

Low-Level Linear Memory Management

Chris Hawblitzel Edward Wei Heng Huang Eric Krupski Lea Wittie

Department of Computer Science, Dartmouth College

ABSTRACT

Efficient low-level systems need more control over memory than safe high-level languages usually provide. As a result, run-time systems are typically written in unsafe languages such as C. This paper extends previous work on linear types, alias types, regions, and typed garbage collection to give type-safe code more control over memory. The approach is truly low-level: memory consists of a single linear array of words, with load and store operations but no built-in notion of an object. The paper constructs lists and arrays out of the basic linear memory primitives, and then introduces *type sequences* for building regions of nonlinear data. It then describes a Cheney queue typed garbage collector, implemented safely over regions.

1. INTRODUCTION

Modern computers rely on the correctness and security of low-level systems, such as garbage collectors, device drivers, and embedded system code. Often, a small amount of low-level code (say, in a firewall, a network interface device driver, or secure coprocessor[24]) is all that stands between a computer system and a hacker trying to break into the system. Even in the absence of malicious outsiders, buggy device drivers often cause annoying system crashes. Given their role as the foundation for higher-level services, one might expect low-level systems to benefit from the static and run-time checks provided by type-safe languages. However, systems programmers usually lean towards assembly language, C, or C++ rather than Java, ML, or Haskell, because they need efficient low-level control of memory. Java's array bounds checking, for example, prevents the buffer overflow attacks that so many C programs are susceptible to, but often imposes extra run-time overhead. Automatic garbage collection prevents dangling pointers, but *implementing* a garbage collector requires the ability to explicitly free heap objects, a privilege not usually granted to safe language programs.

This paper extends previous work on linear types, alias

types, regions, and typed garbage collection to give type-safe code more control over memory. It models memory as a linear array of words, indexed by integer memory addresses. Since linearity prevents aliasing, store operations can change the type of a memory word without the danger of a subsequent load operation reading the word with the wrong type. To support integer addressing, the type system includes arithmetic based on singleton types, in the style of Xi et al[32]. The paper then extends these basic tools with *coercion functions*, which are essentially proofs to modify the types of values, and *type sequences*, which introduce non-linearity into the system without imposing a specific memory management strategy for the nonlinear data. These are sufficient to implement region-based memory management[26][8] and garbage collection.

To demonstrate the practicality of these ideas, we have developed a safe C-like language called *Clay*, which implements the type system in this paper, along with some heuristic-based type inference and static detection of possible 32-bit integer overflow. The Clay type checker uses the Omega test software[20], modified to support arbitrary-precision arithmetic, to check arithmetic constraints. The type-checked code is compiled to C++. All the examples in this paper, including the garbage collectors and other run-time system code, are included with the Clay distribution¹.

Rather than focusing on the specifics of Clay, this paper presents the type system as part of an abstract machine called λ^{low} , based on the standard call-by-value higher-order polymorphic lambda calculus, F^ω . Although it may seem strange to describe low-level operations from within a high-level lambda calculus framework, the generality of the lambda calculus allows us to apply the ideas from λ^{low} to many settings. For example, a closure-converted subset of λ^{low} mimics C-like languages such as Clay, while a CPS- and closure-converted subset of λ^{low} mimics register transfer languages or assembly-languages. Because of this, we believe that λ^{low} 's type system can contribute to the development of systems for *proof-carrying code* (PCC)[18] and *typed assembly language* (TAL)[17]. Since these systems run untrusted code directly on a trusted computer, they are particularly sensitive to bugs or vulnerabilities in the run-time system, and considerable work has focused on moving memory management out of the "trusted computing base" (TCB) for PCC and TAL. The most ambitious effort to reduce the TCB is *foundational PCC*[2][10], which attempts to build an entire PCC/TAL system up from a small number of axioms. Because linear memory types are based on a simple,

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

Copyright 200X ACM X-XXXXX-XX-X/XX/XX ...\$5.00.

¹available at <http://www.cs.dartmouth.edu/~hawblitz/>

low-level model of memory, they may be suitable for foundational PCC systems.

2. BACKGROUND

This paper draws from previous research on typed memory management, including linear types, alias types, regions, and typed garbage collection. This section provides a brief background on these techniques.

The primary obstacle to type-safe memory management is dangling pointers (pointers to deallocated memory). In the presence of aliasing (multiple pointers to the same object), it is difficult to prevent dangling pointers. The simplest strategy for safe deallocation is to disallow aliasing completely: if all data structures are trees, rather than arbitrary graphs, then each object has only one pointer to it, and the type system can ensure safety by invalidating the one pointer to an object when the program deallocates the object. In the terminology of linear types [27], a pointer P to an object O is “consumed” when O dies. For example, if the variable x contains the pointer P , then the expression “`deallocate(x)`” removes x from the scope of the rest of the program, so that any subsequent reference to x fails to type-check.

While pure linear data structures have been implemented in LISP [4] and designed for typed assembly language [6], they tend to be clumsy in practice. Many useful data structures, such as circular and doubly-linked lists, aren’t linear. Even for linear data structures, useful operations on the data often violate the linearity. For example, a purely linear program cannot traverse a list while simultaneously holding a pointer to the head of the list, because this would result in two pointers to the list. As a result, programs often make extra copies of data structures, one copy for every pointer to the data structure.

Alias types [23][28] allow limited forms of aliasing rather than prohibiting aliasing completely, so that a program can express some nonlinear data structures. Alias types describe the current state of a set of objects in a *constraint*. For example, the constraint $\{\ell_1 \mapsto \langle \text{int}, \text{int} \rangle, \ell_2 \mapsto \langle \text{int}, \text{int}, \text{int} \rangle\}$ indicates that there are two distinct objects at memory locations ℓ_1 and ℓ_2 , one of which is a pair of integers and one of which is a triplet of integers. Allocation, deallocation, loads, and stores alter the state of the current constraint. For example, storing a floating point number into the second element of ℓ_2 produces a new constraint $\{\ell_1 \mapsto \langle \text{int}, \text{int} \rangle, \ell_2 \mapsto \langle \text{int}, \text{float}, \text{int} \rangle\}$. Constraints are treated linearly, so that the old constraint is consumed when the new constraint is produced, which prevents the program from trying to use the second field of ℓ_2 as an integer after it has been changed to a float. In contrast to pure linear types, the program may keep many pointers to ℓ_2 simultaneously—the type “pointer-to- ℓ_2 ” is nonlinear and may be freely copied. The linearity of the constraint ensures that any use of a “pointer-to- ℓ_2 ” value respects the current state of ℓ_2 .

While alias types are more flexible than pure linear data types, the linear constraint disallows many nonlinear data structures and nonlinear operations on data. In practice, the allowed data structures are “almost linear”—they look like linear data augmented with carefully specified extra pointers. For example, Walker et al [28] express a circular list as a singly-linked list plus one special pointer from the tail to head. However, the authors do not present standard operations on circular lists, such as traversing a list in a circle from the head to tail and back around to the head.

One can extend the power of linear and alias types by adding ever more sophisticated ways to control aliasing (see shape types, for example [9]), but most programming languages allow unrestricted pointer aliasing in the heap, and a compiler from a standard source programming language to typed assembly language or proof-carrying code must deal with this unrestricted aliasing. One strategy is to allow arbitrary nonlinear data inside a *region* [26][8], whose lifetime is controlled linearly (or almost linearly). A pointer to an object O inside a region R is assigned a type “ τ at R ”, where τ describes the data inside O . Loads and stores to a “ τ at R ” value are only type-correct if R is still alive. When R is destroyed, all the objects in it are deallocated together. Deallocation of individual objects inside a live region is usually disallowed, because arbitrary aliasing is possible inside the region, which means that individual object deallocation could leave dangling pointers.

Any program can be trivially rewritten to take advantage of regions: simply create one big region, allocate all objects in the region, and keep the region alive until the program finishes. This is usually not an efficient strategy, though, so either automated inference [26] or manual region management is necessary to break data up into smaller regions that the program can deallocate as soon as possible to minimize memory consumption. Wang and Appel [30] observed that a program can copy live data from one region into a second region and then deallocate the first region, effectively constructing a copying garbage collector written entirely with type-safe language features. Such a *typed* (or *type-preserving*) collector offers the generality of traditional garbage collection without requiring the collector to be trusted. Unfortunately, the details of the copying process are troublesome; the collector needs some way to traverse the data, and it needs to maintain forwarding pointers from the old region into the new region. Several solutions to these problems have been proposed [30][15][16], but each new proposal changes the typing rules for regions, often in *ad hoc* and complex ways. Furthermore, the resulting collectors are often unable to use low-level techniques common in collectors written in C.

The goal of this paper is to develop regions from the ground up, starting with a combination of linear and alias types, in order to provide a stable, flexible, safe, and low-level platform for constructing typed garbage collectors. The rest of paper is as follows:

- Section 3 introduces the core language for manipulating linear memory. This includes basic function, tuple, polymorphic, and recursive types, plus linear memory types and integer arithmetic types. These types are sufficient to implement basic recursive linear data structures.
- Section 4 extends the language with coercion functions, which add flexibility to the data types from section 3. The stack type from section 3 becomes an array type, for example.
- Section 5 describes how to combine nonlinear functions and linear memory types to implement primitive “regions” of nonlinear data. No language extensions are required, but the technique does not allow regions to grow dynamically – all the region’s data must be allocated when the region is allocated.

- Section 6 introduces a new language feature called *type sequences* so that the regions from section 5 can grow as new data is allocated. Because type sequences seem somewhat *ad hoc*, section 6 also describes how to encode type sequences using a more elegant language extension called *delayed types*.
- Section 7 shows how to build a simple Cheney queue copying garbage collector using the regions from section 6.

3. LINEAR MEMORY

This section describes the syntax and semantics of λ^{low} . To make the core language clearer, we delay a couple of features (coercion functions and type sequences) until later sections. Type checking rules for expressions appear throughout the paper; the complete syntax and semantics of λ^{low} is found in [12]. For the portion of the language described in this subsection, two type environments are needed to type-check an expression: Γ maps variables to types, and Ψ maps integer memory word addresses to types. For convenience, we often write a combined context C , defined in this subsection to be $C \triangleq \Psi; \Gamma$. The notation $C = C_1, C_2$ indicates that C , C_1 , and C_2 share the same nonlinear assumptions, but that each of C 's linear assumptions appears in either C_1 or C_2 , not both. The notation \dot{C} denotes a context with no linear assumptions.

The state of a running program consists of memory M , which maps integer memory word addresses to one-word values, and an expression e . The evaluation rules (in [12]) describe how a machine state (M, e) steps to a new state (M', e') . The load and store expressions are the *only* expressions that read or modify M ; even the region and garbage collection examples later in the paper are built entirely on top of simple loads and stores.

The expressions e , values v , and types τ are defined as follows (kinds K , integers I , and booleans B are defined later in this section):

types

$$\begin{aligned} \tau = & \tau_1 \xrightarrow{\phi} \tau_2 \mid \phi(\vec{\tau}) \mid \text{Mem}(I, \tau) \\ & \mid \alpha \mid \lambda \alpha : K. \tau \mid \tau_1 \tau_2 \mid \forall \alpha : K; B. \tau \mid \exists \alpha : K; B. \tau \mid \mu \alpha : K. \tau \\ & \mid I \mid B \mid \text{Int}(I) \mid \text{Bool}(B) \mid \text{if } B \text{ then } \tau_1 \text{ else } \tau_2 \end{aligned}$$

expressions

$$\begin{aligned} e = & i \mid b \mid x \mid e_1 e_2 \mid e \tau \mid \phi(\vec{e}) \mid \lambda x : \tau \xrightarrow{\phi} e \mid e_1 \text{ op } e_2 \mid \neg e \\ & \mid \Lambda \alpha : K; B. v \mid \text{let } \langle \vec{x} \rangle = e_1 \text{ in } e_2 \mid \text{if } e_1 \text{ then } e_2 \text{ else } e_3 \\ & \mid \text{pack}[\tau_1, e] \text{ as } \exists \alpha : K; B. \tau_2 \mid \text{unpack } \alpha, x = e_1 \text{ in } e_2 \\ & \mid \text{roll}[(\mu \alpha : K. \tau_0) \tau_1 \cdots \tau_n](e) \mid \text{unroll}(e) \mid \text{fix } x : \tau. v \\ & \mid \text{load}(e_{\text{ptr}}, e_{\text{Mem}}) \mid \text{store}(e_{\text{ptr}}, e_{\text{Mem}}, e_v) \mid \text{fact} \end{aligned}$$

values

$$\begin{aligned} v = & i \mid b \mid \Lambda \alpha : K; B. v \mid \text{pack}[\tau_1, v] \text{ as } \exists \alpha : K; B. \tau_2 \mid \phi(\vec{v}) \\ & \mid \text{roll}[(\mu \alpha : K. \tau_0) \tau_1 \cdots \tau_n](v) \mid \lambda x : \tau \xrightarrow{\phi} e \mid \text{fact} \end{aligned}$$

Most of the types are borrowed from other languages, as discussed below. The one new type is the the “linear memory type”, $\text{Mem}(I, \tau)$, which is loosely based on alias types. As discussed in the previous section, alias types maintain a linear context that maps locations to types. Alias types

deliberately abstract away the details of how the objects are allocated and where in memory objects reside: in the context $\{\ell_1 \mapsto \langle \tau_1, \tau_2 \rangle, \ell_2 \mapsto \langle \tau_3, \tau_4, \tau_5 \rangle\}$, the “locations” ℓ_1 and ℓ_2 are opaque abstractions, not integers. Linear memory types, on the other hand, are designed to implement object allocation from the ground up, and therefore they must expose the details of the underlying memory. Rather than mapping opaque locations to types, linear memory types map *integer word addresses* to the types of individual memory words. If the two objects above reside at memory locations 500 and 700, then five distinct mappings describe the state of the objects:

$$\{500 \mapsto \tau_1, 501 \mapsto \tau_2, 700 \mapsto \tau_3, 701 \mapsto \tau_4, 702 \mapsto \tau_5\}$$

Our examples do not require some of the more advanced features of alias types, such as explicit store polymorphism and nonlinear constraints. Because of this, it is easiest to express the type of each memory word as an individual first-class linear value, rather than as a member of a constraint set. Traditional linear data structuring mechanisms[27] then suffice to hold the types of multiple memory words. For example, the memory location 500 is described by an object of linear type “ $\text{Mem}(500, \tau_1)$ ”, and the state of the pair of integers at locations 500 and 501 can be stored in a linear tuple of type $\wedge \langle \text{Mem}(500, \tau_1), \text{Mem}(501, \tau_2) \rangle$.

Following Walker et al[29], the type system includes nonlinear and linear functions ($\tau_1 \xrightarrow{\phi} \tau_2$) and nonlinear and linear tuples ($\phi(\vec{\tau})$), where ϕ is \wedge to indicate linear data, and ϕ is \cdot to indicate nonlinear data. A linear functions must be called exactly once, while a nonlinear function can be called arbitrarily often. Linear tuples are consumed when their fields are extracted. Linear data structures may hold nonlinear data (a linear tuple type $\wedge \langle \cdot \rangle, \cdot \langle \cdot \rangle$ containing empty nonlinear tuples is legal), but nonlinear data structures cannot contain linear data (the nonlinear tuple type $\cdot \langle \wedge \rangle, \wedge \langle \cdot \rangle$ containing empty linear tuples is illegal), which ensures that linear data is never aliased. To enforce this restriction, we assign kinds to types that distinguish between linear and nonlinear types. The kind \hat{n} describes all linear types of size n words, while \dot{n} describes all nonlinear types of size n words. For example, the nonlinear empty tuple type $\cdot \langle \cdot \rangle$ has kind $\dot{0}$.

We use the n in the kind \hat{n} to restrict operations on memory: values stored in memory must be exactly 1 word long. The type of singleton booleans $\text{Bool}(\tau)$ has kind $\dot{1}$, for example, so that a boolean value fits in one word of memory. A triplet of booleans of type $\cdot \langle \text{Bool}(\tau), \text{Bool}(\tau), \text{Bool}(\tau) \rangle$, though, has kind $\dot{3}$, and is too large to fit into a single memory word. The kinding rules for tuple types and linear memory types illustrate the size and linearity rules for data (the environment $\Phi; \Delta$ is explained later):

$$\frac{\forall j. (\Phi; \Delta \vdash \tau_j : n_j) \quad n = \sum_j n_j}{\Phi; \Delta \vdash \cdot \langle \vec{\tau} \rangle : \dot{n}}$$

$$\frac{\forall j. (\Phi; \Delta \vdash \tau_j : \hat{n}_j) \quad n = \sum_j n_j}{\Phi; \Delta \vdash \wedge \langle \vec{\tau} \rangle : \hat{n}}$$

$$\frac{\Phi; \Delta \vdash I : \text{int} \quad \Phi; \Delta \vdash \tau : \dot{1}}{\Phi; \Delta \vdash \text{Mem}(I, \tau) : \hat{0}}$$

The type $\text{Mem}(I, \tau)$ has kind $\hat{0}$, so that it is linear and occupies no space. Operations on memory consume values of type $\text{Mem}(I, \tau)$ and produce new values of type $\text{Mem}(I, \tau')$, so that at all times there is exactly one value of type $\text{Mem}(i, \tau)$ for each memory location i . For example, the following function swaps two words of memory, consuming a pair of values of type $\wedge \langle \text{Mem}(500, \tau_1), \text{Mem}(501, \tau_2) \rangle$ and producing a pair of values of type $\wedge \langle \text{Mem}(501, \tau_1), \text{Mem}(500, \tau_2) \rangle$:

$$\begin{aligned} \lambda x : \wedge \langle \text{Mem}(500, \tau_1), \text{Mem}(501, \tau_2) \rangle \longrightarrow \\ \text{let } \langle x_1, x_2 \rangle = x \text{ in} \\ \text{let } \langle y_1, x'_1 \rangle = \text{load}(500, x_1) \text{ in} \\ \text{let } \langle y_2, x'_2 \rangle = \text{load}(501, x_2) \text{ in} \\ \text{let } x''_1 = \text{store}(500, x'_1, y_2) \text{ in} \\ \text{let } x''_2 = \text{store}(501, x'_2, y_1) \text{ in } \wedge \langle x''_2, x''_1 \rangle \end{aligned}$$

When x_1 has type $\text{Mem}(500, \tau_1)$, the expression $\text{load}(500, x_1)$ consumes x_1 (so that x_1 is no longer in scope in the remainder of the function), and produces a pair whose first element contains the contents of memory word 500, and whose second element has type $\text{Mem}(500, \tau_1)$. The typing rule for the expression $\text{load}(e_{\text{ptr}}, e_{\text{Mem}})$ requires that the first argument e_{ptr} be a *singleton integer* of type $\text{Int}(I)$, and the second argument e_{Mem} have type $\text{Mem}(I, \tau)$ for some τ .

$$\frac{C_1 \vdash e_{\text{ptr}} : \text{Int}(I) \quad C_2 \vdash e_{\text{Mem}} : \text{Mem}(I, \tau)}{C_1, C_2 \vdash \text{load}(e_{\text{ptr}}, e_{\text{Mem}}) : \wedge \langle \tau, \text{Mem}(I, \tau) \rangle}$$

Unlike the more usual type “int”, which is the set of all integers, the singleton type $\text{Int}(I)$ is the set containing only one integer I , so that the integer constant 500 has type $\text{Int}(500)$:

$$\dot{C} \vdash i : \text{Int}(i)$$

For the load expression, the I in $\text{Int}(I)$ and $\text{Mem}(I, \tau)$ must match, so that $\text{load}(500, x_1)$ type-checks in the example above, but $\text{load}(499, x_1)$ would not type-check.

The expression $\text{store}(500, x'_1, y_2)$ consumes x'_1 and produces a value of type $\text{Mem}(500, \tau_2)$, documenting the change to the type of the value at memory word 500 from τ_1 to τ_2 :

$$\frac{C_2 \vdash e_{\text{Mem}} : \text{Mem}(I, \tau_1) \quad C_3 \vdash e_v : \tau_2 \quad C_1 \vdash e_{\text{ptr}} : \text{Int}(I) \quad C_1, C_2, C_3 \vdash \tau_2 : \dot{1}}{C_1, C_2, C_3 \vdash \text{store}(e_{\text{ptr}}, e_{\text{Mem}}, e_v) : \text{Mem}(I, \tau_2)}$$

Notice that the store expression overwrites whatever τ_1 value was previously at memory location 500, while the load expression makes a copy of a value from memory. Because only nonlinear values may be discarded and copied, the type $\text{Mem}(I, \tau)$ is only legal if τ is nonlinear.

While the examples above show how to manipulate values of type $\text{Mem}(i, \tau)$, they leave one question unanswered: what exactly is the value that has type $\text{Mem}(i, \tau)$? The syntax for the abstract machine defines a special value, “fact”, which is sort of a universal unit value for all types $\text{Mem}(i, \tau)$. The fact value carries no run-time information, though, and is not represented at all in a real machine. The type-checking rule states that fact is well-typed if the current context contains exactly one linear assumption, saying that memory location i holds a value of type τ :

$$\dot{C}, i \mapsto \tau \vdash \text{fact} : \text{Mem}(i, \tau)$$

3.1 Integer and Boolean Polymorphism

The swap function described earlier only works for fixed types τ_1 and τ_2 , and the particular memory addresses 500 and 501. The type system allows polymorphism over base types (kind \hat{n}), integers (kind “int”), booleans (kind “bool”), and type operators (kind $K_1 \rightarrow K_2$), so the swap function can be rewritten with type:

$$\begin{aligned} \forall \alpha_1 : \text{int}. \forall \alpha_2 : \text{int}. \forall \beta_1 : \dot{1}. \forall \beta_2 : \dot{1}. \\ \wedge \langle \text{Int}(\alpha_1), \text{Int}(\alpha_2), \text{Mem}(\alpha_1, \beta_1), \text{Mem}(\alpha_2, \beta_2) \rangle \longrightarrow \\ \wedge \langle \text{Mem}(\alpha_2, \beta_1), \text{Mem}(\alpha_1, \beta_2) \rangle \end{aligned}$$

Following the approach of Xi et al[32], the singleton integers of type $\text{Int}(\alpha_1)$ and $\text{Int}(\alpha_2)$ are the one-word run-time representations of the integer types α_1 and α_2 . This avoids the need for full-blown dependent types; the type system only needs to deal with a subset of integer arithmetic that can be solved easily by a standard constraint solver: universally quantified variables, addition, subtraction, comparison, and multiplication by constants. We use the symbol I to refer to integer types, and the symbol B to refer to boolean types: $(0 \leq \alpha_1 \wedge \alpha_1 \leq \alpha_2)$ is a legal boolean type, while $\text{Int}(\alpha_1 + 2 * \alpha_2 + 1)$ is a legal singleton integer type that is the run-time representation of the integer type $\alpha_1 + 2 * \alpha_2 + 1$.

arithmetic

$$\begin{aligned} i &= \dots \mid -2 \mid -1 \mid 0 \mid 1 \mid 2 \mid \dots \\ b &= \text{true} \mid \text{false} \\ I &= \alpha \mid i \mid I_1 \text{ iop } I_2 \\ B &= \alpha \mid b \mid \neg B \mid B_1 \text{ bop } B_2 \mid I_1 \text{ cmp } I_2 \\ \text{iop} &= + \mid - \mid * \quad \text{bop} = \wedge \mid \vee \\ \text{cmp} &= < \mid > \mid \leq \mid \geq \mid = \mid \neq \\ \text{op} &= \text{iop} \mid \text{bop} \mid \text{cmp} \end{aligned}$$

In addition to the universal polymorphic type $\forall \alpha : K.\tau$, the type system supports existential polymorphic types $\exists \alpha : K.\tau$. For example, the type $\exists \alpha : \text{int}. \text{Int}(\alpha)$ is the type of a singleton integer, for some unknown integer α ; this type imitates the traditional type “int” (the type of all integers). As another example, the type

$$\exists \alpha : \text{int}. \wedge \langle \text{Int}(\alpha), \text{Mem}(\alpha, \tau_1), \text{Mem}(\alpha + 1, \tau_2) \rangle$$

says that some unknown integer α is an address where two words of type τ_1 and τ_2 are stored, and also provides the singleton integer $\text{Int}(\alpha)$ needed to load and store from these memory addresses; this type imitates a pointer to a pair of τ_1 and τ_2 .

Universal and existential polymorphic types may contain constraints that must be satisfied when types are substituted in for type variables. For example, if swap has the following type:

$$\forall \alpha_1 : \text{int}. \forall \alpha_2 : \text{int}; \alpha_1 \leq \alpha_2. \forall \beta_1 : \dot{1}. \forall \beta_2 : \dot{1} \dots$$

Then (*swap* 3 6 ...) would be legal, but (*swap* 10 6 ...) would be illegal. When type checking the body of the swap function, the fact that $\alpha_1 \leq \alpha_2$ is kept inside an environment B , while the kinds of the type variables are kept in an environment Δ , so that the combined type checking context is now $C \triangleq \Psi; \Delta; \Gamma; B$.

3.2 Type Equivalence

Following Xi’s approach, we define two integer types I_1 and I_2 to be equivalent ($I_1 \equiv I_2$) if our constraint solver can show that for all substitutions of integer constants i for integer variables α in I_1 and I_2 , the types I_1 and I_2 simplify to equal integer constants. For example, the type system considers $\alpha_1 + 2 * \alpha_2 + 1$ to be equivalent to $2 * \alpha_2 + 2 + \alpha_1 - 1$. However, our language contains two types that do not interact gracefully with Xi’s approach: type operator application “ $\tau_1 \tau_2$ ”, which is used for polymorphic abstract data types, and the conditional type “if B then τ_1 else τ_2 ”, which is used to create unions². Conceivably, either of these types could have kind `int`, so that they may be substituted for an integer variable α inside an integer type I during type-checking (the type-checking rule for type application $e \tau$ performs this sort of substitution, for instance). If this happens, the constraint solver may be faced with an integer type like $2 + (\text{if } \beta \text{ then } 3 \text{ else } 4) + (\alpha (\lambda \gamma : \text{int} . \gamma + 3))$, which is difficult to deal with. Nevertheless, we valued the convenience of Xi’s approach in our implementation (and in the examples in this paper) enough to resort to a hack: the kind system bans the kinds `int` and `bool` from appearing as the return type of a type operator (e.g. $\mathbb{1} \rightarrow \text{int}$), so that integer and boolean types do not contain operator applications, and the type system disallows conditional types from having kind `int` or `bool`:

kinds:

$$K = J \mid \text{int} \mid \text{bool} \quad J = \overset{\phi}{n} \mid K \rightarrow J$$

$$\frac{\Phi; \Delta, \alpha : K_a \vdash \tau : J_b}{\Phi; \Delta \vdash \lambda \alpha : K_a . \tau : K_a \rightarrow J_b}$$

$$\frac{\Phi; \Delta \vdash \tau_f : K_a \rightarrow J_b \quad \Phi; \Delta \vdash \tau_a : K_a}{\Phi; \Delta \vdash \tau_f \tau_a : J_b}$$

$$\frac{\Phi; \Delta \vdash B : \text{bool} \quad \Phi; \Delta \vdash \tau_1 : J \quad \Phi; \Delta \vdash \tau_2 : J}{\Phi; \Delta \vdash \text{if } B \text{ then } \tau_1 \text{ else } \tau_2 : J}$$

Once these type and kind restrictions are in place, the type equivalence rules for type application and conditional types no longer interfere with the type equivalence rules for integers and booleans:

$$\Phi; B \vdash (\lambda \alpha : K . \tau_b) \tau_a \equiv [\alpha \mapsto \tau_a] \tau_b$$

$$\Phi; B \vdash (\text{if true then } \tau_1 \text{ else } \tau_2) \equiv \tau_1$$

$$\Phi; B \vdash (\text{if false then } \tau_1 \text{ else } \tau_2) \equiv \tau_2$$

In a typed assembly language or proof-carrying code system, it might make more sense to require the program to provide proofs of arithmetic equivalence using fundamental arithmetic axioms, rather than to rely on a constraint solver[7][25]; in this case, we wouldn’t need to hack the kind system.

²Our technical report uses a different type for unions that requires more explicit programmer annotation, in order to simplify the implementation’s type inference algorithm. The conditional type in this paper is more elegant, though, and is essential for section 6.1’s use of *delayed types*.

3.3 Recursive Types, Lists, and Stacks

As described earlier, an existential type like

$$\exists \alpha : \text{int} . \wedge \langle \text{Int}(\alpha), \text{Mem}(\alpha, \tau_1), \text{Mem}(\alpha + 1, \tau_2) \rangle$$

can imitate a pointer to data in memory. When this idea is combined with recursive types, the type system is powerful enough to express pointer-based data structures in the style of recursive alias types[28]. For example, the following type defines a simple linked list terminated by 0, containing data elements of type τ :

$$\text{List} = \mu \beta : (\text{int} \rightarrow \hat{0}) . \lambda \alpha_1 : \text{int} . \text{if } \alpha_1 = 0 \text{ then } \wedge \langle \rangle \text{ else} \\ \exists \alpha_2 : \text{int} . \wedge \langle \text{Mem}(\alpha_1, \text{Int}(\alpha_2)), \text{Mem}(\alpha_1 + 1, \tau), \beta \alpha_2 \rangle$$

For clarity, we’ll often use the $\hat{=}$ symbol to define data types, rather than using the abstract machine’s official, but more cryptic, recursive type $\mu \alpha : K . \tau$:

$$\text{List}(\alpha_1 : \text{int}) : \hat{0} \hat{=} \text{if } \alpha_1 = 0 \text{ then } \wedge \langle \rangle \text{ else} \\ \exists \alpha_2 : \text{int} . \wedge \langle \text{Mem}(\alpha_1, \text{Int}(\alpha_2)), \text{Mem}(\alpha_1 + 1, \tau), \text{List } \alpha_2 \rangle$$

The type “if B then τ_1 else τ_2 ” is equivalent to τ_1 if B is true and to τ_2 if B is false, so the *List* type contains an empty tuple if $\alpha_1 = 0$ and a three-word tuple for non-zero α_1 . The first two words of the tuple assert that memory locations α_1 and $\alpha_1 + 1$ contain values of type $\text{Int}(\alpha_2)$ (the pointer to the next element) and τ (the data in the current element). By itself, the $\text{Int}(\alpha_2)$ is just an arbitrary integer, but together with the *Mem* type in $\text{List } \alpha_2$ it may be used to load the fields of the next list element. Based on this approach, the language easily supports tree-like data structures. For example, free lists (or lists of free lists of various size objects) provide a simple memory management strategy for linear data.

By allowing simple integer arithmetic, the type system expresses more than just traditional link-based structures. For example, a slight modification of the linked list example produces a stack type, held in a contiguous sequence of memory locations $\alpha_1 \dots \alpha_2 - 1$:

$$\text{Stack}(\alpha_1 : \text{int})(\alpha_2 : \text{int}) : \hat{0} \hat{=} \text{if } \alpha_1 = \alpha_2 \text{ then } \wedge \langle \rangle \text{ else} \\ \wedge \langle \text{Mem}(\alpha_1, \tau), \text{Stack}(\alpha_1 + 1) \alpha_2 \rangle$$

4. COERCION FUNCTIONS

The previous section’s stack type is fine for simple pushes and pops, but it does not support constant-time random access into the middle of the stack. Getting to the middle of the stack is hard to do: the program must unroll the data type n times to get to the n th element, performing tuple extractions with each unroll. Since $O(n)$ abstract machine operations are required, there’s no obvious way to perform a constant time access of an inner element. And yet, unroll operations and tuple extractions for size-zero data aren’t supposed to impose any run-time cost; they are purely compile-time coercions. What is needed is a way to combine n such coercions into a single $O(1)$ time operation. Suppose that an expression e satisfies the following three constraints:

- It has no effect on the memory M
- If it evaluates to a value, then that value’s type has size 0

- It is guaranteed to terminate (by evaluating to a value in a finite time)

Then there is no reason to execute e on a real machine; a compiler can simply erase e in the same way that it erases the pack/unpack and roll/unroll coercions when compiling typed code to untyped machine language code. Such an expression can still perform many coercion operations, such as traversing a stack data structure to retrieve a Mem value from deep inside the stack, but now with no run-time cost, since the expression is erased by the compiler. For example, consider the following data type for arrays:

$$LArray (\alpha_1 : \text{int}) (\alpha_2 : \text{int}) (\beta : \text{int} \rightarrow \hat{0}) : \hat{0} \triangleq \\ \text{if } \alpha_1 = \alpha_2 \text{ then } \hat{\langle} \rangle \text{ else } \hat{\langle} \beta \alpha_1, LArray (\alpha_1 + 1) \alpha_2 \beta \rangle$$

The $LArray$ type is similar to the $Stack$ type, except for extra flexibility it is polymorphic over an operator β so that each array element α_1 can have data of a different type ($\beta \alpha_1$). The following coercion function splits an array into two adjacent arrays, the left containing elements $\alpha_1 \dots \alpha_2 - 1$ and the right containing elements $\alpha_2 \dots \alpha_3 - 1$:

$$\text{fix } split : \tau_{split}. \Lambda \alpha_1 : \text{int}. \Lambda \alpha_2 : \text{int}. \Lambda \alpha_3 : \text{int}. \Lambda \beta : \text{int} \rightarrow \hat{0} \\ ; \alpha_1 \leq \alpha_2 \leq \alpha_3. \lambda arr : (LArray \alpha_1 \alpha_3) \cdot^{(\alpha_2 - \alpha_1)} \\ \text{if } \alpha_1 = \alpha_2 \\ \text{then } \hat{\langle} \text{roll}[LArray \alpha_1 \alpha_2 \beta] (\hat{\langle} \rangle), arr \rangle \\ \text{else let } \langle head, tail \rangle = \text{unroll}(arr) \text{ in} \\ \text{let } \langle left, right \rangle = split (\alpha_1 + 1) \alpha_2 \alpha_3 \beta tail \text{ in} \\ \hat{\langle} \text{roll}[LArray \alpha_1 \alpha_2 \beta] (\hat{\langle} head, left \rangle), right \rangle \\ \tau_{split} \triangleq \forall \alpha_1 : \text{int}. \forall \alpha_2 : \text{int}. \forall \alpha_3 : \text{int}. \\ \forall \beta : \text{int} \rightarrow \hat{0}; \alpha_1 \leq \alpha_2 \leq \alpha_3. (LArray \alpha_1 \alpha_3) \cdot^{(\alpha_2 - \alpha_1)} \\ \hat{\langle} (LArray \alpha_1 \alpha_2), (LArray \alpha_2 \alpha_3) \rangle$$

In the case where $\alpha_1 = \alpha_2$, the left array is empty and the right array is the whole array. Otherwise, the function pops the head of the array, recurses, and pushes the head back onto the left array returned from the recursive call.

For the compiler to omit a call to this coercion function at run-time, it must know that it satisfies the three conditions above: no stores to memory, return type of size 0, and definite termination. Since the language already supports integer arithmetic, an easy (though heavy-handed) way to ensure termination is to annotate the type of each coercion function ($\lambda x : \tau \xrightarrow{\phi, limit} e$) with a nonnegative integer $limit$:

$$\frac{\Phi; \Delta \vdash \tau_1 : \overset{\phi_1}{n_1} \quad \Phi; \Delta \vdash \tau_2 : \overset{\phi_2}{n_2} \\ (\Phi; \Delta \vdash limit : \text{int} \wedge n = 0) \text{ or } (limit = \infty \wedge n = 1)}{\Phi; \Delta \vdash \tau_1 \xrightarrow{\phi, limit} \tau_2 : \overset{\phi}{n}}$$

The type system only allows coercion functions to call functions with lower limits than their own limit. The recursive call above is legal because the caller has limit $(\alpha_2 - \alpha_1)$, as indicated in its annotation, while the called function ($split (\alpha_1 + 1) \alpha_2 \alpha_3 \beta$) has limit $(\alpha_2 - (\alpha_1 + 1))$, which is smaller than $(\alpha_2 - \alpha_1)$. The current limit is stored in the context C , whose definition is extended to be $C = \Psi; \Delta; \Gamma; B; limit$ where $limit$ is either an integer type I (for coercion functions) or the symbol ∞ (for normal functions; we usually omit the ∞ annotation on function types). In a context with an integer limit, expressions are not allowed to store anything to memory, although loads are allowed.

The $split$ function performs one rather unrealistic operation: it tests to see whether two type variables α_1 and α_2 are equal to each other. This *intentional type analysis* would be problematic for a compiler to implement for ordinary functions: if types are erased before run-time, α_1 and α_2 won't be available to use in the run-time equality test. Since coercion functions won't actually be executed at run-time, though, the language lets them use a boolean type test expression “if B then e_1 else e_2 ”.

Extensions to λ^{low} for coercion functions

limit = $I \mid \infty$

types

$$\tau = \dots \mid \tau_1 \xrightarrow{\phi, limit} \tau_2$$

expressions

$$e = \dots \mid \lambda x : \tau \xrightarrow{\phi, limit} e \mid \text{if } B \text{ then } e_1 \text{ else } e_2 \mid \text{coerce}(e)$$

values

$$v = \dots \mid \lambda x : \tau \xrightarrow{\phi, limit} e$$

It's easy to define other coercions for the $LArray$ type: combining two adjacent arrays into a single array, constructing and deconstructing empty arrays, and constructing and deconstructing single-element arrays. These coercions suffice to implement familiar operations to get and set array elements. For example, given an expression of type $(LArray 0 10 (\lambda \alpha : \text{int}. \text{Mem}(500 + \alpha, \tau)))$, describing an array of 10 elements of type τ stored in memory words 500...509, element 5 of the array is accessed by splitting the array into three arrays 0...4, 5...5, and 6...9, retrieving the $\text{Mem}(505, \tau)$ fact from the single-element array 5...5, performing a load or store to memory word 505, and then recombining the three arrays into the original array 0...9.

Viewed from another perspective (courtesy of the Curry-Howard isomorphism), the $split$ function takes a proof as an argument (say, a proof that memory words 500...509 all contain values of type τ) and returns two proofs as a result. The language of coercion functions is then nothing more than a strongly normalizing proof language, although it happens to look suspiciously similar to the original programming language. In the abstract machine and in Clay, the only differences between the proof and programming languages is the treatment of stores (disallowed in the proof language), intentional type analysis (disallowed in the programming language), and function calls (restricted in the proof language). This similarity has the advantage that programmers need not learn a second language to implement “proofs”, and the programming and proof languages interact seamlessly. On the other hand, a specialized, dedicated proof language would be more elegant, powerful, more amenable to automated proof generation tools, and easier to connect to existing libraries of proofs. The LTT system[7], for example, embeds LF[11] proofs in a separate programming language. The TL system[22] also embeds proofs in programs.

5. NONLINEAR DATA STRUCTURES

The previous sections described types for linear data structures, such as linear lists and linear arrays. Later sections

will describe extensions to the type system for building regions of nonlinear data structures. Interestingly, even without extending the type system, linear memory types can express nonlinear data structures. Consider the following C code that creates a circular list with 3 elements, and calls a function to get the 5th list element, which it then splices into the beginning of the list:

```

struct List {
  struct List *next;
};
struct List *nth(int n, struct List *x) {
  if(n==0) return x;
  else return nth(n-1, x->next);
}
struct List a;
struct List b;
struct List c;
a.next = &b;
b.next = &c;
c.next = &a;
a.next = nth(5, &a);

```

This code creates enough unpredictable aliases to thwart a direct representation of the list with alias types or linear memory types. However, rather than making the list a linear type, we can encode the list as a nonlinear function that fetches data from a linear area of memory represented by an abstract type β . The linear area of memory β acts much like a region, while the nonlinear type $List\ \beta$ acts like a nonlinear pointer into the region (something like the type “List at β ” in standard region terminology). As long as β stays around, any $List\ \beta$ value can be used to load and store from the region:

$$List(\beta : \hat{0}) : i \triangleq \exists \alpha : \text{int} . \langle \text{Int}(\alpha), \beta \xrightarrow{\cdot, 0} \wedge \langle \text{Mem}(\alpha, List\ \beta), \text{Mem}(\alpha, List\ \beta) \xrightarrow{\wedge, 0} \beta \rangle \rangle$$

Each list value contains one singleton integer $\text{Int}(\alpha)$, which is the address of the list element. To get the data from the list element, the program calls the nonlinear function of type $\beta \xrightarrow{\cdot, 0} \wedge \langle \text{Mem}(\alpha, List\ \beta), \dots \rangle$. This function consumes the linear “region” β , and returns a linear memory type describing the list element’s data (in this example, the element contains only one word, the field “next”). With the linear memory type, the program is free to load and store the element’s data. When it is finished loading and storing, it calls the linear $\text{Mem}(\alpha, List\ \beta) \xrightarrow{\wedge, 0} \beta$ function to relinquish the linear memory type and reconstruct the region β .

The $\beta \xrightarrow{\cdot, 0} \wedge \langle \text{Mem}(\alpha, List\ \beta), \dots \rangle$ function in $List\ \beta$ manipulates linear data, but is not linear itself; it is merely a middleman that passes linear data back and forth between the region and the program. Thus, the $List\ \beta$ type is nonlinear, and the program can freely copy and discard lists. For example, here is an implementation of function “nth”:

```

nth =  $\Lambda\beta : \hat{0} . \text{fix}$ 
  f : ( $\forall \alpha : \text{int} . \text{Int}(\alpha) \xrightarrow{\cdot} (List\ \beta) \xrightarrow{\cdot} \beta \xrightarrow{\cdot} \wedge \langle List\ \beta, \beta \rangle$ ).
     $\Lambda \alpha : \text{int} . \lambda n : \text{Int}(\alpha) \xrightarrow{\cdot} \lambda x : List\ \beta \xrightarrow{\cdot} \lambda r : \beta \xrightarrow{\cdot}$ 
      if n = 0 then  $\wedge \langle x, r \rangle$  else
      let  $\langle x', r' \rangle = \text{getnext}\ \beta\ x\ r$  in
      f ( $\alpha - 1$ ) (n - 1) x' r'

```

The n th function in λ^{low} is similar to the C implementation, except that it takes a third argument r of type β , and returns r along with the desired list element (in the tuple $\wedge \langle x, r \rangle$) when the function terminates. The r value is needed to load the $next$ field from the list x , as implemented in the following helper function:

```

getnext =  $\Lambda\beta : \hat{0} . \lambda x : List\ \beta \xrightarrow{\cdot} \lambda r : \beta \xrightarrow{\cdot}$ 
  unpack  $\alpha, y = \text{unroll}(x)$  in
  let  $\langle yptr, yf \rangle = y$  in
  let  $\langle mem, zf \rangle = yf\ r$  in
  let  $\langle x', mem' \rangle = \text{load}(yptr, mem)$  in
   $\wedge \langle x', zf\ mem' \rangle$ 

```

The $getnext$ function extracts the list address $yptr$ and the list function yf , which when called yields the linear memory value mem , of type $\text{Mem}(\alpha, List\ \beta)$. With the address and linear memory value, $getnext$ loads the $next$ field from memory and returns it, along with the linear β value. It’s easy to define a $setnext$ function in the same way, using a store instead of a load:

```

setnext =
 $\Lambda\beta : \hat{0} . \lambda x : List\ \beta \xrightarrow{\cdot} \lambda x' : List\ \beta \xrightarrow{\cdot} \lambda r : \beta \xrightarrow{\cdot}$ 
  unpack  $\alpha, y = \text{unroll}(x)$  in
  let  $\langle yptr, yf \rangle = y$  in
  let  $\langle mem, zf \rangle = yf\ r$  in
  zf (store(yptr, mem, x'))

```

With these functions, the C code

```
a.next = nth(5, &a);
```

is implemented with a function:

```

fifth =  $\Lambda\beta : \hat{0} . \lambda a : List\ \beta \xrightarrow{\cdot} \lambda r : \beta \xrightarrow{\cdot}$ 
  let  $\langle b, r' \rangle = \text{nth}\ \beta\ 5\ 5\ a\ r$  in
  setnext  $\beta\ a\ b\ r'$ 

```

All that remains is to allocate and initialize the circular list. For this, we need a concrete type for β , and we choose a type that describes three words of memory, at addresses 1, 2, and 3, each containing a list. Let’s call this type $Three$:

$$Three : \hat{0} \triangleq \wedge \langle \text{Mem}(1, List\ Three), \text{Mem}(2, List\ Three), \text{Mem}(3, List\ Three) \rangle$$

To create a list of type $List\ Three$, we need to build a pair of an address and a function. Let’s start with address 2, corresponding to the object “b” from the C code:

```

b = roll[List\ Three](
  pack[2,  $\cdot(2, f_b)$ ] as  $\exists \alpha : \text{int} . \langle \text{Int}(\alpha), Three \xrightarrow{\cdot, 0} \dots \rangle$ )

```

The function f_b must extract a $\text{Mem}(2, List\ Three)$ value from a $Three$ value, and return this value along with a function that reconstructs the original $Three$ value:

```

f_b =  $\lambda r : Three \xrightarrow{\cdot, 0}$ 
  let  $\langle m_1, m_2, m_3 \rangle = r$  in
   $\wedge \langle m_2, \lambda m'_2 : \text{Mem}(2, List\ Three) \xrightarrow{\wedge, 0} \wedge \langle m_1, m'_2, m_3 \rangle \rangle$ 

```

In a similar way, we can construct a value a for address 1 and a value c for address 3. Finally, we implement the C code for initialization

```

a.next = &b;
b.next = &c;
c.next = &a;

```

using three store operations, assuming we are given memory words 1, 2, and 3 initialized to some arbitrary junk types τ_1 , τ_2 , and τ_3 , which we overwrite:

```

λm : ^⟨Mem(1, τ1), Mem(2, τ2), Mem(3, τ3)⟩ →
  let ⟨m1, m2, m3⟩ = m in
  let ⟨m'1, m'2, m'3⟩ = ^⟨
    store(1, m1, b),
    store(2, m2, c),
    store(3, m3, a)⟩ in
  fifth Three a ^⟨m'1, m'2, m'3⟩

```

We chose an example based on a single, minimal list type for simplicity. The techniques above extend to more complex data types and mixtures of data types, though. Consider a type

```

struct LTree {
  struct List *list;
  struct LTree *left;
  struct LTree *right;
};

```

This is represented as:

$$LTree(\beta : \hat{0}) : \hat{1} \triangleq$$

$$\begin{aligned} & \exists \alpha : \text{int} . \langle \text{Int}(\alpha), \beta \xrightarrow{\cdot 0} \wedge \langle \tau_{mem}, \tau_{mem} \xrightarrow{\wedge, 0} \beta \rangle \rangle \\ \tau_{mem} = & \wedge \langle \text{Mem}(\alpha, List \beta), \\ & \text{Mem}(\alpha + 1, LTree \beta), \\ & \text{Mem}(\alpha + 2, LTree \beta) \rangle \end{aligned}$$

6. TYPE SEQUENCES AND REGIONS

The previous section’s encoding of nonlinear data structures in “regions” suffers from a serious limitation: the regions cannot grow dynamically. Suppose we wanted to add a fourth list to the region defined by *Three*. We’d need to change the region type from *Three* to *Four*:

$$\begin{aligned} \text{Three} : \hat{0} \triangleq & \wedge \langle \text{Mem}(1, List \text{Three}), \\ & \text{Mem}(2, List \text{Three}), \\ & \text{Mem}(3, List \text{Three}) \rangle \\ \text{Four} : \hat{0} \triangleq & \wedge \langle \text{Mem}(1, List \text{Four}), \\ & \text{Mem}(2, List \text{Four}), \\ & \text{Mem}(3, List \text{Four}), \\ & \text{Mem}(4, List \text{Four}) \rangle \end{aligned}$$

Unfortunately, the type *Three* is embedded in the types of all the existing lists. It’s not clear how coerce a list of type *List Three* to type *List Four*. This section combines the intuition behind the previous section’s encoding with a new language mechanism called *type sequences* to encode regions that grow dynamically. Type sequences allow a program to refine an existing type τ at run-time; a list of type *List* τ stays valid as τ is refined, so that the program can add new data to a region without needing to coerce the old data to a different type. We need to be very careful about our notion of type refinement, though: changing an existing type arbitrarily is certainly unsafe. However, if we underspecify τ to start with, by leaving an abstract hole “?” in the type:

$$\tau = \wedge \langle \text{Mem}(1, List \tau), \wedge \langle \text{Mem}(2, List \tau), \wedge \langle \text{Mem}(3, List \tau), ? \rangle \rangle \rangle$$

then we can later specify the hole “?” to be a particular type, such as $\text{Mem}(4, List \tau)$, as long as we don’t try to specify “?” in two different, conflicting ways. We use linearity to ensure that a hole is filled in at most once.

To make the encoding of regions more obvious, we start by extending the type system to support not just individual holes, but infinite vectors of holes. Afterwards, we show how to simplify the type system by encoding vectors of holes using individual holes.

We define a *type sequence* to be an infinite vector of holes, which are filled in one at a time, in order. For example, if F is a type sequence, then $(F 0)$ is specified first, followed by $(F 1)$, then $(F 2)$, and so on.

The key properties of type sequences are:

- Large, dynamic namespace: F supplies an unbounded number of names for types: $(F 0)$, $(F 1)$, $(F 2)$, ..., and F grows at run-time as new elements are added to the sequence.
- Indelibility: once a type τ is assigned to a name $(F I)$, then $(F I)$ will always refer to τ , and no other type.

Using these properties, we can define a region as a mapping from memory words to elements of a type sequence:

$$(LArray\ 0\ N\ (\lambda I : \text{int} . \text{Mem}(I, F I)))$$

As F is filled in, the region’s array grows to include more and more linear memory types $\text{Mem}(I, F I)$.

Extensions to λ^{low} for type sequences

types

$$\tau = \dots \mid F^K \mid \text{Gen}(\tau, I) \mid \text{Eq}(\tau_1, \tau_2) \mid \text{InDomain}(I, \tau)$$

expressions

$$\begin{aligned} e = \dots \mid & \text{make_eq}(\tau) \mid \text{apply_eq}(\tau, e_1, e_2) \mid \text{new_seq}(J) \\ & \mid \text{discard_seq}(e) \mid \text{define_seq}(e, \tau) \mid \text{in_domain}(I_1, I_2, e_1, e_2) \end{aligned}$$

The extensions to λ^{low} for handling type sequences are shown above. First, we define an equality type $\text{Eq}(\tau_1, \tau_2)$, where a value of type $\text{Eq}(\tau_1, \tau_2)$ is evidence that τ_1 is equivalent to τ_2 . The $\text{make_eq}(\tau)$ expression creates a new equality value of type $\text{Eq}(\tau, \tau)$, for any type τ , and the expression $\text{apply_eq}(\tau_f, e_{eq}, e_f)$ uses an equality value e_{eq} of type $\text{Eq}(\tau_a, \tau_b)$ to substitute τ_b for τ_a in selected locations inside the type of e_f , effectively coercing e_f to a different, but equivalent, type:

$$\frac{C \vdash \tau : K}{C \vdash \text{make_eq}(\tau) : \text{Eq}(\tau, \tau)}$$

$$\frac{C_1, C_2 \vdash \tau_f : K \rightarrow J \quad C_1 \vdash e_{eq} : \text{Eq}(\tau_a, \tau_b) \quad C_1, C_2 \vdash \tau_a : K \quad C_1, C_2 \vdash \tau_b : K \quad C_2 \vdash e_f : \tau_f \tau_a}{C_1, C_2 \vdash \text{apply_eq}(\tau_f, e_{eq}, e_f) : \tau_f \tau_b}$$

The `new_seq(J)` expression creates a new type sequence F of kind $\text{int} \rightarrow J$ for any non-integer, non-boolean kind J :

$$C \vdash \text{new_seq}(J) : \exists F : (\text{int} \rightarrow J). \text{Gen}(F, 0)$$

The allocation of new types in the sequence is controlled by a size-zero linear *generator* of type $\text{Gen}(F, I)$. The expression `define_seq(e_{gen}, τ)`, where e_{gen} has type $\text{Gen}(F, I)$, adds a new type τ to the sequence. It consumes the old generator and returns three values, all of size zero:

- a new generator, of type $\text{Gen}(F, I + 1)$
- a nonlinear proof that $(F I) = \tau$, of type $\text{Eq}(F I, \tau)$
- a nonlinear proof that I will always be in the domain of F (since sequences grow but don't shrink), of type $\text{InDomain}(I, F)$.

$$\frac{C \vdash e : \text{Gen}(\tau_f, I) \quad C \vdash \tau_f : \text{int} \rightarrow J \quad C \vdash \tau_a : J}{C \vdash \text{define_seq}(e, \tau_a) : \wedge \langle \text{Gen}(\tau_f, I+1), \text{Eq}(\tau_f I, \tau_a), \text{InDomain}(I, \tau_f) \rangle}$$

The $\text{Eq}(F I, \tau)$ type is used to substitute τ for $(F I)$ and vice-versa, via the `apply_eq(τ_f, e_{eq}, e_f)` expression described above. The $\text{InDomain}(I_1, F)$ type, when combined with the current generator of type $\text{Gen}(F, I_2)$ produces evidence that $0 \leq I_1 < I_2$, which is useful when I_1 is an index into an array of length I_2 :

$$\frac{C_1, C_2 \vdash I_1 : \text{int} \quad C_1, C_2 \vdash I_2 : \text{int} \quad C_1 \vdash e_1 : \text{InDomain}(I_1, \tau_f) \quad C_2 \vdash e_2 : \text{Gen}(\tau_f, I_2)}{C_1, C_2 \vdash \text{in_domain}(I_1, I_2, e_1, e_2) : \wedge \langle \text{Know}(0 \leq I_1 \wedge I_1 < I_2), \text{Gen}(\tau_f, I_2) \rangle}$$

Here, $\text{Know}(B)$ is an abbreviation for $\exists \alpha : \text{bool}; B. \cdot \langle \rangle$.

When the program is finished adding elements to a sequence, it may use the `discard_seq(e)` expression to discard the linear generator:

$$\frac{C \vdash e : \text{Gen}(\tau, I)}{C \vdash \text{discard_seq}(e) : \cdot \langle \rangle}$$

The nonlinearity of the $\text{Eq}(F I, \tau)$ and $\text{InDomain}(I, F)$ values is the basis for building nonlinear data structures. Suppose a region consists of a single block of memory, starting at address $Base$ and containing $Size$ words (for clarity, we'll often use italicized, capitalized letters and words for type variables rather than Greek letters). Then a nonlinear pointer to a type τ in the region is a triplet of type $\cdot \langle \text{Int}(Base + I), \text{Eq}(F I, \tau), \text{InDomain}(I, F) \rangle$. The region itself contains an array of linear facts, each of the type $\text{Mem}(Base + I, F I)$:

$$\begin{aligned} & \text{Region} (F : \text{int} \rightarrow \mathbb{1}) (Base : \text{int}) (Alloc : \text{int}) \\ & (Size : \text{int}) : \hat{0} \triangleq \wedge \langle \\ & \quad \text{Gen}(F, Alloc), \\ & \quad (LArray\ 0\ Alloc\ (\lambda I : \text{int}. \text{Mem}(Base + I, F I))), \\ & \quad (LArray\ Alloc\ Size\ (\lambda I : \text{int}. \tau_{free})) \rangle \\ & \tau_{free} \triangleq \exists \beta : \mathbb{1}. \text{Mem}(Base + I, \beta) \end{aligned}$$

The first element of the *Region* tuple is the sequence generator, the second is the array of allocated memory, which grows over time, and the third is the array of free memory, which shrinks over time. Given a region and a pointer, loading the pointed-to word consists of three steps:

- use the $\text{InDomain}(I, F)$ value from the pointer together with the $\text{Gen}(F, Alloc)$ value from the region to conclude that $0 \leq I < Alloc$.

- Now that I is known to be within the bounds of the $(LArray\ 0\ Alloc\ \dots)$ array, use element I of the array, of type $\text{Mem}(Base + I, F I)$, to load a value of type $(F I)$. Call this loaded value x .

- Finally, use the $\text{Eq}(F I, \tau)$ value from the pointer says that coerce x from type $(F I)$ to type τ .

Storing a word follows a similar sequence of steps. Notice that the middle step (using array element I to perform a load) is very similar to the operation performed by the *getnext* function from the previous section; both approaches temporarily borrow a linear memory type from a region, use it for a load, and then return the linear memory type to the region.

Each new allocation in the region grows the region's allocation array and shrinks the region's free array. When the program is done using the region, it merges the allocation array back into the free array, and the connection between memory and the $(F I)$ types is lost. The pointer types $\cdot \langle \text{Int}(Base + I), \text{Eq}(F I, \tau), \text{InDomain}(I, F) \rangle$ are still legal types, but without the $\text{Mem}(Base + I, F I)$ facts in the allocation array, they are no longer useful.

At this point, the program can use the free array to allocate other objects. In particular, it can use the free array to create a new region in place of the old region. Although it would be possible to continue using the old generator, appending new types on the end of the old sequence, it seems easier to discard the old generator and create a new sequence, with a new generator. Therefore, the abstract machine environment tracks the state of a *set* of type sequences, which grows as new sequences are allocated:

$$\begin{aligned} \Phi = \{ & F_1^{K_1} \xrightarrow{\phi_1} \{0 \mapsto \tau_{1,0}, 1 \mapsto \tau_{1,1}, \dots\}, \\ & F_2^{K_2} \xrightarrow{\phi_2} \{0 \mapsto \tau_{2,0}, 1 \mapsto \tau_{2,1}, \dots\}, \\ & \dots \} \end{aligned}$$

Each sequence is assigned an identifier $F_i^{K_i}$, which acts as a type operator of kind $(\text{int} \rightarrow K_i)$. We've been a little slippery about what exactly a sequence is in the abstract machine up to this point; now we can state more precisely that a sequence identifier is a type operator, and type variables (such as the F used in the region example above) may be bound to sequence identifiers. The new (and final) abstract machine environment is defined as $C = \Psi; \Phi; \Delta; \Gamma; B; \text{limit}$. One tricky point: each $F_i^{K_i}$ in Φ must suffice to type-check both $\text{Gen}(F_i^{K_i}, I)$ expressions, which are linear, and other expressions containing the type operator $F_i^{K_i}$, which may be nonlinear. The $\text{Gen}(F_i^{K_i}, I)$ expression is only valid in a context $\{\dots, F_i^{K_i} \xrightarrow{\phi_i} \dots\}$ where $\phi_i = \wedge$. Proofs of the following theorems are found in [12]:

Preservation: If $C \vdash (M, e : \tau)$ and (M, e) steps to (M', e') , then there is some C' so that $C' \vdash (M', e' : \tau)$

Progress: If $\Psi; \Phi; \emptyset; \emptyset; \text{true}; \text{limit} \vdash (M, e : \tau)$, and e is not a value, then (M, e) steps

Strong normalization: If $\Psi; \Phi; \emptyset; \emptyset; \text{true}; i \vdash (M, e : \tau)$, then (M, e) steps to a value (M, v) in a finite number of steps

Erasure: If $\Psi; \Phi; \emptyset; \emptyset; true; limit \vdash (M, e : \tau)$, then

(i) If (M, e) steps to (M', e') , then $erase((M, e))$ steps to $erase((M', e'))$ in zero or one steps,

(ii) If $erase(e)$ is a value, then (M, e) steps to some (M', v) such that $erase((M, e)) = erase((M', v))$ in zero or more steps,

(iii) If $erase((M, e))$ steps to some (M'_{erase}, e'_{erase}) , then (M, e) steps in one or more steps to some (M', e') that erases to (M'_{erase}, e'_{erase}) .

Our definition of $erase((M, e))$ erases types and calls to coercion functions, just as the real Clay compiler does.

6.1 Delayed types

Type sequences are sufficient to implement regions, but they seem uncomfortably complicated and *ad hoc*. This section shows that type sequences can be implemented using simpler constructs. Rather than adding types $Eq(\tau_1, \tau_2)$, $Gen(\tau, I)$, and $InDomain(I, \tau)$ to the type system directly, this section builds these types out of a single new type, $Delayed(\tau)$. Using the encodings from this section, we reimplemented regions in Clay without relying on built-in $Eq(\tau_1, \tau_2)$, $Gen(\tau, I)$, and $InDomain(I, \tau)$ types.

First, it's easy to emulate the $Eq(\tau_1, \tau_2)$ type with coercion functions. Following the encoding of $Eq(\tau_1, \tau_2)$ with functions presented in Baars and Swierstra [3], we write:

$$\begin{aligned} Eq_K(\alpha : K)(\beta : K) &: \hat{0} \triangleq \forall F : K \rightarrow \hat{0} . F \alpha \xrightarrow{\hat{0}} F \beta \\ make_eq_K &= \Lambda \alpha : K . \Lambda F : K \rightarrow \hat{0} . \lambda x : F \alpha \xrightarrow{\hat{0}} x \\ apply_eq_{K, \hat{0}} &= \Lambda \alpha : K . \Lambda \beta : K . \Lambda F : K \rightarrow \hat{0} . \\ \lambda eq : (Eq_K \alpha \beta) &\xrightarrow{\hat{0}} \lambda x : F \alpha \xrightarrow{\hat{0}} eq F x \\ apply_eq_{K, \hat{0}} &= \Lambda \alpha : K . \Lambda \beta : K . \Lambda F : K \rightarrow \hat{0} . \\ \lambda eq : (Eq_K \alpha \beta) &\xrightarrow{\hat{0}} \lambda x : F \alpha \xrightarrow{\hat{0}} \\ let \langle y \rangle = eq(\lambda \gamma : K . \wedge \langle F \gamma \rangle) \wedge \langle x \rangle &in y \end{aligned}$$

One drawback to this encoding is that a coercion function only returns size-zero values, so $apply_eq$ can only act on arguments x with type of kind $\hat{0}$ or $\hat{0}$. In practice, this has only been a minor inconvenience. For example, instead of using $apply_eq$ to modify the type of a one-word value loaded from a region, as in the previous section, our new region implementation uses $apply_eq$ to modify the type of the $Mem(Base + I, FI)$ value to $Mem(Base + I, \tau)$ before performing the load.

To implement the $Gen(\tau, I)$ and $InDomain(I, \tau)$ types, we introduce *delayed types*, which implement individual holes instead of infinite vectors of holes. Only two new expressions are needed, one to create a delayed type (creating one new hole named α), and the other to fill in the hole (specifying that α is in fact equal to a new type τ_b):

Extensions to λ^{low} for delayed types

types

$$\tau = \dots \mid F^K \mid Delayed(\tau)$$

expressions

$$e = \dots \mid delay_type(K) \mid specify_type(e_{delayed}, \tau_b, \tau_f, e_f)$$

kinding

$$\frac{\Phi; \Delta \vdash \tau : K}{\Phi; \Delta \vdash Delayed(\tau) : \hat{0}}$$

typing

$$\dot{C} \vdash delay_type(K) : \exists \alpha : K . Delayed(\alpha)$$

$$\frac{C_1 \vdash e_{delayed} : Delayed(\tau_a) \quad C_1, C_2 \vdash \tau_a : K \quad C_1, C_2 \vdash \tau_b : K \quad C_1, C_2 \vdash \tau_f : K \rightarrow J \quad C_2 \vdash e_f : \tau_f \tau_a}{C_1, C_2 \vdash specify_type(e_{delayed}, \tau_b, \tau_f, e_f) : \tau_f \tau_b}$$

The $delay_type$ expression is similar to new_seq , while $specify_type$ is similar to $define_seq$. Note that $specify_type$ consumes its $Delayed(\tau)$ argument, so no separate expression is needed to discard $Delayed(\tau)$ values. Furthermore, $specify_type$ substitutes τ_b for τ_a directly, rather than relying on Eq and $apply_eq$.

Consider a type sequence F that grows to include types $F_0 = \tau_0$, $F_1 = \tau_1$, $F_2 = \tau_2$, and so on. Delayed types can emulate this with a series of functions F_0, F_1, F_2, \dots , which are defined so that:

$$\begin{aligned} F_0 M &= (\text{if } M = 0 \text{ then } \tau_0 \text{ else } F_1 M) \\ F_1 M &= (\text{if } M = 1 \text{ then } \tau_1 \text{ else } F_2 M) \\ F_2 M &= (\text{if } M = 2 \text{ then } \tau_2 \text{ else } F_3 M) \\ &\dots \end{aligned}$$

Each function F_i fills in one element of the type sequence, but defers the rest of the sequence to the next function F_{i+1} . The result is that the function F_0 fulfills the obligations of the type sequence F : $F_0 0 = \tau_0$, $F_0 1 = \tau_1$, $F_0 2 = \tau_2$, and so on, but we don't have to specify all the τ_i at once, because the specification of each function F_i is delayed until τ_i is known. Here is an implementation of this idea that almost implements $Gen(F, I)$ (call it $PreGen F I$):

$$\begin{aligned} PreGen_J(F_0 : \text{int} \rightarrow J)(N : \text{int}) &: \hat{0} \triangleq \\ \exists F_N : \text{int} \rightarrow J . \wedge \langle Delayed(F_N), \\ \forall M : \text{int}; M \geq N . \langle \cdot \rangle &\xrightarrow{\hat{0}} Eq(F_0 M, F_N M) \rangle \end{aligned}$$

For any particular N , $(PreGen F_0 N)$ describes a state where F_N is yet to be specified, but all F_0, F_1, \dots, F_{N-1} are already specified, so that for any $M \geq N$, we know

$$\begin{aligned} F_0 M &= (\text{if } M = 0 \text{ then } \tau_0 \text{ else } F_1 M) \\ &= F_1 M = (\text{if } M = 1 \text{ then } \tau_1 \text{ else } F_2 M) \\ &\dots \\ &= F_{N-1} M = (\text{if } M = N - 1 \text{ then } \tau_{N-1} \text{ else } F_N M) \\ &= F_N M \end{aligned}$$

This invariant is maintained by a coercion function that returns $Eq(F_0 M, F_N M)$ for any $M \geq N$. Suppose that $(PreGen F_0 N)$ exists, and we want to add τ_N to the sequence, producing a $(PreGen F_0 N+1)$ value and an $Eq(F_0 N, \tau_N)$ value. We allocate a new delayed type $Delayed(F_{N+1})$, and consume the old value $Delayed(F_N)$ by specifying

$$F_N = \lambda M : \text{int} . (\text{if } M = N \text{ then } \tau_N \text{ else } F_{N+1} M)$$

This definition immediately implies that $Eq(F_N N, \tau_N)$. The invariant that $Eq(F_0 M, F_N M)$ for all $M \geq N$ implies that $Eq(F_0 N, F_N N)$, so we can combine the two Eq values to conclude $Eq(F_0 N, \tau_N)$. The definition of F_N also implies

that for any $M \geq N + 1$, $\text{Eq}(F_N M, F_{N+1} M)$. Since we already know that for any $M \geq N$, $\text{Eq}(F_0 M, F_N M)$, we can combine Eq values to conclude that for any $M \geq N + 1$, $\text{Eq}(F_0 M, F_{N+1} M)$. Together with the $\text{Delayed}(F_{N+1})$ value, this gives us a value of type $(\text{PreGen } F_0 N + 1)$.

$\text{Gen}(F, N)$ supports one feature that $(\text{PreGen } F N)$ lacks: its sequence extension expression returns not only $\text{Gen}(F, N + 1)$ and $\text{Eq}(F N, \tau_N)$ values, but also an $\text{InDomain}(N, F)$ value, which is essential to the region encoding. To implement this, we define a $(\text{Gen } F G N)$ type that uses two PreGen values, one to handle the type sequence F , and one dedicated to producing $(\text{InDomain } N G)$ values. The functions F and G are defined so that $F 0 = \tau_0$, $F 1 = \tau_1$, ..., $F N - 1 = \tau_{N-1}$, and $G 0 = \cdot\langle \rangle$, $G 1 = \cdot\langle \rangle$, ..., $G N - 1 = \cdot\langle \rangle$.

$$\begin{aligned} \text{InDomain}(M : \text{int})(G : \text{int} \rightarrow \dot{0}) : \dot{0} &\triangleq G M \\ \text{Gen}_J(F : \text{int} \rightarrow J)(G : \text{int} \rightarrow \dot{0})(N : \text{int}) : \hat{0} &\triangleq \\ &\wedge \langle \text{PreGen}_J F N, \text{PreGen}_G G N \rangle \end{aligned}$$

A $(\text{PreGen } G N)$ value is evidence that types $G 0, G 1, \dots, G N - 1$ have all been specified, but that types $G N, G N + 1, G N + 2, \dots$ are still unspecified. An $(\text{InDomain } M G)$ value is evidence that $G M$ has been specified (because there's no way to construct a value of an unspecified type). Logically, then, we would conclude from a $(\text{PreGen } G N)$ value and an $(\text{InDomain } M G)$ value that M cannot be in the range $N, N + 1, N + 2, \dots$ because $G M$ cannot be both specified and unspecified. Thus, we conclude that $M < N$. In our type system, a program derives this conclusion using the expression “if $M < N$ then e_1 else e_2 ”. In e_1 , $M < N$ is known to be true. In e_2 , $M \geq N$ is known to be true, but this leads to a contradiction: $(\text{PreGen } G N)$ lets e_2 specify $G M$ any way it wants if $M \geq N$; instead of making $G M = \cdot\langle \rangle$, e_2 defines $G M = \forall \alpha : \dot{0} . \alpha$, so that $\text{Eq}(G M, \forall \alpha : \dot{0} . \alpha)$ holds. But this lets e_2 change the type of the $(\text{InDomain } M G)$ value from $(G M)$ to $\forall \alpha : \dot{0} . \alpha$, and suddenly e_2 holds a “contradiction”: a value of type $\forall \alpha : \dot{0} . \alpha$, which can be used to prove anything, including $M < N$. Thus, both e_1 and e_2 prove $M < N$. Note that this “contradiction” does not mean that the type system is unsafe; instead, it simply means that e_2 is never evaluated – in any expression of the form “if true then e_1 else e_2 ”, e_2 can derive a contradiction because it assumes true to be false, but this is harmless at run-time, because e_1 is executed, not e_2 .

7. SIMPLE COPYING COLLECTION

We have used λ^{low} 's type system to implement several typed garbage collectors, including an incremental mark-sweep collector and two polymorphic copying collectors. The collectors follow the ideas of Wang et al[30] and Monnier et al[15], but adds several new improvements:

- The copying collectors are genuine Cheney queues, in which all the collection state is stored in to-space. In contrast to the previously described typed garbage collectors, our collectors do not need an auxiliary stack to implement a recursive descent through the live data.
- The approach of Wang et al requires that the heap have a monomorphic type, which imposes inefficiencies on the compiled code. The approach of Monnier et al uses intentional type analysis to handle polymorphism, but this still leaves a subtle inefficiency in the

collector: the collector must perform run-time analysis and processing to pack, unpack, roll, and unroll expressions. In addition, there must be some sort of run-time tag bits or tag words to identify existentials, rolled types, and so on. λ^{low} can hide these operations inside coercion functions, which impose no run-time space or time cost.

- The collectors supports cyclic data structures, correctly handling mixtures of pointers that may be null and pointer types that don't allow null.
- The collectors make all data layout and tag information explicit, down to the last bit. This demonstrates that it is possible to implement efficient representations of header words and forwarding pointers in a typed collector.

More information about these collectors, including complete implementations in Clay, is available in Wei[31] and at the URL <http://www.cs.dartmouth.edu/~hawblitz/>. The technical report[12] also contains a translation from λ^C , a CPS and closure-converted intermediate language developed for typed assembly language by Morrisett et al[17], to λ^{low} , using one of our copying collectors for memory management. Since [17] already provides a translation from a variant of Girard's System-F (the polymorphic lambda calculus) to λ^C , this forms a complete translation from a high-level polymorphic language to λ^{low} . The details of the collectors are beyond the scope of this paper, though; for brevity's sake, this section describes only a simple monomorphic Cheney queue typed garbage collector, based on the regions from the previous section.

A Cheney queue collector starts with a root pointer into a region of memory (from-space), and makes a breadth-first traverse through all the data reachable from that root. The collector blindly copies each traversed object into the to-space region, even though the object still contains pointers that point to from-space, and then later goes back to fix up the pointers so that to-space objects points to other to-space objects. After all live data is traversed, the collector deallocates the from-space region and, for the next collection, allocates a new to-space region (which is usually just the old from-space region, recycled).

Ideally, the type of each pointer in from-space would say that it points to another object in from-space, and the type of each pointer in to-space would say that it points to another object in to-space; in this case, each object need only concern itself with its own region. Unfortunately, the blind copy leaves temporary objects in the to-space Cheney queue pointing back to from-space. Furthermore, the collector tags the copied from-space object with a forwarding pointer into to-space, so that it won't attempt to copy the same object more than once. Since there are pointers into to-space from from-space and vice-versa, the type of each object must be aware not only of its own region, but of the previous and next regions.

To equip the objects with information about multiple regions, we build a type sequence of type sequences: the outer type sequence R contains one type sequence (RE) for each region number E (the E is supposed to stand for “epoch”), and each (REI) describes the I th word of memory in region number E . If region E is from-space, then a from-space object can name the type of the I th word in to-space using the

type $(R(E+1)I)$. To add even more flexibility, each word descriptor (REI) is actually a type operator that takes a region number as an argument: $(REIE')$ is the type of the word stored in region E , word I , configured with the extra information E' . This extra information describes the types of the pointers contained in an object: even though the object *lives* in region E , it might *point* to a different region E' . This captures the state of the objects in the Cheney queue, which have been copied to region E , but still point to region $E-1$. The type of region E now contains three arrays: one for the finished objects that point to region E , one for the Cheney queue objects that still point to region $E-1$, and one for the free space.

$$\begin{aligned} GcRegion(R : K_R)(E : \text{int})(Scan : \text{int})(Alloc : \text{int}) : \hat{0} \triangleq & \\ \wedge \langle \text{Gen}(R E, Alloc), & \\ (LArray\ 0\ Scan\ (\lambda I : \text{int}. \text{Mem}(Addr(E, I), REIE)), & \\ (LArray\ Scan\ Alloc\ (\lambda I : \text{int}. & \\ \text{Mem}(Addr(E, I), REI(E-1))), & \\ (LArray\ Alloc\ SIZE\ (\lambda I : \text{int}. & \\ \exists \beta : \hat{1}. \text{Mem}(Addr(E, I), \beta))) \rangle & \\ K_R \triangleq \text{int} \rightarrow \text{int} \rightarrow \text{int} \rightarrow \hat{1} & \end{aligned}$$

Before the types get too complicated, we introduce a simplifying assumption: all regions have the same size $SIZE$, all even numbered regions use memory words $BASE \dots (BASE+SIZE-1)$, and all odd numbered regions use memory words $(BASE+SIZE) \dots (BASE+2*SIZE-1)$. The $(Base+I)$ from the previous definition of regions is changed to

$$Addr(E, I) \triangleq (BASE + SIZE * (E \text{ mod } 2) + I)$$

The type language doesn't yet contain a *mod* operator, so the real Clay implementation of the collector uses two variables E_{hi} and E_{lo} to form a region number $2 * E_{hi} + E_{lo}$, where $0 \leq E_{lo} \leq 1$. Since this would add much clutter and little illumination to the presentation below, this paper uses just the single variable E .

The simple collector in this section only supports one object type, $GcObject$, defined to contain a forwarding pointer, one floating point data field, and two possibly-null pointers to other $GcObjects$. The $GcObject$ type describes the non-linear type of each of the 4 words of the object:

$$\begin{aligned} GcObject(R : K_R)(E : \text{int})(This : \text{int}) : \hat{0} \triangleq \cdot \langle & \\ (GcWord\ R\ E\ (This + 0)\ (GcFwd\ R)), & \\ (GcWord\ R\ E\ (This + 1)\ (GcPrim\ \text{float})), & \\ (GcWord\ R\ E\ (This + 2)\ (GcPtr\ R)), & \\ (GcWord\ R\ E\ (This + 3)\ (GcPtr\ R)) \rangle & \\ GcWord(R : K_R)(E : \text{int})(This : \text{int})(A : \text{int} \rightarrow \hat{1}) : \hat{0} \triangleq \cdot \langle & \\ \text{Eq}(R\ E\ This, A), & \\ \text{InDomain}(This, R\ E) \rangle & \end{aligned}$$

The abbreviation $GcWord$ describes the type A of a single word $This$ in a region E , where the type A is actually a type operator $(A : \text{int} \rightarrow \hat{1})$ that takes a region number E' as an argument. For non-pointer data types, like float, the E' argument is irrelevant, and the A operator always returns the same type:

$$GcPrim(T : \hat{1})(E' : \text{int}) : \hat{1} \triangleq T$$

Pointer types, on the other hand, use the E' to identify the region containing the pointed-to object:

$$\begin{aligned} GcPtr(R : K_R)(E' : \text{int}) : \hat{1} \triangleq \exists \alpha : \text{int}. \cdot \langle \text{Int}(Addr(E', \alpha)), & \\ \text{if } Addr(E', \alpha) = 0 \text{ then } \cdot \langle \rangle \text{ else } GcObject\ R\ E'\ \alpha \rangle & \end{aligned}$$

A non-null pointer is a pair of the singleton integer containing the real memory address of the target object, and a $GcObject$ describing the state of the target object. A null pointer is just the singleton integer $\text{Int}(0)$. Forwarding pointers are just like ordinary pointers, except that they point into the next region:

$$GcFwd(R : K_R)(E' : \text{int}) : \hat{1} \triangleq GcPtr\ R\ (E' + 1)$$

Garbage collection starts with to-space entirely free. Suppose that from-space is region number 5, to-space is region number 6, and root from-space pointer points to a $GcObject$ in words 100...103 of region 5. Then the root pointer has type $(GcPtr\ R\ 5)$, and from-space has four linear facts:

$$\begin{aligned} \text{Mem}(Addr(5, 100), GcFwd\ R\ 5), \dots & \\ \dots, \text{Mem}(Addr(5, 103), GcPtr\ R\ 5) & \end{aligned}$$

The collector copies the root object to to-space with load and store expressions; the root object goes into the front of the Cheney queue, at words 0...3 of region 6:

$$\begin{aligned} \text{Mem}(Addr(6, 0), GcFwd\ R\ 5), \dots & \\ \dots, \text{Mem}(Addr(6, 3), GcPtr\ R\ 5) & \end{aligned}$$

The types of the words in to-space are initially exactly the same as the types in from-space; the collector really does perform a blind copy from one-space to the other. Since each Mem fact is a mixture of region number 6 and region number 5, to-space's linear array takes care of the mismatch between the 6 and the 5:

$$\begin{aligned} (LArray\ Scan\ Alloc\ (\lambda I : \text{int}. & \\ \text{Mem}(Addr(E, I), REI(E-1))) \rangle & \end{aligned}$$

At this point, $Scan = 0$, $Alloc = 4$, $E = 6$, and $E-1 = 5$, while the type sequence $(R\ 6)$ is defined as:

$$((R\ 6\ 0) = (GcFwd\ R)), \dots, ((R\ 6\ 3) = (GcPtr\ R))$$

The collector constructs a $(GcObject\ R\ 6\ 0)$ with this sequence, which in turn is used to form a to-space pointer of type $(GcPtr\ R\ 6)$. In from-space, the root object's forwarding pointer has type $(GcFwd\ R\ 5)$, which is equivalent to $(GcPtr\ R\ 6)$, allowing the forwarding pointer to point to the to-space object.

Next, the collector starts to scan the Cheney queue, shifting words from the Cheney queue linear array to the finished object linear array. Since the finished object linear array uses facts of the form $\text{Mem}(Addr(6, I), R\ 6\ I\ 6)$ instead of $\text{Mem}(Addr(6, I), R\ 6\ I\ (6-1))$, the collector must figure out a way to convert each value of type $(R\ 6\ I\ 5)$ into a value of type $(R\ 6\ I\ 6)$:

- The collector sets the forwarding pointer to null with a store operation; it is easy to build a null pointer to any region number.
- For the float value, the collector merely observes that $\text{Mem}(Addr(6, 1), GcPrim\ \text{float}\ 5)$ is equivalent to $\text{Mem}(Addr(6, 1), GcPrim\ \text{float}\ 6)$, since $GcPrim$ ignores the region number argument. No loads or stores are necessary.

- For each of the two pointers, the collector must convert a $(GcPtr\ R5)$ to a $(GcPtr\ R6)$. This is exactly the problem that the collector solved when it copied the root object. Thus, the collector just makes a blind copy of whatever from-space objects the $(GcPtr\ R5)$ pointers point to.

The collector continues scanning and copying until the Cheney queue is empty. As it progresses, it must keep track of the types of the objects in the Cheney queue, so that it can repeatedly pop the front object off the queue and process it. The fact that the queue is full of objects, each of size 4 words, is stored in a nonlinear array of type:

$$(NArray\ (Scan \div 4)\ (Alloc \div 4)\ (\lambda\alpha : \text{int}.(GcObject\ R\ 6\ (4 * \alpha))))$$

$NArray$ is exactly like $LArray$, except that it is nonlinear and holds nonlinear size-0 elements. As was the case with mod , the abstract machine doesn't actually have a division operation, but it simplifies the presentation to use one. The real Clay implementation changes the definitions to make $Scan$ and $Alloc$ object indices rather than word offsets, so that it multiplies them by 4 to compute the word offset rather than dividing by 4 to compute the object index.

8. RELATED WORK

The previous section compared our collectors to some other typed garbage collectors. Particularly close to our work, although not as low-level, is Monnier et al's second region calculus and collector[16]. The region calculus extends ordinary regions with a type operator that the program updates linearly at run-time; the types of the object are filtered by this type operator. The linear control over the type operator is similar to the way our linear arrays control the types of words in a region.

Linear TAL [6] shares our goal of building a memory management system from low-level linear memory primitives. However, their approach uses copying to transform nonlinear data into linear data, eliminating all aliasing, while our approach maintains the aliasing present in the nonlinear data. Since purely linear data doesn't need a tracing garbage collector, their type system need not be complicated enough to implement type-safe collection. On the other hand, the data copying imposes a steep run-time cost.

Separation logics[14][21] share our goal of using techniques derived from linear logic to describe data structures. In [12], we add an expression to λ^{low} that deduces $I_1 \neq I_2$ if a value of type $\text{Mem}(I_1, \tau_1)$ and a value of type $\text{Mem}(I_2, \tau_2)$ exist simultaneously; this is similar to an axiom in [21], and we've used it to develop a usable linear circular list type.

Birkedal et al [5] use separation logic to prove the correctness of a simple Cheney queue copying collector. Like our approach, they describe the intermediate state of the collector with a linear invariant, partitioned into $0 \dots Scan - 1$, $Scan \dots Alloc - 1$, $Alloc \dots SIZE - 1$ portions for to-space. Unlike our approach, they also track the changing state in from-space. They have not yet integrated their verified collector with a richly typed programming language; this may require a separate proof that a well-typed expression's type is maintained after the collector reorganizes the heap.

Walker[1] explores a variant of separation logic that describes memory location adjacency, rather than treating locations as integers. It would be interesting to see how much

of our λ^{low} and Clay code could be rewritten with adjacency rules rather than integer arithmetic rules.

Igarashi and Kobayashi[13] developed a hybrid memory management strategy combining garbage collection and linear data. However, the collector was not written in a safe language; rather, the focus was on making sure that linear data deallocation and traditional garbage collection did not interfere with each other. Petersen et al[19] also assume an existing garbage collector, but use an ordered variation of linearity to expose the memory allocation process to high level programs in a safe way.

9. CONCLUSIONS AND FUTURE WORK

This paper has demonstrated that linear types, augmented with support for simple arithmetic types, coercion functions, and type sequences or delayed types, are sufficient to express a variety of type-safe data structures, ranging from simple lists to Cheney queues of nonlinear objects. Many issues remain. First, λ^{low} forces programs to contain many explicit coercions (pack, unpack, split, combine, type equality coercions, etc.) to set up simple loads and stores, making programming tedious. Second, our proof language is something of a hack, but it isn't yet clear what to replace it with. In particular, we're concerned that recursive types and type sequences will invalidate the strong normalization property of a typical proof language based on the Curry-Howard correspondence. On the other hand, if the nonlinear data structures from section 5 could be rewritten without recursive types, then maybe we could achieve growable regions with neither recursive types nor type sequences.

Acknowledgments

The authors would like to thank Fred Smith and the many anonymous reviewers for their comments on various drafts of this paper.

10. REFERENCES

- [1] Amal Ahmed and David Walker. The logical approach to stack typing. In *2003 ACM SIGPLAN Workshop on Types in Language Design and Implementation*, 2003.
- [2] Andrew W. Appel. Foundational proof-carrying code. In *Logic in Computer Science*, 2001.
- [3] Arthur I. Baars and S. Doaitse Swierstra. Typing dynamic typing. In *International Conference on Functional Programming*, 2002.
- [4] Henry G. Baker. Lively linear Lisp — 'Look Ma, no garbage!'. *ACM SIGPLAN Notices*, 27(9):89–98, 1992.
- [5] L. Birkedal, N. Torp-Smith, and J.C. Reynolds. Local reasoning about a copying garbage collector. In *Symposium on Principles of Programming Languages (POPL)*, 2004.
- [6] James Cheney and Greg Morrisett. A linearly typed assembly language. Technical report, Department of Computer Science, Cornell University.
- [7] Karl Crary and Joseph C. Vanderwaart. An expressive, scalable type theory for certified code. In *Proceedings of the seventh ACM SIGPLAN international conference on Functional programming*, pages 191–205. ACM Press, 2002.
- [8] Karl Crary, David Walker, and Greg Morrisett. Typed memory management in a calculus of capabilities. In

Proceedings of the 26th ACM SIGPLAN-SIGACT symposium on Principles of programming languages, pages 262–275. ACM Press, 1999.

- [9] Pascal Fradet and Daniel Le Metayer. Shape types. In *Proceedings of the 24th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 27–39. ACM Press, 1997.
- [10] N. Hamid, Z. Shao, V. Trifonov, S. Monnier, and Z. Ni. A syntactic approach to foundational proof carrying-code. In *Proc. Seventeenth Annual IEEE Symposium on Logic In Computer Science (LICS'02)*, 2002.
- [11] Robert Harper, Furio Honsell, and Gordon Plotkin. A framework for defining logics. *Journal of the ACM (JACM)*, 40(1):143–184, 1993.
- [12] Heng Huang, Lea Wittie, and Chris Hawblitzel. Formal properties of linear memory types. Technical Report TR2003-468, Dartmouth College, 2003.
- [13] A. Igarashi and N. Kobayashi. Garbage collection based on a linear type system, 2000.
- [14] Samin S. Ishtiaq and Peter W. O’Hearn. BI as an assertion language for mutable data structures. In *Symposium on Principles of Programming Languages*, pages 14–26, 2001.
- [15] Stefan Monnier, Bratin Saha, and Zhong Shao. Principled scavenging. In *Proceedings of the ACM SIGPLAN’01 conference on Programming language design and implementation*, pages 81–91. ACM Press, 2001.
- [16] Stefan Monnier and Zhong Shao. Typed regions. Technical Report YALEU/DCS/TR-1242, Department of Computer Science, Yale University, 2002.
- [17] Greg Morrisett, David Walker, Karl Crary, and Neal Glew. From system F to typed assembly language. In *ACM Transactions on Programming Languages and Systems (TOPLAS)*, volume 21, pages 527–568. ACM Press, 1999.
- [18] George C. Necula and Peter Lee. Safe kernel extensions without run-time checking. In *2nd Symposium on Operating Systems Design and Implementation (OSDI ’96), October 28–31, 1996. Seattle, WA*, pages 229–243, 1996.
- [19] L. Petersen, R. Harper, K. Crary, and F. Pfenning. A type theory for memory allocation and data layout. In *symposium on Principles of Programming Languages*, 2003.
- [20] William Pugh. The omega test: a fast and practical integer programming algorithm for dependence analysis. In *Proceedings of the 1991 ACM/IEEE conference on Supercomputing*, pages 4–13. ACM Press, 1991.
- [21] J. Reynolds. Separation logic: a logic for shared mutable data structures, 2002.
- [22] Z. Shao, B. Saha, V. Trifonov, and N. Papaspyrou. A type system for certified binaries. In *ACM Symposium on Principles of Programming Languages*, 2002.
- [23] Frederick Smith, David Walker, and Greg Morrisett. Alias types. In *In European Symposium on Programming*, 2000.
- [24] S.W. Smith and S.H. Weingart. Building a high-performance, programmable secure coprocessor.

In *Computer Networks (Special Issue on Computer Network Security)*, volume 31, pages 831–860, 1999.

- [25] David Teller and Zhong Shao. Algorithm-independent framework for verifying integer constraints. Technical Report YALEU/DCS/TR-1195, Department of Computer Science, Yale University, 2000.
- [26] Mads Tofte and Jean-Pierre Talpin. Region-based memory management. *Information and Computation*, 132(2):109–176, 1997.
- [27] P. Wadler. Linear types can change the world! In M. Broy and C. Jones, editors, *IFIP TC 2 Working Conference on Programming Concepts and Methods, Sea of Galilee, Israel*, pages 347–359. North Holland, 1990.
- [28] David Walker and Greg Morrisett. Alias types for recursive data structures. *Lecture Notes in Computer Science (volume 2071)*, 2001.
- [29] David Walker and Kevin Watkins. On regions and linear types (extended abstract). In *Proceedings of the sixth ACM SIGPLAN international conference on Functional programming*, pages 181–192. ACM Press, 2001.
- [30] Daniel C. Wang and Andrew W. Appel. Type-preserving garbage collectors. In *Proceedings of the 28th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 166–178. ACM Press, 2001.
- [31] Ed Wei. Using low level linear memory management for type-preserving mark-sweep garbage collector (undergraduate thesis). Technical Report TR2003-465, Dartmouth College, 2003.
- [32] Hongwei Xi and Frank Pfenning. Eliminating array bound checking through dependent types. In *Proceedings of the ACM SIGPLAN ’98 conference on Programming language design and implementation*, pages 249–257. ACM Press, 1998.

Appendix A: Formal Definition of λ^{low}

To save space, congruence rules are omitted here (they are included in [12], as are the type erasure rules). Because the proofs in [12] are based on type sequences rather than delayed types, we present type sequences here but omit delayed types. The rules follow notational conventions from Wadler[27], Walker et al[29], and [32].

abbreviations

$$\begin{aligned} \forall\alpha : K.\tau &\triangleq \forall\alpha : K; \text{true}.\tau \\ \exists\alpha : K.\tau &\triangleq \exists\alpha : K; \text{true}.\tau \\ \text{Know}(B) &\triangleq \exists\alpha : \text{bool}; B.\cdot \langle \rangle \\ \text{let } x = e_1 \text{ in } e_2 &\triangleq \text{let } \langle x \rangle = \wedge \langle e_1 \rangle \text{ in } e_2 \\ \tau_1 \xrightarrow{\phi} \tau_2 &\triangleq \tau_1 \xrightarrow{\phi, \infty} \tau_2 \end{aligned}$$

$$\text{linearity: } \quad \phi = \cdot \mid \wedge$$

$$\text{kinds: } \quad K = J \mid \text{int} \mid \text{bool} \quad J = \overset{\phi}{n} \mid K \rightarrow J$$

arithmetic

$$\begin{aligned} i &= \dots \mid -2 \mid -1 \mid 0 \mid 1 \mid 2 \mid \dots \quad b = \text{true} \mid \text{false} \\ I &= \alpha \mid i \mid I_1 \text{ iop } I_2 \quad B = \alpha \mid b \mid \neg B \mid B_1 \text{ bop } B_2 \mid I_1 \text{ cmp } I_2 \\ \text{iop} &= + \mid - \mid * \quad \text{bop} = \wedge \mid \vee \\ \text{cmp} &= < \mid > \mid \leq \mid \geq \mid = \mid \neq \\ \text{op} &= \text{iop} \mid \text{bop} \mid \text{cmp} \end{aligned}$$

types

$$\begin{aligned} \text{limit} &= I \mid \infty \\ \tau &= \tau_1 \xrightarrow{\phi, \text{limit}} \tau_2 \mid \phi(\vec{\tau}) \mid \alpha \mid I \mid B \mid \lambda \alpha : K. \tau \mid \forall \alpha : K; B. \tau \\ &\mid \exists \alpha : K; B. \tau \mid \tau_1 \tau_2 \mid \mu \alpha : K. \tau \mid \text{Int}(I) \mid \text{Bool}(B) \\ &\mid \text{if } B \text{ then } \tau_1 \text{ else } \tau_2 \mid \text{Mem}(I, \tau) \mid \text{Gen}(\tau, I) \mid \text{Eq}(\tau_1, \tau_2) \\ &\mid F^K \mid \text{InDomain}(I, \tau) \end{aligned}$$

expressions

$$\begin{aligned} e &= i \mid b \mid x \mid e_1 e_2 \mid e \tau \mid \phi(\vec{e}) \mid \lambda x : \tau \xrightarrow{\phi, \text{limit}} e \mid e_1 \text{ op } e_2 \mid \neg e \\ &\mid \Lambda \alpha : K; B. v \mid \text{let } (\vec{x}) = e_1 \text{ in } e_2 \\ &\mid \text{if } e_1 \text{ then } e_2 \text{ else } e_3 \mid \text{if } B \text{ then } e_1 \text{ else } e_2 \\ &\mid \text{pack}[\tau_1, e] \text{ as } \exists \alpha : K; B. \tau_2 \mid \text{unpack } \alpha, x = e_1 \text{ in } e_2 \\ &\mid \text{roll}[(\mu \alpha : K. \tau_0) \tau_1 \cdots \tau_n](e) \mid \text{unroll}(e) \mid \text{fix } x : \tau. v \\ &\mid \text{load}(e_{\text{ptr}}, e_{\text{Mem}}) \mid \text{store}(e_{\text{ptr}}, e_{\text{Mem}}, e_v) \mid \text{coerce}(e) \\ &\mid \text{make_eq}(\tau) \mid \text{apply_eq}(\tau, e_1, e_2) \mid \text{new_seq}(K) \mid \text{discard_seq}(e) \\ &\mid \text{define_seq}(e, \tau) \mid \text{in_domain}(I_1, I_2, e_1, e_2) \mid \text{fact} \end{aligned}$$

values

$$\begin{aligned} v &= i \mid b \mid \Lambda \alpha : K; B. v \mid \text{pack}[\tau_1, v] \text{ as } \exists \alpha : K; B. \tau_2 \\ &\mid \text{roll}[(\mu \alpha : K. \tau_0) \tau_1 \cdots \tau_n](v) \mid \lambda x : \tau \xrightarrow{\phi, \text{limit}} e \\ &\mid \phi(\vec{v}) \mid \text{fact} \end{aligned}$$

environments

$$\begin{aligned} M &= \{1 \mapsto v_1, \dots, n \mapsto v_n\} \\ C &= \Psi; \Phi; \Delta; \Gamma; B; \text{limit} \\ \Psi &= \{1 \mapsto \tau_1, \dots, n \mapsto \tau_n\} \\ \Phi &= \{F_1^{K_1} \xrightarrow{\phi_1} \delta_1, \dots, F_n^{K_n} \xrightarrow{\phi_n} \delta_n\} \\ \delta &= \{0 \mapsto \tau_0, \dots, n \mapsto \tau_n\} \\ \Delta &= \{\alpha_1 \mapsto K_1, \dots, \alpha_n \mapsto K_n\} \\ \Gamma &= \{x_1 \mapsto \tau_1 \cdots, x_n \mapsto \tau_n\} \end{aligned}$$

judgments

$$\begin{aligned} \Phi; \Delta \vdash \tau : K \\ B \vdash B_1 \doteq B_2 \\ B \vdash I_1 \doteq I_2 \\ \Phi; B \vdash \tau_1 \equiv \tau_2 \\ C \vdash e : \tau \\ C \vdash (M, e : \tau) \\ (M, e) \rightarrow (M', e') \end{aligned}$$

evaluation rules

$$\begin{aligned} (M, \text{load}(i, \text{fact})) &\rightarrow (M, \wedge \langle M(i), \text{fact} \rangle) \\ (M, \text{store}(i, \text{fact}, v)) &\rightarrow ([i \mapsto v]M, \text{fact}) \end{aligned}$$

$$\frac{e \rightarrow e'}{(M, e) \rightarrow (M, e')}$$

$$(\lambda x : \tau \xrightarrow{\phi, I} e_1) v_2 \rightarrow \text{coerce}([x \mapsto v_2]e_1)$$

$$(\lambda x : \tau \xrightarrow{\phi, \infty} e_1) v_2 \rightarrow [x \mapsto v_2]e_1$$

$$\text{coerce}(v) \rightarrow v$$

$$(\Lambda \alpha : K; B. v) \tau \rightarrow [\alpha \mapsto \tau]v$$

$$v_1 \text{ op } v_2 \rightarrow \text{simplify}(v_1 \text{ op } v_2)$$

$$\neg v \rightarrow \text{simplify}(\neg v)$$

$$\text{if } \text{true} \text{ then } e_1 \text{ else } e_2 \rightarrow e_1$$

$$\text{if } \text{false} \text{ then } e_1 \text{ else } e_2 \rightarrow e_2$$

$$\frac{\vdash B \doteq \text{true}}{\text{if } B \text{ then } e_1 \text{ else } e_2 \rightarrow e_1}$$

$$\frac{\vdash B \doteq \text{false}}{\text{if } B \text{ then } e_1 \text{ else } e_2 \rightarrow e_2}$$

$$\text{unroll}(\text{roll}[\tau](v)) \rightarrow v$$

$$\text{fix } x : \tau. v \rightarrow [x \mapsto \text{fix } x : \tau. v]v$$

$$\text{make_eq}(\tau) \rightarrow \text{fact}$$

$$\text{apply_eq}(\tau, \text{fact}, v) \rightarrow v$$

$$\text{discard_seq}(\text{fact}) \rightarrow \cdot \langle \rangle$$

$$\text{define_seq}(\text{fact}, \tau) \rightarrow \wedge \langle \text{fact}, \text{fact}, \text{fact} \rangle$$

$$\text{let } \langle x_1, \dots, x_n \rangle = \phi \langle v_1, \dots, v_n \rangle \text{ in } e \rightarrow [x_1 \mapsto v_1, \dots, x_n \mapsto v_n]e$$

$$\text{unpack } \alpha, x = (\text{pack}[\tau_1, v_1] \text{ as } \tau_2) \text{ in } e_2 \rightarrow [\alpha \mapsto \tau_1, x \mapsto v_1]e_2$$

$$\text{in_domain}(I_1, I_2, \text{fact}, \text{fact}) \rightarrow \wedge \langle \text{know}(0 \leq I_1 \wedge I_1 < I_2), \text{fact} \rangle$$

$$\text{new_seq}(K) \rightarrow \text{pack}[F^K, \text{fact}] \text{ as } \exists \alpha : \text{int} \rightarrow K. \text{FunGen}(\alpha, 0), \text{ where } F^K \text{ is fresh}$$

kinding rules

$$\Phi; \Delta \vdash i : \text{int}$$

$$\Phi; \Delta, \alpha : K \vdash \alpha : K$$

$$\Phi; \Delta \vdash b : \text{bool}$$

$$\frac{\Phi; \Delta \vdash I_1 : \text{int} \quad \Phi; \Delta \vdash I_2 : \text{int}}{\Phi; \Delta \vdash I_1 \text{ iop } I_2 : \text{int}}$$

$$\frac{\Phi; \Delta \vdash I_1 : \text{int} \quad \Phi; \Delta \vdash I_2 : \text{int}}{\Phi; \Delta \vdash I_1 \text{ cmp } I_2 : \text{bool}}$$

$$\frac{\Phi; \Delta \vdash B_1 : \text{bool} \quad \Phi; \Delta \vdash B_2 : \text{bool}}{\Phi; \Delta \vdash B_1 \text{ bop } B_2 : \text{bool}}$$

$$\frac{\Phi; \Delta \vdash B : \text{bool}}{\Phi; \Delta \vdash \neg B : \text{bool}}$$

$$\frac{\Phi; \Delta, \alpha : K_a \vdash \tau : J_b}{\Phi; \Delta \vdash \lambda \alpha : K_a. \tau : K_a \rightarrow J_b}$$

$$\frac{\Phi; \Delta \vdash \tau_f : K_a \rightarrow J_b \quad \Phi; \Delta \vdash \tau_a : K_a}{\Phi; \Delta \vdash \tau_f \tau_a : J_b}$$

$$\frac{\Phi; \Delta, \alpha : K \vdash B : \text{bool} \quad \Phi; \Delta, \alpha : K \vdash \tau : \overset{\phi}{n}}{\Phi; \Delta \vdash \forall \alpha : K; B. \tau : \overset{\phi}{n}}$$

$$\frac{\Phi; \Delta, \alpha : K \vdash B : \text{bool} \quad \Phi; \Delta, \alpha : K \vdash \tau : \overset{\phi}{n}}{\Phi; \Delta \vdash \exists \alpha : K; B. \tau : \overset{\phi}{n}}$$

$$\frac{\Phi; \Delta, \alpha : K \vdash \tau : K}{\Phi; \Delta \vdash \mu \alpha : K. \tau : K}$$

$$\frac{(\Phi; \Delta \vdash \text{limit} : \text{int} \wedge n = 0) \text{ or } (\text{limit} = \infty \wedge n = 1)}{\Phi; \Delta \vdash \tau_1 \xrightarrow{\phi, \text{limit}} \tau_2 : \overset{\phi}{n}}$$

$$\frac{\forall j. (\Phi; \Delta \vdash \tau_j : n_j) \quad n = \sum_j n_j}{\Phi; \Delta \vdash \cdot \langle \vec{\tau} \rangle : \overset{\phi}{n}}$$

type checking rules

$$\frac{\forall j. (\Phi; \Delta \vdash \tau_j : \overset{\phi_j}{n_j}) \quad n = \sum_j n_j}{\Phi; \Delta \vdash \wedge \langle \vec{\tau} \rangle : \hat{n}}$$

$$\frac{\Phi; \Delta \vdash I : \text{int}}{\Phi; \Delta \vdash \text{Int}(I) : \dot{1}}$$

$$\frac{\Phi; \Delta \vdash B : \text{bool}}{\Phi; \Delta \vdash \text{Bool}(B) : \dot{1}}$$

$$\frac{\Phi; \Delta \vdash B : \text{bool} \quad \Phi; \Delta \vdash \tau_1 : J \quad \Phi; \Delta \vdash \tau_2 : J}{\Phi; \Delta \vdash \text{if } B \text{ then } \tau_1 \text{ else } \tau_2 : J}$$

$$\frac{\Phi; \Delta \vdash I : \text{int} \quad \Phi; \Delta \vdash \tau : \dot{1}}{\Phi; \Delta \vdash \text{Mem}(I, \tau) : \hat{0}}$$

$$\frac{\Phi; \Delta \vdash I : \text{int} \quad \Phi; \Delta \vdash \tau : \text{int} \rightarrow J}{\Phi; \Delta \vdash \text{Gen}(\tau, I) : \hat{0}}$$

$$\frac{\Phi; \Delta \vdash I : \text{int} \quad \Phi; \Delta \vdash \tau : \text{int} \rightarrow J}{\Phi; \Delta \vdash \text{InDomain}(I, \tau) : \hat{0}}$$

$$\frac{\Phi; \Delta \vdash \tau_1 : K \quad \Phi; \Delta \vdash \tau_2 : K}{\Phi; \Delta \vdash \text{Eq}(\tau_1, \tau_2) : \hat{0}}$$

$$\Phi; F^K \mapsto \delta; \Delta \vdash F^K : \text{int} \rightarrow K$$

type equivalence rules

$$\Phi; B \vdash \tau \equiv \tau$$

$$\frac{\Phi; B \vdash \tau_1 \equiv \tau_2}{\Phi; B \vdash \tau_2 \equiv \tau_1}$$

$$\frac{\Phi; B \vdash \tau_1 \equiv \tau_2 \quad \Phi; B \vdash \tau_2 \equiv \tau_3}{\Phi; B \vdash \tau_1 \equiv \tau_3}$$

$$\Phi; B \vdash (\lambda \alpha : K. \tau_b) \tau_a \equiv [\alpha \mapsto \tau_a] \tau_b$$

$$\frac{B \vdash I \doteq i}{\Phi; F^K \mapsto \text{fun}; B \vdash F^K I \equiv \text{fun}(i)}$$

$$\frac{B \vdash I_1 \doteq I_2}{\Phi; B \vdash I_1 \equiv I_2}$$

$$\frac{B \vdash B_1 \doteq B_2}{\Phi; B \vdash B_1 \equiv B_2}$$

$$\Phi; B \vdash \text{Eq}(\tau_1, \tau_2) \equiv \text{Eq}(\tau_2, \tau_1)$$

$$\Phi; B \vdash \text{if true then } \tau_1 \text{ else } \tau_2 \equiv \tau_1$$

$$\Phi; B \vdash \text{if false then } \tau_1 \text{ else } \tau_2 \equiv \tau_2$$

$$\frac{\Psi = \Psi_{\text{spare}}, \Psi_e \quad \Phi = \dot{\Phi}_{\text{spare}}, \Phi_e \quad \forall i \in \text{dom}(\Psi). (\emptyset; \dot{\Phi}; \emptyset; \emptyset; \text{true}; \infty \vdash M(i) : \Psi(i)) \quad \Psi_e; \Phi_e; \Delta; \Gamma; B; \text{limit} \vdash e : \tau}{\Psi; \Phi; \Delta; \Gamma; B; \text{limit} \vdash (M, e : \tau)}$$

$$\dot{C}, x : \tau \vdash x : \tau$$

$$\frac{C, \alpha : K, B \vdash v : \tau}{C \vdash \Lambda \alpha : K; B.v : \forall \alpha : K; B.\tau}$$

$$C \vdash e : \forall \alpha : K; B.\tau_1$$

$$C \vdash \tau_2 : K$$

$$C \vdash [\alpha \mapsto \tau_2] B$$

$$C \vdash e \tau_2 : [\alpha \mapsto \tau_2] \tau_1$$

$$\Psi; \Phi; \Delta; \Gamma; B; I_2 \vdash e : \tau$$

$$\frac{(\text{limit}_1 = \infty \wedge \tau : \hat{0}) \text{ or } (\text{limit}_1 = I_1 \wedge B \vdash I_1 > I_2)}{\Psi; \Phi; \Delta; \Gamma; B; \text{limit}_1 \vdash \text{coerce}(e) : \tau}$$

$$C \vdash e : \tau_1 \quad C \vdash \tau_1 \equiv \tau_2$$

$$\Phi; \Delta \vdash \tau_1 : K_1 \quad \Phi; \Delta \vdash \tau_2 : K_1$$

$$C \vdash e : \tau_2$$

$$C = C_1, \dots, C_n$$

$$\forall i. (C_i \vdash e_i : \tau_i) \quad C \vdash \phi(\vec{\tau}) : K$$

$$C \vdash \phi(\vec{e}) : \phi(\vec{\tau})$$

$$\frac{\overset{\phi}{\Phi}; \Delta \vdash \tau_1 : K \quad \overset{\phi}{\Psi}; \overset{\phi}{\Phi}; \Delta; \overset{\phi}{\Gamma}, x : \tau_1; B; \text{limit}_2 \vdash e : \tau_2 \quad (\text{limit}_2 = \infty) \text{ or } (\overset{\phi}{\Phi}; \Delta \vdash \text{limit}_2 : \text{int} \quad B \vdash \text{limit}_2 \geq 0)}{\overset{\phi}{\Psi}; \overset{\phi}{\Phi}; \Delta; \overset{\phi}{\Gamma}; B; \text{limit}_1 \vdash \lambda x : \tau_1 \overset{\phi, \text{limit}_2}{\longrightarrow} e : \tau_1 \overset{\phi, \text{limit}_2}{\longrightarrow} \tau_2}$$

$$C_1, C_2 = \Psi; \Phi; \Delta; \Gamma; B; \text{limit}_C$$

$$C_1 \vdash e_1 : \tau_a \overset{\phi, \text{limit}_f}{\longrightarrow} \tau_b \quad C_2 \vdash e_2 : \tau_a$$

$$(\text{limit}_C = \text{limit}_f = \infty) \text{ or } (B \vdash \text{limit}_f < \text{limit}_C)$$

$$\text{or } (\text{limit}_C = \infty, \text{limit}_f = I, \overset{\phi_b}{\Phi}; \Delta \vdash \tau_b : \hat{0})$$

$$C_1, C_2 \vdash e_1 e_2 : \tau_b$$

$$C_a \vdash e_a : \phi(\vec{\tau}) \quad C_b, x : \vec{\tau} \vdash e_b : \tau_b$$

$$C_a, C_b \vdash \text{let} \langle \vec{x} \rangle = e_a \text{ in } e_b : \tau_b$$

$$C \vdash \tau : \hat{n}$$

$$C, x : \tau \vdash v : \tau$$

$$C \vdash (\text{fix } x : \tau.v) : \tau$$

$$C \vdash \tau_1 : \hat{n} \quad C \vdash \exists \alpha : K; B.\tau_2 : \hat{n}$$

$$C \vdash e : [\alpha \mapsto \tau_1] \tau_2 \quad C \vdash [\alpha \mapsto \tau_1] B$$

$$C \vdash \text{pack}[\tau_1, e] \text{ as } \exists \alpha : K; B.\tau_2 : \exists \alpha : K; B.\tau_2$$

$$C_1 \vdash e_1 : \exists \alpha : K; B.\tau_1$$

$$C_2, \alpha : K, x : \tau_1, B \vdash e_2 : \tau_2$$

$$C_1, C_2 \vdash \text{unpack } \alpha, x = e_1 \text{ in } e_2 : \tau_2$$

$$\frac{\begin{array}{c} \tau = (\mu\alpha : K.\tau_0)\tau_1 \cdots \tau_n \\ C \vdash \tau : K \\ C \vdash e : ([\alpha \mapsto \mu\alpha : K.\tau_0]\tau_0)\tau_1 \cdots \tau_n \end{array}}{C \vdash \text{roll}[\tau](e) : \tau}$$

$$\frac{\begin{array}{c} \tau = (\mu\alpha : K.\tau_0)\tau_1 \cdots \tau_n \\ C \vdash \tau : K \quad C \vdash e : \tau \end{array}}{C \vdash \text{unroll}(e) : ([\alpha \mapsto \mu\alpha : K.\tau_0]\tau_0)\tau_1 \cdots \tau_n}$$

$$\dot{C}, i \mapsto \tau \vdash \text{fact} : \text{Mem}(i, \tau)$$

$$\dot{C}, F^K \hat{\mapsto} \text{fun} \vdash \text{fact} : \text{Gen}(F^K, i)$$

$$\dot{C} \vdash \text{fact} : \text{Eq}(\tau, \tau)$$

$$\frac{i \in \text{dom}(\delta)}{\dot{C}, F^K \mapsto \delta \vdash \text{fact} : \text{InDomain}(i, F^K)}$$

$$\frac{C_1 \vdash e_{\text{ptr}} : \text{Int}(I) \quad C_2 \vdash e_{\text{Mem}} : \text{Mem}(I, \tau)}{C_1, C_2 \vdash \text{load}(e_{\text{ptr}}, e_{\text{Mem}}) : \wedge \langle \tau, \text{Mem}(I, \tau) \rangle}$$

$$\frac{\begin{array}{c} C = C_1, C_2, C_3 = \Psi; \Phi; \Delta; \Gamma; B; \infty \\ C_2 \vdash e_{\text{Mem}} : \text{Mem}(I, \tau_1) \quad C_3 \vdash e_v : \tau_2 \\ C_1 \vdash e_{\text{ptr}} : \text{Int}(I) \quad C \vdash \tau_2 : \dot{1} \end{array}}{C \vdash \text{store}(e_{\text{ptr}}, e_{\text{Mem}}, e_v) : \text{Mem}(I, \tau_2)}$$

$$\frac{C \vdash e : \text{Gen}(\tau_f, I) \quad C \vdash \tau_f : \text{int} \rightarrow J \quad C \vdash \tau_a : J}{C \vdash \text{define_seq}(e, \tau_a) : \wedge \langle \text{Gen}(\tau_f, I + 1), \text{Eq}(\tau_f I, \tau_a), \text{InDomain}(I, \tau_f) \rangle}$$

$$\frac{\begin{array}{c} C_1, C_2 \vdash I_1 : \text{int} \quad C_1, C_2 \vdash I_2 : \text{int} \\ C_1 \vdash e_1 : \text{InDomain}(I_1, \tau_f) \quad C_2 \vdash e_2 : \text{Gen}(\tau_f, I_2) \end{array}}{C_1, C_2 \vdash \text{in_domain}(I_1, I_2, e_1, e_2) : \wedge \langle \text{Know}(0 \leq I_1 \wedge I_1 < I_2), \text{Gen}(\tau_f, I_2) \rangle}$$

$$\frac{\dot{C} \vdash \tau : K}{\dot{C} \vdash \text{make_eq}(\tau) : \text{Eq}(\tau, \tau)}$$

$$\frac{\begin{array}{c} C_1, C_2 \vdash \tau_f : K \rightarrow J \quad C_1 \vdash e_{\text{eq}} : \text{Eq}(\tau_a, \tau_b) \\ C_1, C_2 \vdash \tau_a : K \quad C_1, C_2 \vdash \tau_b : K \quad C_2 \vdash e_f : \tau_f \tau_a \end{array}}{C_1, C_2 \vdash \text{apply_eq}(\tau_f, e_{\text{eq}}, e_f) : \tau_f \tau_b}$$

$$\dot{C} \vdash \text{new_seq}(J) : \exists \alpha : \text{int} \rightarrow J.\text{Gen}(\alpha, 0)$$

$$\frac{C \vdash e : \text{Gen}(\tau, I)}{C \vdash \text{discard_seq}(e) : \cdot \langle \rangle}$$

$$\dot{C} \vdash i : \text{Int}(i)$$

$$\dot{C} \vdash b : \text{Bool}(b)$$

$$\frac{C_1 \vdash e_1 : \text{Int}(I_1) \quad C_2 \vdash e_2 : \text{Int}(I_2)}{C_1, C_2 \vdash e_1 \text{ iop } e_2 : \text{Int}(I_1 \text{ iop } I_2)}$$

$$\frac{C_1 \vdash e_1 : \text{Int}(I_1) \quad C_2 \vdash e_2 : \text{Int}(I_2)}{C_1, C_2 \vdash e_1 \text{ cmp } e_2 : \text{Bool}(I_1 \text{ cmp } I_2)}$$

$$\frac{C_1 \vdash e_1 : \text{Bool}(B_1) \quad C_2 \vdash e_2 : \text{Bool}(B_2)}{C_1, C_2 \vdash e_1 \text{ bop } e_2 : \text{Bool}(B_1 \text{ bop } B_2)}$$

$$\frac{C \vdash e : \text{Bool}(B)}{C \vdash \neg e : \text{Bool}(\neg B)}$$

$$\frac{C_a \vdash e_1 : \text{Bool}(B) \quad C_b, B \vdash e_2 : \tau \quad C_b, \neg B \vdash e_3 : \tau}{C_a, C_b \vdash \text{if } e_1 \text{ then } e_2 \text{ else } e_3 : \tau}$$

$$\frac{\begin{array}{c} C = \Psi; \Phi; \Delta; \Gamma; B_C; I \\ C \vdash B : \text{bool} \quad C, B \vdash e_1 : \tau \quad C, \neg B \vdash e_2 : \tau \end{array}}{C \vdash \text{if } B \text{ then } e_1 \text{ else } e_2 : \tau}$$