.
- MOGGI, E. 1991. Notions of computation and monads. *Information and Computation* 93, 55–92.
- MORRISSETT, G., CRARY, K., GLEW, N., GROSSMAN, D., SAMUELS, R., SMITH, F., WALKER, D., WEIRICH, S., AND ZDANCEWIC, S. 2000. Typed Assembly Language for the Intel IA32 architecture. See <http://www.cs.cornell.edu/talc> for the latest implementation.
- MORRISSETT, G., FELLEISEN, M., AND HARPER, R. 1995. Abstract models of memory management. In *ACM Conference on Functional Programming and Computer Architecture*. ACM Press, La Jolla, 66–77.
- MORRISSETT, G. AND HARPER, R. 1997. Semantics of memory management for polymorphic languages. In *Higher-Order Operational Techniques in Semantics*, A. Gordon and A. Pitts, Eds. Publications of the Newton Institute. Cambridge University Press, Cambridge, UK.
- MORRISSETT, G., WALKER, D., CRARY, K., AND GLEW, N. 1998. From System F to Typed Assembly Language. In *Twenty-Fifth ACM Symposium on Principles of Programming Languages*. ACM Press, San Diego, 85–97.
- MORRISSETT, G., WALKER, D., CRARY, K., AND GLEW, N. 1999. From System F to Typed Assembly Language. *ACM Transactions on Programming Languages and Systems* 21, 3 (May), 528–569.
- NECULA, G. 1997. Proof-carrying code. In *Twenty-Fourth ACM Symposium on Principles of Programming Languages*. ACM Press, Paris, 106–119.
- NECULA, G. AND LEE, P. 1996. Safe kernel extensions without run-time checking. In *Proceedings of Operating System Design and Implementation*. USENIX assoc., Seattle, 229–243.
- NECULA, G. AND LEE, P. 1998. The design and implementation of a certifying compiler. In *ACM Conference on Programming Language Design and Implementation*. ACM Press, Montreal, 333–344.
- PEYTON JONES, S. L. AND WADLER, P. 1993. Imperative functional programming. In *Twentieth ACM Symposium on Principles of Programming Languages*. ACM Press, Charleston, South Carolina.
- PLOTKIN, G. D. 1975. Call-by-name, call-by-value, and the lambda calculus. *Theoretical Computer Science* 1, 125–159.
- REYNOLDS, J. C. 1972. Definitional interpreters for higher-order programming languages. In *Conference Record of the 25th National ACM Conference*. Boston, 717–740.
- REYNOLDS, J. C. 1978. Syntactic control of interference. In *Fifth ACM Symposium on Principles of Programming Languages*. ACM Press, Tucson, Arizona, 39–46.
- REYNOLDS, J. C. 1989. Syntactic control of interference, part 2. In *Automata, Languages and Programming: 16th International Colloquium*. Lecture Notes in Computer Science, vol. 372. Springer-Verlag, Stresa, Italy, 704–722.
- SABRY, A. AND FELLEISEN, M. 1993. Reasoning about programs in continuation-passing style. *LISP and Symbolic Computation* 6, 3/4, 289–360.
- SCHNEIDER, F. 2000. Enforceable security policies. *ACM Transactions on Information and System Security* 3, 1 (Feb.).
- SMITH, F., WALKER, D., AND MORRISSETT, G. 2000. Alias types. In *Lecture Notes in Computer Science*, G. Smolka, Ed. Vol. 1782. Springer-Verlag, Berlin, 366–381. Appeared in the *Ninth European Symposium on Programming*.
- TOFTE, M. AND BIRKEDAL, L. 1998. A region inference algorithm. *Transactions on Programming Languages and Systems* 20, 4 (Nov.), 734–767.
- TOFTE, M. AND TALPIN, J.-P. 1994. Implementation of the typed call-by-value λ -calculus using a stack of regions. In *Twenty-First ACM Symposium on Principles of Programming Languages*. ACM Press, Portland, Oregon, 188–201.
- TOFTE, M. AND TALPIN, J.-P. 1997. Region-based memory management. *Information and Computation* 132, 2, 109–176.
- ACM Transactions on Programming Languages and Systems, Vol. TBD, No. TDB, Month Year.

- WADLER, P. 1990. Linear types can change the world! In *Programming Concepts and Methods*, M. Broy and C. Jones, Eds. North Holland, Sea of Galilee, Israel. IFIP TC 2 Working Conference.
- WADLER, P. 1993. A taste of linear logic. In *Mathematical Foundations of Computer Science*. Lecture Notes in Computer Science, vol. 711. Springer-Verlag, Gdansk, Poland.
- WAHBE, R., LUCCO, S., ANDERSON, T., AND GRAHAM, S. 1993. Efficient software-based fault isolation. In *Fourteenth ACM Symposium on Operating Systems Principles*. ACM Press, Asheville, 203–216.
- WALKER, D. 2000. A type system for expressive security policies. In *Twenty-Seventh ACM Symposium on Principles of Programming Languages*. ACM Press, Boston, 254–267.
- WALKER, D. AND MORRISSETT, G. 2000. Alias types for recursive data structures. In *Workshop on Types in Compilation*. Montreal.
- WILSON, P. R. 1992. Uniprocessor garbage collection techniques. In *International Workshop on Memory Management*, Y. Bekkers and J. Cohen, Eds. Number 637 in Lecture Notes in Computer Science. Springer-Verlag, St. Malo, 1–42.
- WRIGHT, A. K. AND FELLEISEN, M. 1994. A syntactic approach to type soundness. *Information and Computation* 115, 1, 38–94.
- WULF, W. A., LEVIN, R., AND HARBISON, S. P. 1981. *Hydra/C.mmp: An Experimental Computer System*. McGraw-Hill, New York, NY.

Received February 2000; accepted May 2000