

# Safe Transferable Regions

**Gowtham Kaki**

Purdue University, USA<sup>1</sup>  
[gkaki@purdue.edu](mailto:gkaki@purdue.edu)

**G. Ramalingam**

Microsoft Research, India  
[grama@microsoft.com](mailto:grama@microsoft.com)

---

## Abstract

There is an increasing interest in alternative memory management schemes that seek to combine the convenience of garbage collection and the performance of manual memory management in a single language framework. Unfortunately, ensuring safety in presence of manual memory management remains as great a challenge as ever. In this paper, we present a C#-like object-oriented language called BROOM that uses a combination of region type system and lightweight runtime checks to enforce safety in presence of user-managed memory regions called *transferable regions*. Unsafe transferable regions have been previously used to contain the latency due to unbounded GC pauses. Our approach shows that it is possible to restore safety without compromising on the benefits of transferable regions. We prove the type safety of BROOM in a formal framework that includes its C#-inspired features, such as higher-order functions and generics. We complement our type system with a type inference algorithm, which eliminates the need for programmers to write region annotations on types. The inference algorithm has been proven sound and relatively complete. We describe a prototype implementation of the inference algorithm, and our experience of using it to enforce memory safety in dataflow programs.

**2012 ACM Subject Classification** Software and its engineering → Allocation / deallocation strategies; Software and its engineering → Object oriented languages; Software and its engineering → Data flow languages; Software and its engineering → Software verification and validation;

**Keywords and phrases** Memory Safety, Formal Methods, Type System, Type Inference, Regions, Featherweight Java

**Digital Object Identifier** [10.4230/LIPIcs.ECOOP.2018.11](https://doi.org/10.4230/LIPIcs.ECOOP.2018.11)

**Acknowledgements** We are grateful to Kapil Vaswani, Dimitrios Vytiniotis, Michael Isard, and Steven Hand for their valuable inputs during the initial stages of this project. We would also like to thank multiple anonymous reviewers for their careful scrutiny and feedback on various versions of this work.

## 1 Introduction

Computations performed by a network of concurrent communicating actors often involve data transfer between producers and consumers. Consider for example the `SELECT` query operator shown in Fig. 1, which functions as an actor in a dataflow computation involving a network of other such query operators. `SELECT` receives a stream of input messages, each associated with a time window  $t$ , processed by method `onReceive`. Each input message contains a list of inputs, each processed by applying a user-defined function to create a

---

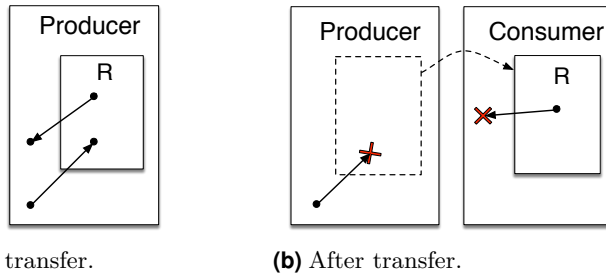
<sup>1</sup>Work done during an internship at Microsoft Research, India.



## 11:2 Safe Transferable Regions

```
1 class SelectVertex<TIn, TOut> {
2   Func<TIn, TOut> selector;
3   Dictionary<Time, List<TOut>> map;
4   ...
5   void onReceive(Time t, List<TIn> inList) {
6     if (!map.ContainsKey(t)) map[t] = new List<TOut>();
7     foreach (TIn input in inList) {
8       TOut output = selector(input);
9       map[t].add(output); } }
10  void onNotify(Time t) {
11    List<TOut> outList = map[t];
12    map.Remove(t);
13    transfer(successorId, t, outList); }
14 }
```

■ **Figure 1** SELECT dataflow operator



■ **Figure 2** References in and out of a transferable region ( $R$ ) become invalid after transfer.

corresponding output. Multiple messages with the same timestamp are permitted and messages with different timestamps may arrive out of order. An invocation of method `OnNotify` indicates that no more input messages with a timestamp  $t$  will be subsequently delivered. At this point, the operator completes the processing for time window  $t$  and sends a corresponding output message to its successor.

**Performance.** When the transferred data structures are large, as is the case with many streaming big-data analysis systems, garbage collection overhead becomes significant [7]. Further, in a distributed dataflow system, the GC pause at one node can have a cascading adverse effect on the performance of other nodes [7, 14]: a GC pause at an upstream actor can block downstream actors that are waiting for messages. However, much of the GC overhead results from the collector performing avoidable or unproductive work. For the example in Fig. 1, GC might repeatedly traverse the `map`, although its objects cannot be collected until a suitable timing message arrives. There has been increasing interest in off-heap memory management for better performance in the context of several systems and languages (e.g., Spark [25], Scala [18], Rust). Several systems (e.g., [3]) resort to using “free object pools” to partially alleviate this problem, in the absence of language support.

**Safety.** Several of these systems are designed to allow the producer and consumer to execute in the same address space or in different address spaces (as determined by the compiler and runtime). We wish to ensure memory safety in the presence of such transferred data. A transfer operation, with no additional checks, may cause memory safety violations, both at the producer of the data structure, and its (possibly remote) consumer. At the producer, any existing references into the transferred data structure become invalid post transfer. If the data structure contains references to objects outside (the transferred data), then such references become invalid in the context of the consumer. Both scenarios are depicted in Fig. 2. While references from within the transferable data to outside can be

disallowed in the interest of safety, references into transferable data needs to be allowed before the data is transferred. Allowing such references is crucial, as any non-trivial program creates temporary references to the internal objects of a data-structure.

**Safe Transferable Regions.** In this paper, we present a language-based approach to cleanly encapsulate transferable data and a safe and efficient implementation of such data based on regions. A region is a block of memory that is allocated and freed in one shot, often in constant time. A region may contain one or more contiguous range of memory locations, and individual objects may be dynamically allocated within the region over time, while they are deallocated en masse when the region is freed. Thus, a region is a good fit for a transferable data-structure. In Fig. 1, the output to be constructed for each time window  $t$  (i.e., `map[τ]`) can be a separate region that is allocated when the first message with timestamp  $t$  arrives, and deallocated after `map[τ]` is transferred in `onNotify`.

The use of regions alleviates the performance concern related to garbage collection described earlier. (It also enables more efficient serialization/deserialization of data-structures across address spaces, which is also a significant performance bottleneck in such systems.) However, manual region memory management introduces memory safety issues such as dangling pointers. We present a single type system that simultaneously addresses this safety concern, as well as those described above related to the use of *transferable data*.

Safe region-based memory management using types was pioneered by Tofte and Talpin [20, 21], who use only lexically scoped regions. At runtime, the set of all regions (existing at a point in time) forms a stack. Thus, the lifetimes of all regions must be well-nested: it is not possible to have two regions whose lifetimes overlap, with neither one's lifetime contained within the other. Unfortunately, the data structures in the above example do not satisfy this restriction (as the output messages for multiple time windows may be simultaneously live, without any containment relation between their lifetimes). We refer to regions with lexically scoped lifetimes as *stack regions* and to regions that do not have such a lexically scoped region as *dynamic regions*.

Our focus, in this paper, is on first-class *memory-safe* dynamic regions that can be safely transferred across address spaces. We refer to such dynamic regions as *transferable* regions. Our approach is based on a variation of ideas introduced by [23, 10] that combine linear types and regions to support dynamic regions. Unlike the prior work, we ensure memory safety in the presence of transferable regions through a *combination of a static typing discipline and lightweight runtime checks in place of linear types*.

**Language.** The cornerstone of this approach is an `open` lexical block for transferable regions, that “opens” a transferable region and guarantees that the region won't be transferred/freed while it is open. By nesting a Tofte-Talpin style `letregion` lexical block, that delimits the lifetime of a stack region, inside an `open` lexical block for a transferable region, we can guarantee that the transferable region will remain live as long as the stack region is live. We say that the former *outlives* the latter, and any references from the stack region to the transferable region are therefore safe. This is particularly useful, since such stack regions serve as temporary working storage while working with the (opened) transferable regions.

By controlling the outlives relationships between various regions, we only allow safe cross-region references, while prohibiting unsafe ones. In the above example, an outlives relationship *from* the stack region *to* the transferable region means that the references in that direction are allowed, but not the references in the opposite direction. In contrast, if an `open` block of a transferable region  $R_0$  is nested inside an `open` block of another transferable region  $R_1$ , we do not establish any outlives relationships, thus declaring our intention to not allow any cross-region references between  $R_0$  and  $R_1$ . Finally, we observe that outlives

relationships are established based on the lexical structure of the program, hence a static type system can enforce them effectively. By assigning region types to objects, which capture the regions such objects are allocated in, and by maintaining outlives relationships between various regions, we can statically decide the safety of all references in the program.

**Type System.** Formally, our type system introduces *region parameters* (for both classes and methods) and uses constrained parametric polymorphism over these parameters, where the constraints capture outlives constraints between the region parameters. The type system may be seen as a form of ownership type system [5], with a region being the owner of all objects allocated in that region.

**Lightweight Runtime Checks.** Ensuring memory safety using this approach requires ensuring that the use of transferable regions satisfies certain temporal properties. Firstly, a transferable region should not be transferred/freed inside an open block of that region (i.e., while it is still open). Secondly, a transferred/freed region should not be opened. These are tpestate invariants on the transferable region objects, which are hard to enforce statically in the presence of unrestricted aliasing. Techniques like linear types and unique pointers can be used to restrict aliasing, but the constraints they impose are often hard to program around. We therefore enforce tpestate invariants at runtime via lightweight checks. In particular, we define an acceptable state transition discipline for transferable regions (Fig. 4), and check, at runtime, whether a given transition of a transferable region (*e.g.*, from *open* state to *freed* state) is valid or not. The check is lightweight since it only involves checking a single tag that captures the current state. We believe that this is a reasonable choice since regions are coarse-grained objects manipulated infrequently, when compared to the fine-grained objects that are present inside these regions, for which safety is enforced statically.

**Type Inference.** One of the key contributions of this paper is a type inference algorithm that eliminates the need for users to write region type annotations. The users write programs in the underlying language that provides primitives for the users to manipulate regions (create, open, free, transfer) and to allocate objects in regions, but has no region type annotations.

Our inference algorithm proceeds in three stages: in the first stage, we elaborate the program by introducing region parameters (for classes and methods); in the second stage we generate a set of constraints that must be satisfied for the program to type check; in the third stage, we solve the set of constraints, inferring the preconditions of all classes and methods (in terms of the expected outlives-constraints between the region parameters). We show that the algorithm is sound, and the stages of constraint generation and solving are complete (i.e., the algorithm is complete relative to the first stage of elaboration).

**Evaluation.** Our work was inspired by [7], which presents evidence that realistic programs can be written using transferable regions and that this can yield significant performance improvements. While [7] addresses the engineering challenges in extending a managed runtime with transferable regions, it adopts an ad hoc approach insofar as language design is concerned, exposing the region functionality through an unsafe API, and in the process losing the safety guarantees. Our contribution is an alternative approach that is grounded in sound theory and restores the language safety guarantees. The utility of our approach is demonstrated by a prototype implementation of our type inference algorithm that was able to identify unsafe memory accesses among the benchmarks extracted from [7].

**Contributions.** The paper makes the following contributions:

- We present BROOM, a C#-like typed object-oriented language that supports programmer-managed memory regions. BROOM extends its core language, which includes *lambdas* (higher-order functions) and *generics* (parametric polymorphism), with constructs to create, manage and free static and transferable memory regions. Transferable regions are

first-class values in BROOM.

- BROOM is equipped with a region type system that statically guarantees safety of all memory accesses in a well-typed program, provided that certain typestate invariants on regions hold. The latter invariants are enforced via simple runtime checks.
- We define an operational semantics for BROOM, and a type safety result that clearly defines and proves safety guarantees described above.
- We describe a region type inference algorithm for BROOM that (a). completely eliminates the need to annotate BROOM programs with region types, and (b). enables seamless interoperability between region-aware BROOM programs and legacy standard library code that is region-oblivious. Our inference algorithm is based on a novel constraint solver that performs abduction in a partial-order constraint domain to infer weakest solutions to recursive constraints.
- We establish the soundness and relative completeness of the type inference algorithm.
- We describe an implementation of BROOM frontend in OCaml, along with case studies where the region type system was able to identify unsafe memory accesses statically.

## 2 An Informal Overview of BROOM

BROOM enriches a simple object-oriented language (supporting parametric polymorphism and lambdas) with a set of region-specific constructs. In this section, we present an informal overview of these region-specific constructs.

### 2.1 Using Regions in BROOM

**Stack Regions** The “`letregion R { S }`” construct creates a new stack region, with a static identifier `R`, whose scope is restricted to the statement `S`. The semantics of `letregion` is similar to Tofte and Talpin [20]’s `letregion` expression: objects can be allocated by `S` in the newly created region while `R` is in scope, but the region and all objects allocated within it are freed at the end of `S`.

**Object Allocation** The “`new@R T ()`” construct creates a new object of type `T` in the region `R`. The specification of the allocation region `R` in this construct is optional. At runtime, BROOM maintains a stack of *active* regions, and we refer to the region at the top of the stack as the *allocation context*. The statement `new T ()` allocates the newly created object in the current allocation context. This is important as it enables BROOM applications to use existing region-oblivious `C#` libraries, as explained soon.

**Transferable Regions** Transferable regions are first class values of BROOM: they are objects of the class `Region`, they are created using the `new` keyword, and can be passed as arguments, stored in data structures, and returned from methods. A transferable region is intended to encapsulate a single data-structure, consisting of a collection of objects with a distinguished root object of some type `T`, which we refer to as the region’s *root* object. The class `Region` is parametric over the type `T` of this root object.

The `Region` constructor takes as a parameter a function that constructs the root object: it creates a new region and invokes this function, with the new region as the allocation context, to create the root object of the region. The following code illustrates the creation of a transferable region, whose root is an object of type `A`.

```
Region<A> rgn = new Region<A>(() => new A())
```

In the above code, `rgn` is called the *handler* to the newly created region, and is required to read the contents of the region, or change its state. The class `Region` offers two methods:

```

1 class SelectVertex<TIn, TOut> {
2     Func<TIn, TOut> selector;
3     Dictionary<Time, Region<List<TOut>>> map;
4     ...
5     void onReceive(Time t, Region<List<TIn>> inRgn) {
6         if (!map.ContainsKey(t))
7             map[t] = new Region<List<TOut>> (() => new List<TOut>());
8         open inRgn as inList {
9             letregion R0 {
10                foreach (TIn input in inList) {
11                    open map[t] as outList {
12                        TOut output = selector(input);
13                        outList.add(output); } } }
14                inRgn.free();
15            }
16        void onNotify(Time t) {
17            Region<List<TOut>> outRgn = map[t];
18            map.Remove(t);
19            outRgn.transfer(successorId);
20        }
21    }

```


■ **Figure 3** SELECT dataflow operator in BROOM

a free method that deallocates the region (and all the objects allocated within it), and transfer method that transfers the region to a (possibly remote) consumer process.

**Open and Closed Regions** A transferable region must be explicitly *opened* using BROOM’s open construct in order to either read or update or allocate objects in the region. Specifically, the construct “open rgn as v@R { S }” does the following: (a). It opens the transferable region handled by rgn for allocation (i.e., makes it the current allocation context), (b). binds the identifier R to this open region, and (c). initializes the newly introduced local variable v to refer to the root object of the region. The @R part of the statement is optional and may be omitted. The open construct is intended to simplify the problem of ensuring memory safety, as will be explained soon. We refer to a transferable region that has not been opened as a *closed* region. A transferable region can only be transferred or freed when it is in closed state. The acceptable state transition discipline over the lifetime of a transferable region is described in Fig. 4. Enforcement of this discipline is done at runtime.

Opening a region also makes it the current allocation context. Thus, given a C# library function f (that makes no use of BROOM’s region constructs), opening a region R and invoking f has the effect that all objects created by this invocation are allocated in the region R.

**Motivating Example** Fig 3 shows how the motivating example of Fig. 1 can be written in BROOM. The onReceive method receives its input message in a *transferred* region (i.e., a *closed* region whose ownership is transferred to the recipient). On Line 7, we create a new region to store the output for time t, initializing it to contain an empty list. On Lines 8 and 9, we open the input region inRgn followed by creating a new stack region R0. The temporary objects created by the iteration on line 10, for example, will be allocated in the stack region R0 that lives just long enough. We open the desired output region on line 11, so that the new output objects created by the invocation of selector on line 12 are allocated in the output region. Finally, the input region is freed on line 14. The output region at map[t] stays as long as input messages with timestamp t keep arriving. When the timing message for t arrives, the onNotify method transfers the outRgn at map[t] to a downstream actor.



■ **Figure 4** The lifetime of a dynamic (transferable) region in BROOM

|              |                              |   |                           |
|--------------|------------------------------|---|---------------------------|
|              | $\pi \in$ Region identifiers | $\rho \in$ Region variables   | $a, b \in$ Type variables |
|              | $m \in$ Method names         | $x, y, f \in$ Variables and fields  |                           |
| $cn$         | $\in$ Class names            | $::=$ Object   Region   A   B   |                           |
| $K$          | $\in$ FGJ class types        | $::=$ $cn\langle\bar{T}\rangle$   |                           |
| $T$          | $\in$ FGJ types              | $::=$ a   K   unit   $\bar{T} \rightarrow T$  |                           |
| $r$          | $\in$ Region annotations     | $::=$ $\rho$   $\pi$  |                           |
| $N$          | $\in$ Annotated class types  | $::=$ $cn\langle\bar{T}\rangle\langle\bar{r}\rangle$  |                           |
| $\tau$       | $\in$ types                  | $::=$ $T@r$   N   unit   $\langle\bar{\rho} \phi\rangle\bar{\tau} \xrightarrow{r} \tau$   |                           |
| $C$          | $\in$ Class definitions      | $::=$ class $cn\langle\bar{a} \triangleleft \bar{K}\rangle\langle\bar{\rho} \phi\rangle \triangleleft N\{\bar{\tau} \bar{f}; \bar{d}\}$   |                           |
| $d$          | $\in$ Methods                | $::=$ $\tau m\langle\bar{\rho} \phi\rangle(\bar{\tau} \bar{x})\{\text{return } e;\}$  |                           |
| $\phi, \Phi$ | $\in$ Region constraints     | $::=$ true   $r \succeq r$   $r = r$   $\phi \wedge \phi$   |                           |
| $e$          | $\in$ Expressions            | $::=$ ()   x   e.f   e.m( $\bar{r}$ )( $\bar{e}$ )   new N( $\bar{e}$ )<br>  $\lambda@r\langle\bar{\rho} \phi\rangle(\bar{x}:\bar{\tau}).e$   e( $\bar{r}$ )( $\bar{e}$ )   let $x = e$ in e<br>  letregion $\pi$ in e   open e as $y@r$ in e |                           |

■ **Figure 5** FEATHERWEIGHT BROOM: syntax

**Cloning** Note that in the example from Fig. 3 the object returned by the selector (on Line 12) should not contain any references to the input object, since the input region, where the object resides, will be freed at the end of the method. If there is a need for the output object to point to subobjects of the input object, such subobjects must be cloned (to copy them from the input region to the output region). Fortunately, BROOM’s region type system (§ 3) is capable of capturing such nuances in the type of selector and the type checker will ensure correctness. Furthermore, the type can be automatically inferred by BROOM’s region type inference (§ 4), which can perform the above reasoning on behalf of the programmer.

### 3 FEATHERWEIGHT BROOM

The purpose of BROOM’s region type system is to enforce the key invariant required for memory safety, namely that an object  $o_1$  in a region  $R_1$  contains a reference to an object  $o_2$  in  $R_2$ , only if  $R_2$  is guaranteed to outlive  $R_1$ . While the invariant is easily stated, enforcing it in the presence of first-class dynamic regions, parametric polymorphism (generics), and higher-order functions requires new reasoning principles that we formally develop in this section. We introduce FEATHERWEIGHT BROOM (FB), an explicitly typed core language (with region types) that incorporates the features introduced in the previous section. FEATHERWEIGHT BROOM builds on the Featherweight Generic Java (FGJ) [13] formalism, and reuses notations and various definitions from [13], such as the definition of type well-formedness for the core

(region-free) language. (The language used in Section 2 is essentially a version of FGJ without region types and with some syntactic sugar.)

### 3.1 Syntax

Fig 5 describes the syntax of FB. We refer to the class types of FGJ as *core types*. The following definition of `Pair` class in FB illustrates some of the key elements of the formal language (the symbol  $\triangleleft$  should be read *extends*, and the symbol  $\succeq$  stands for *outlives*):

```
class Pair<a  $\triangleleft$  Object, b  $\triangleleft$  Object>
    < $\rho_0, \rho_1, \rho_2$  |  $\rho_1 \succeq \rho_0 \wedge \rho_2 \succeq \rho_0$ >  $\triangleleft$  Object< $\rho_0$ > {
    a@ $\rho_1$  fst;
    b@ $\rho_2$  snd;
    a@ $\rho_1$  getFst() { return this.fst; }
}
```

Note that we elide showing constructors since they are uninteresting from the type system's standpoint; their behavior in FB is same as in FGJ.

A class in FB is parametric over zero or more type variables (as in FGJ) as well as one or more region variables  $\rho$ . We refer to the first region parameter ( $\rho_0$  in the above example) as the *allocation region* of the class: it serves to identify the region where an instance of the class is allocated<sup>2</sup>. An object in FB can contain fields referring to objects allocated in regions ( $\bar{\rho}$ ) other than its own allocation region ( $\rho$ ), provided that the former outlive the latter (i.e.,  $\bar{\rho} \succeq \rho$ ). In such case, the definition of object's class needs to be parametric over allocation regions of its fields (i.e., their classes). Furthermore, the constraint that such regions must outlive the allocation region of the class needs to be made explicit in the definition, as the `Pair` class does in the above definition. We say that the `Pair` class exhibits *constrained region polymorphism*.

To construct objects of the `Pair` class, its type and region parameters need to be instantiated with core types ( $T$ ) and region annotations<sup>3</sup> ( $r$ ), respectively. For example:

```
letregion  $\pi_0$  in
  let snd = new Object< $\pi_0$ >() in
    letregion  $\pi_1$  in
      let fst = new Object< $\pi_1$ >() in
        let p = new Pair<Object, Object>< $\pi_1, \pi_1, \pi_0$ > (fst, snd);
```

In the above code, the instantiation of  $\rho_0$  and  $\rho_1$  with  $\pi_1$ , and  $\rho_2$  with  $\pi_0$  is allowed because (a)  $\pi_0$  and  $\pi_1$  are live during the instantiation, and (b).  $\pi_0 \succeq \pi_1$  and  $\pi_1 \succeq \pi_1$  (since outlives is reflexive). Observe that the region type of `p` conveys the fact that (a). it is allocated in region  $\pi_1$ , and (b). it holds references to objects allocated in region  $\pi_0$  and  $\pi_1$ . In contrast, if we choose to allocate the `snd` object also in  $\pi_1$ , then `p` would be contained in  $\pi_1$ , and its region type would be `Pair<Object, Object>< $\pi_1, \pi_1, \pi_1$ >`, which we abbreviate as `Pair<Object, Object>@ $\pi_1$` . In general, we treat  $B\langle\bar{T}\rangle@ $\pi$$  as being equivalent to  $B\langle\bar{T}\rangle\langle\bar{\pi}\rangle$ . Region annotation on type  $a$ , where  $a$  is a type variable, assumes the form  $a@ $\pi$$ . If  $a$  is instantiated with `Pair<Object, Object>`, the result is the type of a `Pair` object contained in  $\pi$ .

<sup>2</sup>In general,  $\bar{\rho}$  denotes the sequence  $\rho_1\rho_2\dots$  of region parameters of a class or a method. In some cases, we also represent region parameters as  $\rho\bar{\rho}$  to clearly distinguish between the allocation region parameter ( $\rho$ ) and the rest.

<sup>3</sup>Region annotations ( $r$ ) include region variables ( $\rho$ ) and region identifiers ( $\pi$ ). Region identifiers are to region variables, as types ( $T$ ) are to type variables ( $a$ )



Classes in FB are independently parameterized over types and regions. While this design decision has a downside in that it allows type variables to only denote the types of objects contained in a single region, it yields benefits that outweigh the costs. In particular, it lets us support region-polymorphic higher-order functions as class fields. This allows, for example, a generic class, whose type parameters are  $a$  and  $b$ , to contain a region-polymorphic function of type  $\langle \rho_0, \rho_1, \rho_2 \mid \rho_1 \succeq \rho_0 \wedge \rho_2 \succeq \rho_0 \rangle (a@_{\rho_1}, b@_{\rho_2}) \rightarrow \text{Pair} \langle a, b \rangle \langle \rho_0, \rho_1, \rho_2 \rangle$  as its field. Such region-polymorphic higher-order fields are used frequently by the dataflow operators, which apply them in the context of various regions (*e.g.*, see Fig. 3). Keeping type and region parameterizations separate also simplifies the type system so that inference becomes practical (Sec. 4). The need for polymorphism at the field level is also why FB treats function closures specially, rather than as objects of type `Func` as in C#.

Like classes, methods can also exhibit constrained region polymorphism. A method definition in FB is necessarily polymorphic over its allocation context (§ 2.1), and optionally polymorphic with respect to the regions containing its arguments (*i.e.*, a method has at least one region parameter). Region parameters, like those on classes, are qualified with constraints ( $\phi$ ). If a method is not intended to be polymorphic with respect to its allocation context (for example, if its allocation context needs to be same as the allocation region of its *this* argument), then the required monomorphism can be captured as an equality constraint in  $\phi$ .

FB extends FGJ’s expression language with a lambda expression and an application expression ( $e(\bar{r})(\bar{e})$ ) to define and apply functions (lambdas). Functions, like methods, exhibit constrained region polymorphism, as evident in their arrow region type  $(\langle \bar{\rho} \mid \phi \rangle \bar{r} \xrightarrow{\tau})$ . A function, like a method, is necessarily polymorphic w.r.t its allocation context. Since a function closure can escape the context in which it is created, it is important to keep track of the region in which it is created in order to avoid unsafe dereferences. The  $r$  annotation above the arrow in the arrow type denotes the allocation region of the corresponding closure.

Note that mutable object fields are conspicuously absent from the FB model (and also the FGJ model in general), but this isn’t a major shortcoming considering that constructor application, which happens during a new object creation, includes assignments to the object fields (see FGJ [13]). Thus the type system is already obligated to handle unsafe assignments.

## 3.2 Types and Well-formedness

Well-formedness and typing rules of FEATHERWEIGHT BROOM establish the conditions under which a region type is considered well-formed, and an expression is considered to have a certain region type, respectively. Fig. 6 contains an illustrative subset of such rules<sup>4</sup>. The rules refer to a context ( $\mathcal{A}$ ), which is a tuple of:

- A set ( $\Delta \in 2^r$ ) of regions that are estimated to be live,
- A finite map ( $\Theta \in a \mapsto K$ ) of type variables to their *bounds*, *i.e.*, classes they are declared to extend (this artifact is inherited from FGJ), and
- A constraint formula ( $\Phi$ ) that captures the outlives constraints on regions in  $\Delta$ .

In addition, the context for the expression typing judgment also includes (a). a type environment ( $\Gamma \in x \mapsto \tau$ ) that contains the type bindings for variables in scope, and (b). the region ( $r$ ) that serves as the allocation context for the expression being type checked. Like the judgments in FGJ [13], all the judgments defined by the rules in Fig. 6 are implicitly

<sup>4</sup>Full formal development can be found in the appendix

**Type Well-formedness**  $\boxed{\mathcal{A} \vdash \tau \text{ ok}}$ 

$$\frac{CT(B) = \text{class } B\langle \bar{a} \triangleleft \bar{K} \rangle \langle \bar{\rho} \mid \phi \rangle \triangleleft N\{\dots\} \quad \bar{r} \in \Delta \quad \Theta \Vdash B\langle \bar{T} \rangle \text{ ok} \quad \Phi \vdash [\bar{r}/\bar{\rho}](\phi)}{(\Delta, \Theta, \Phi) \vdash B\langle \bar{T} \rangle \langle \bar{r} \rangle \text{ ok}}$$

$$\frac{\Theta \Vdash T \text{ ok} \quad r \in \Delta \quad \Theta \Vdash T <: \text{Object}}{(\Delta, \Theta, \Phi) \vdash T@r \text{ ok}}$$

$$\frac{\Theta \Vdash T \text{ ok} \quad \mathcal{A} = (\Delta, \Theta, \Phi)}{\mathcal{A} \vdash \text{Region}\langle T \rangle \langle \pi_{\top} \rangle \text{ ok}}$$

**Expression Typing**  $\boxed{\mathcal{A}, \Gamma, r \vdash e : \tau}$ 

$$\frac{\mathcal{A}, \Gamma, r \vdash \bar{e} : \bar{\tau} \quad \mathcal{A} \vdash N \text{ ok} \quad \text{fields}(N) = \bar{f} : \bar{\tau}}{\mathcal{A}, \Gamma, r \vdash \text{new } N(\bar{e}) : N}$$

$$\frac{\mathcal{A}, \Gamma, r \vdash e : \langle \rho \rangle \text{unit} \xrightarrow{r} T@r}{\mathcal{A}, \Gamma, r \vdash \text{new Region}\langle T \rangle \langle \pi_{\top} \rangle (e) : \text{Region}\langle T \rangle \langle \pi_{\top} \rangle}$$

$$\frac{\mathcal{A}, \Gamma, r \vdash e : \tau' \quad \bar{f} : \bar{\tau} = \text{fields}(\text{bound}_{\mathcal{A}, \Theta}(\tau'))}{\mathcal{A}, \Gamma, r \vdash e.f_i : \tau_i}$$

$$\frac{\mathcal{A} = (\Delta, \Theta, \Phi) \quad \mathcal{A}' = (\Delta \cup \{\pi\}, \Theta, \Phi \wedge \phi) \quad \pi \notin \Delta \quad \phi = \Delta \succeq \pi \quad \mathcal{A}', \Gamma, \pi \vdash e : \tau \quad \mathcal{A} \vdash \tau \text{ ok}}{\mathcal{A}, \Gamma, r \vdash \text{letregion } \pi \text{ in } e : \tau}$$

$$\frac{\mathcal{A} = (\Delta, \Theta, \Phi) \quad \pi \notin \Delta \quad \mathcal{A}, \Gamma, r \vdash e_a : \text{Region}\langle T \rangle \langle \pi_{\top} \rangle \quad \mathcal{A} \vdash \tau \text{ ok} \quad \mathcal{A}' = (\Delta \cup \{\pi\}, \Theta, \Phi) \quad \Gamma = \Gamma[y \mapsto T@_{\pi}] \quad \mathcal{A}', \Gamma', \pi \vdash e_b : \tau}{\mathcal{A}, \Gamma, r \vdash \text{open } e_a \text{ as } y@_{\pi} \text{ in } e_b : \tau}$$

$$\frac{\mathcal{A} = (\Delta, \Theta, \Phi) \quad r \in \Delta \quad \Delta \succeq r \quad \rho, \bar{\rho} \notin \Delta \quad \mathcal{A}' = (\Delta \cup \{\rho, \bar{\rho}\}, \Theta, \Phi \wedge \phi) \quad \Delta \cup \{\rho, \bar{\rho}\} \vdash \phi \text{ ok} \quad \mathcal{A}' \vdash \bar{\tau}^1 \text{ ok} \quad \mathcal{A}' \vdash \tau^2 \text{ ok} \quad \mathcal{A}', \Gamma[\bar{x} \mapsto \bar{\tau}^1], \rho \vdash e : \tau^2}{\mathcal{A}, \Gamma, r \vdash \lambda@r \langle \rho \bar{\rho} \mid \phi \rangle (\bar{x} : \bar{\tau}^1). e : \langle \rho \bar{\rho} \mid \phi \rangle \bar{\tau}^1 \xrightarrow{r} \tau^2}$$

$$\frac{\mathcal{A} = (\Delta, \Theta, \Phi) \quad \mathcal{A}, \Gamma, r \vdash e_0 : \tau \quad \text{mtype}(m, \text{bound}_{\Theta}(\tau)) = \langle \rho \bar{\rho} \mid \phi \rangle \bar{\tau}^1 \rightarrow \tau^2 \quad r, \bar{r} \in \Delta \quad \mathcal{A} \vdash \langle \rho \bar{\rho} \mid \phi \rangle \bar{\tau}^1 \rightarrow \tau^2 \text{ ok} \quad \mathcal{A}, \Gamma, r \vdash \bar{e} : [r\bar{r}/\rho\bar{\rho}] \bar{\tau}^1 \quad \Phi \vdash [r\bar{r}/\rho\bar{\rho}] \phi}{\mathcal{A}, \Gamma, r \vdash e_0.m\langle r\bar{r} \rangle (\bar{e}) : [r\bar{r}/\rho\bar{\rho}] \tau^2}$$

$$\frac{\mathcal{A} = (\Delta, \Theta, \Phi) \quad r', r, \bar{r} \in \Delta \quad \mathcal{A}, \Gamma, r \vdash e : \langle \rho \bar{\rho} \mid \phi \rangle \bar{\tau}^1 \xrightarrow{r'} \tau^2 \quad \Phi \vdash [r\bar{r}/\rho\bar{\rho}] \phi \quad \mathcal{A}, \Gamma, r \vdash \bar{e} : [r\bar{r}/\rho\bar{\rho}] \bar{\tau}^1}{\mathcal{A}, \Gamma, r \vdash e\langle r\bar{r} \rangle (\bar{e}) : [r\bar{r}/\rho\bar{\rho}] \tau^2}$$

**Method Well-formedness**  $\boxed{d \text{ ok in } B}$ 

$$\frac{\Delta = \{\bar{\rho}, \rho_m, \bar{\rho}_m\} \quad \mathcal{A} = (\Delta, [\bar{a} \mapsto \bar{K}], \phi_m) \quad \Delta \vdash \phi_m \text{ ok} \quad \mathcal{A} \vdash \bar{\tau}^1, \tau^2 \text{ ok} \quad CT(B) = \text{class } B\langle \bar{a} \triangleleft \bar{K} \rangle \langle \bar{\rho} \mid \phi \rangle \triangleleft N\{\dots\} \quad \mathcal{A}, \Gamma, \rho_m \vdash e : \tau^2 \quad \Gamma = \cdot[\bar{x} \mapsto \bar{\tau}^1][\text{this} \mapsto B\langle \bar{a} \rangle \langle \bar{\rho} \rangle] \quad \mathcal{A} \vdash \text{override}(m, N, \langle \rho_m \bar{\rho}_m \mid \phi_m \rangle \bar{\tau}^1 \rightarrow \tau^2)}{\tau^2 m\langle \rho_m \bar{\rho}_m \mid \phi_m \rangle (\bar{\tau}^1 \bar{x}) \{ \text{return } e; \} \text{ ok in } B}$$

 ■ **Figure 6** FEATHERWEIGHT BROOM: select type rules

parameterized on a class table ( $CT \in cn \mapsto D$ ) that maps class names to their definitions in FB.

The well-formedness judgment on region types ( $\mathcal{A} \vdash \tau \text{ ok}$ ) makes use of the well-formedness and subtyping judgments on core types. We use a double-piped turnstile ( $\Vdash$ ) for judgments in FGJ [13], and a simple turnstile ( $\vdash$ ) for those in FB. The class table ( $\llbracket CT \rrbracket$ ) for FGJ judgments is derived from FB's class table ( $CT$ ) by erasing all region annotations on types, and region arguments in expressions ( $\llbracket \cdot \rrbracket$  denotes the region erasure operation). The well-formedness rule for class types ( $B\langle \bar{T} \rangle \langle \bar{r} \rangle$ ) is responsible for enforcing the safety property that prevents objects from containing unsafe references. It does so by insisting that regions

$\bar{r}$  satisfy the constraints ( $\phi$ ) imposed by the class on its region parameters. The latter is enforced by checking the validity of  $\phi$ , with actual region arguments substituted for formal region parameters, under the conditions ( $\Phi$ ) guaranteed by the context. The semantics of this sequent is straightforward, and follows directly from the properties of outlives and equality relations. For any well-formed core type  $T$ ,  $T@r$  is a well-formed region type if  $r$  is a valid region. The type  $\text{Region } \langle T \rangle \langle r \rangle$  is well-formed only if  $r = \pi_{\top}$ , where  $\pi_{\top}$  is a special immortal region that outlives every other live region. This arrangement allows `Region` handlers to be aliased and referenced freely from objects in various regions, regardless of their lifetimes. On the flip side, this also opens up the possibility of references between transferable regions, which become unsafe in context of the recipient's address space. Fortunately, such references are explicitly prohibited by the type rule of `Region` objects, as described below.

The type rules distinguish between the new expressions that create objects of the `Region` class, and new expressions that create objects of other classes. The rule for the latter relies on an auxiliary definition<sup>5</sup> called `fields` (undefined for `Region` class) that returns the sequence of type bindings for fields (instance variables) of a given class type. Like in FGJ, the names and types of a constructor's arguments in `FB` are same as the names and types of its class's fields. The type rule for the field access expression ( $e.f_i$ ) also uses `fields` and another definition called `bound`, which returns the bound of a type variable (`bound` is an identity function for concrete types).

The type rule for the new `Region` expression expects the `Region` class's constructor to be called with a nullary function that returns a value in its allocation context. This ensures that the value returned by the function stores no references to objects allocated elsewhere, including the top region ( $\pi_{\top}$ ), thus preventing cross-region references originating from transferable regions. The body of the function might however create new regions while execution, but this is not a problem as long as such regions, and objects allocated in them, don't find their way into the result of its evaluation.

The type rule for the `letregion` expression requires that the static identifier introduced by the expression be unique under the current context (i.e.,  $\pi \notin \Delta$ ). This condition is needed in order to prevent the new region from incorrectly assuming existing outlives relationships on an eponymous region. The expression ( $e$ ) under `letregion` is typechecked with the new region as its allocation context, under the assumption that the region is live ( $\Delta \cup \{\pi\}$ ) and that it is outlived by all existing live regions ( $\Delta \succeq \pi$ ). The result of a `letregion` expression must have a type that is well-formed under a context not containing the new region. This ensures that the value obtained by evaluating a `letregion` expression contains no references to the temporary objects inside the region.

The rule for the `open` expression, unlike the rule for `letregion`, does not introduce any outlives relationship between the newly opened region and any pre-existing region while checking the type of the expression ( $e$ ) under `open`. This prevents new objects allocated inside the transferable region from storing references to those outside. The newly opened region becomes the allocation context for  $e$ , which is type checked under an environment ( $\Gamma$ ) extended with the type binding for the root object.

The type rule for the lambda expression typechecks the lambda-bound expression ( $e$ ) under an extended type environment containing bindings for function's arguments, assuming that region parameters are live, and that declared constraints over region parameters hold. The constraints ( $\phi$ ) are required to be well-formed under  $\Delta$  extended with the function's region parameters ( $\rho\bar{\rho}$ ). Unlike a typical object, a function closure might capture references

<sup>5</sup>All auxiliary definitions we use in this exposition originate from the FGJ calculus.

## 11:12 Safe Transferable Regions

$l \in$  Memory locations  
 $r \in$  Region annotations ::=  $l \mid \dots$   
 $s \in$  Region Typestate ::=  $\square \mid \blacksquare \mid \times$   
 $e \in$  Expressions ::=  $\text{letd } l \text{ in } e \mid \text{opened } l(s) \text{ in } e \mid \perp \mid \dots$

■ **Figure 7** FEATHERWEIGHT BROOM: extended syntax to accommodate run-time constructs

that may become unsafe if the closure escapes its allocation context. FB prevents this scenario by requiring the function closure to always be allocated in the current allocation context ( $r$ ).

The type rule for method application uses an auxiliary definition `mtype` that gives the type of a class method. Method application involves instantiation of method's region parameters, and the type rule requires the first region parameter to be instantiated with the current allocation context ( $r$ ). The type rule for function application does similarly. This requirement ensures the safety of closures as demonstrated by the following example.

Consider a method  $m$  with two region parameters -  $\rho_0$  and  $\rho_1$  (i.e.,  $m\langle\rho_0, \rho_1\rangle(\dots)$ ). Suppose the method immediately returns a function closure that contains a reference to (an object in)  $\rho_1$ . Since function closures are always allocated in the current allocation context, the closure is allocated in  $\rho_0$ , the allocation context for the method body (the rule for methods discussed below). Now, consider the following use of the method:

```
letregion R0 in  
  let f = letregion R1 in m<R0, R1> ()  
in f ()
```

Observe that the first region parameter of  $m$  is instantiated with  $R_0$ , while the allocation context for the method call is  $R_1$ . The function  $f$  returned by  $m$  stores a reference to  $R_1$ , which is unsafe when  $f$  is finally called. The type system fortunately disallows this scenario by enforcing certain restrictions during method and function applications as described above.

The method well-formedness rule makes use of an auxiliary definition `override` that judges whether the current method is a valid overriding of any eponymous method from the super class. The type environment is extended with a binding that binds the `this` keyword to the type of the current class. Note that the type of the current method as accessed via `this` (i.e., `this.m`) is region-parametric, thus admitting region-polymorphic recursion (i.e., recursive call to  $m$  can have different region arguments than  $m$ ).

### 3.3 Operational Semantics and Type Safety

The operational semantics of FB describes a non-trivial runtime component that introduces memory regions and performs run-time verification of their typestate. Fig. 7 shows the additional language constructs of FB that manifest only at run-time. A location ( $l$ ) abstracts all the memory locations associated with a region (i.e., each memory region is associated with a single location). A transferable region can be in one of three possible states: closed ( $\square$ ), open/live ( $\blacksquare$ ), and transferred/freed ( $\times$ ). Constructs `letd` and `opened` are the run-time manifestations of `letregion` and `open`; `letregion` reduces to `letd` while allocating a (static) region, and `open` reduces to `opened` while opening a (transferable) region. An `opened` expression is tagged with the typestate ( $s$ ) of the transferable region before it is opened (as Fig. 4 shows, a transferable region can be `open`'d from either an open state or a closed state). Special value  $\perp$  denotes an exception.

Operational semantics of FB defines a four-place small-step reduction relation of the form shown below:

$$(e, \Sigma) \longrightarrow (e', \Sigma')$$

The typestate of regions is tracked by a finite map ( $\Sigma$ ) from locations to typestates (Fig. 4). The reduction relation relates an expression ( $e$ ) and a typestate map ( $\Sigma$ ) to a reduced expression ( $e'$ ) and an updated typestate map ( $\Sigma'$ ). The semantics gets “stuck” if  $e$  attempts to access an object whose allocation region is not present in  $\Delta$ , or if  $e$  tries to open a transferable region that is not mapped to an appropriate typestate by  $\Sigma$ . On the other hand, if  $e$  attempts to commit an operation on a **Region** object that is not sanctioned by the transition discipline in Fig. 4, then it raises an exception value ( $\perp$ ). An illustrative subset of operational semantics that formalize the intuitions described above is shown in Fig. 8. Fig. 15 of the appendix contains rest of the rules.

To help state the type safety theorem, we define the syntactic class of (runtime) values:

$$v \in \text{values} ::= \text{new } B(\overline{T})(\overline{1})(\overline{v}) \mid \lambda @ \mathbf{1}(\overline{p} \mid \phi)(\overline{\tau} \overline{x}).e \mid \text{new Region } \langle T \rangle \langle \mathbf{1}_\top \mathbf{1} \rangle (v)$$

The first two forms are obtained by using locations for region annotations in **new** and lambda expressions. The last form is the value that **new Region** expressions gets evaluated to; the location  $\mathbf{1}_\top$  stands for the region  $\pi_\top$ , and  $\mathbf{1}$  is the location of the newly allocated transferable region. The following type safety theorem shows that a well-typed program will never attempt to dereference an “invalid” reference (a reference to an object in a region that has been transferred or freed):

► **Theorem 1. (Type Safety)**  $\forall e, \tau, \Delta, \Sigma$ , such that  $\text{consistent}(\Delta, \Sigma)$  and  $\Delta \vdash \Phi \text{ ok}$ , if  $(\Delta, \cdot, \Phi), \cdot \vdash e : \tau$ , then either  $e$  is a value, or  $e$  raises an exception  $((e, \Sigma) \longrightarrow \perp)$ , or there exists an  $e'$  and a  $\Sigma'$  such that  $\text{consistent}(\Delta, \Sigma')$  and  $(e, \Sigma) \longrightarrow (e', \Sigma')$  and  $(\Delta, \cdot, \Phi), \cdot \vdash e' : \tau$ .

The relation  $\text{consistent}$  relates  $\Delta$  and  $\Sigma$  only if both make consistent assumptions about liveness of regions. Since  $\Delta$  is consistent with both  $\Sigma$  and  $\Sigma'$ , the theorem also captures the key property of operational semantics that no live region is ever freed or transferred.

**Integrating regions with GC heap.** FB stores the region handlers in a distinguished top region, which abstracts the GC heap. Type system prevents transferable regions from storing references into the GC heap. A static region however can store references to region handles, which need to be taken into account while GC'ing the heap. Traversing the region memory to identify references into the GC heap unfortunately defeats the purpose of regions. The solution we adopt is to disable collecting region handles as long as any static region is open. Since static regions are intended to store temporary objects while populating a dynamic region, their lifetimes are short enough to let region handles to be GC'd.

## 4 Type Inference

BROOM's region type system imposes an annotation burden, and manually annotating C# standard libraries with region types can be tedious. We now present our region type inference algorithm that eliminates the need to write region type annotations. Formally, the type inference algorithm is an elaboration function from programs in  $\llbracket FB \rrbracket$  (i.e., FB without region types, but with `letregion` and `open` expressions, similar to the language introduced in § 2) to programs in FB. For simplicity, we assume that each `letregion` and `open` construct in the input program introduces a distinct region identifier  $\pi$ .

**Overview.** We now present a high-level outline of the type inference algorithm. The algorithm consists of the following steps:

$$\boxed{(e, \Sigma) \longrightarrow (e', \Sigma')}$$

$$\begin{array}{c}
\text{[LETREGIONBEGIN]} \quad \frac{1 \notin \text{dom}(\Sigma) \quad \Sigma' = \Sigma[1 \mapsto \blacksquare]}{(\text{letregion } r \text{ in } e, \Sigma) \longrightarrow (\text{letd } 1 \text{ in } [1/r]e, \Sigma')} \\
\text{[LETREGION]} \quad \frac{(e, \Sigma) \longrightarrow (e', \Sigma')}{(\text{letd } 1 \text{ in } e, \Sigma) \longrightarrow (\text{letd } 1 \text{ in } e', \Sigma')} \\
\text{[LETREGIONEND]} \quad \frac{\Sigma' = \Sigma[1 \mapsto \times]}{(\text{letd } 1 \text{ in } v, \Sigma) \longrightarrow (v, \Sigma')} \\
\text{[NEWREGION]} \quad \frac{1 \notin \text{dom}(\Sigma) \quad \Sigma' = \Sigma[1 \mapsto \square]}{(\text{new Region } \langle T \rangle \langle \pi_{\top} \rangle (v), \Sigma) \longrightarrow (\text{new Region } \langle T \rangle \langle 1_{\top} 1 \rangle (v \langle 1 \rangle ()), \Sigma')} \\
\text{[NEWREGION]} \quad \frac{(e, \Sigma[1 \mapsto \blacksquare]) \longrightarrow (e', \Sigma')}{(\text{new Region } \langle T \rangle \langle 1_{\top} 1 \rangle (e), \Sigma) \longrightarrow (\text{new Region } \langle T \rangle \langle 1_{\top} 1 \rangle (e'), \Sigma')} \\
\text{[OPEN]} \quad \frac{v_a = \text{new Region } \langle T \rangle \langle 1_{\top} 1 \rangle (v) \quad \Sigma(1) = \square \text{ or } \Sigma(1) = \blacksquare \quad \Sigma' = \Sigma[1 \mapsto \blacksquare]}{(\text{open } v_a \text{ as } x@r \text{ in } e_b, \Sigma) \longrightarrow (\text{opened } 1(\Sigma(1)) \text{ in } [v/x][1/r]e_b, \Sigma')} \\
\text{[OPENED]} \quad \frac{(e, \Sigma) \longrightarrow (e', \Sigma')}{(\text{opened } 1(s) \text{ in } e, \Sigma) \longrightarrow (\text{opened } 1(s) \text{ in } e', \Sigma')} \\
\text{[OPENEND]} \quad \frac{\Sigma' = \Sigma[1 \mapsto s]}{(\text{opened } 1(s) \text{ in } v, \Sigma) \longrightarrow (v, \Sigma')} \\
\text{[OPENTRANSFERRED]} \quad \frac{v_a = \text{new Region } \langle T \rangle \langle 1_{\top} 1 \rangle (v) \quad \Sigma(1) \neq \square \text{ and } \Sigma(1) \neq \blacksquare}{(\text{open } v_a \text{ as } x@r \text{ in } e_b, \Sigma) \longrightarrow \perp} \\
\text{[TRANSFER]} \quad \frac{\Sigma(1) = \square \quad \Sigma' = \Sigma[1 \mapsto \times]}{((\text{new Region } \langle T \rangle \langle 1_{\top} 1 \rangle (v)).\text{transfer}(\dots), \Sigma) \longrightarrow ((), \Sigma')} \\
\text{[TRANSFEROPENED]} \quad \frac{\Sigma(1) = \blacksquare}{((\text{new Region } \langle T \rangle \langle 1_{\top} 1 \rangle (v)).\text{transfer}(\dots), \Sigma) \longrightarrow \perp}
\end{array}$$

■ **Figure 8** FEATHERWEIGHT BROOM: a select subset of operational semantics

1. *Region Parameterization.* The first step elaborates the input program by introducing *formal region parameters* (for each class and method), and *region variables* (representing yet undetermined *actual region parameters*). We also introduce for each class and method, a *predicate variable* ( $\varphi$ ) to denote an undetermined set of outlives-constraints over the region parameters of that class/method.
2. *Constraint Generation.* In the second step, we analyze the program to generate a set of constraints (over the region identifiers and the predicate variables) that must hold (as per the static semantics in Fig. 6).
3. *Constraint Solving.* We solve the generated set of constraints using our fixpoint constraint solving algorithm, which reduces the constraint solving problem to an abduction problem.

If the original program in  $\llbracket FB \rrbracket$  contains unsafe references, for example, a reference from a transferable region to a stack region, then the constraints generated during the elaboration are not satisfiable. In such a case, the solver fails to solve the constraints.

4. If the solver succeeds, it returns a solution  $\eta$  consisting of a pair of substitution functions  $\eta_R$  and  $\eta_P$  for free region and predicate variables, respectively, introduced in step 1. We apply these substitutions to the elaborated program to produce the final program.

We later on establish the soundness of the type inference algorithm, and the completeness of the constraint-generation and constraint-solving steps (steps 2 and 3 above).

## 4.1 Region Parameterization for Classes

Region parameterization is an iterative process involving the following three steps, the first two of which are mutually dependent on each other.

*Introduction of Formal Region Parameters.* For every class  $C$ , we identify a sequence of formal region parameters  $\pi_0, \dots, \pi_n$  that  $C$  should be parametric over.

*Introduction of Actual Region Parameters.* We then replace every instance of class  $C$  in the program by an instance  $C \langle \rho_0, \dots, \rho_n \rangle$ , where  $\rho_0, \dots, \rho_n$  are fresh identifiers denoting actual region parameters.

*Predicate Variable Introduction.* For every class  $C$ , we introduce a fresh predicate variable  $\varphi$ , which represents the yet undetermined outlives-constraints between the formal region parameters of class  $C$ .

We identify the region parameters of classes as follows.

*Non-Recursive Classes.* The class `Object` is defined to have a single region parameter  $\pi_0$  (the allocation region). The region parameters for any other non-recursive class  $C$  is determined only after the region parameters of any class that  $C$  depends on have been determined: this includes the base-class  $B$  of  $C$  and the class (type) of any of its fields. We replace every dependee type  $T$  in  $C$  by its instantiated type, using fresh region parameters as needed. The sequence of region parameters for  $C$  is defined to be the sequence of region parameters for the base class  $B$  concatenated with the list of all fresh region parameters introduced while instantiating the types of the fields in the class. (The class inherits its allocation region from its base class. Note that if a class does not specify an explicit base class, it has an implicit base class `Object`.)

This transformation is illustrated below, using a non-generic `Pair` class:

```

class Pair < Object {
  Object fst;
  Object snd;
}
⇒
class Pair < ρ₀, ρ₁, ρ₂ | φ < Object < ρ₀ > {
  Object < ρ₁ > fst;
  Object < ρ₂ > snd;
}

```

*Recursive Classes.* The region parameters for a recursive class is computed in a similar fashion, with the following difference: any recursive field is ignored while instantiating region parameters for the fields of the class, and the region parameters of the recursive class are computed as before. We then do parameter instantiation for all recursive fields, such that their region annotations (the actual region parameters) are exactly the same as the (formal) region parameters of the class. The following example illustrates this for a non-generic `List` class. The resulting class represents a linked list with spine in the region  $\rho_0$  and data objects in the region  $\rho_1$ .

```

class List < Object {
  Object data;
  List next;
}
⇒
class List < ρ₀, ρ₁ | φ < Object < ρ₀ > {
  Object < ρ₁ > data;
  List < ρ₀, ρ₁ > next;
}

```

The above technique can be extended to mutually recursive classes in a straightforward manner, by simultaneously parameterizing them (and then instantiating them).

*Type-Parametric Classes.* The type parameter  $T$  of a class  $C$  is instantiated as  $T@ρ$  using a single region parameter  $ρ$ . (This can be extended to use the bound specified for  $T$ , if any.)

*Function Types.* Since FB is higher-order, fields of function type are allowed. We explain how the parameter instantiation step instantiates function types below, after discussing parameterization for methods.

## 4.2 Region Parameterization for Methods and Function Types

As the next step, we introduce region parameters for every method. We do this by instantiating the types of all parameters and the return value (of the method) using fresh region identifiers (as explained previously), and then generalizing these region identifiers as formal region parameters of the method. In addition, a fresh region identifier is introduced to represent the allocation region. We also introduce a fresh predicate variable  $φ$  for every method, just as we did for each class. Thus, the method

```
Object m (List x) {...}
```

is instantiated as

```
Object⟨ρ3⟩ m⟨ρ0, ρ1, ρ2, ρ3⟩ (List⟨ρ1, ρ2⟩ x) {...}
```

We then consider every method invocation in the program, and introduce fresh region variables representing the (yet unknown) actual region parameters for this particular invocation. We similarly perform instantiation for every constructor invocation of the form  $new@π_0 T(\dots)$ , by instantiating the type  $T$  as before, turning it into  $new T⟨π_0, ρ_1, \dots, ρ_n⟩(\dots)$ , where  $ρ_1, \dots, ρ_n$  are fresh region variables. Note that the programmer typically specifies the region  $π_0$  where the object is to be allocated. If the programmer does not specify this, it is allocated in the allocation-context region by default. Region variables are introduced only for the other region parameters.

Function types (of fields and parameters) are instantiated just like methods above, (and the region where the closure is allocated is determined just as for other objects). For example, the function type  $List \rightarrow Object$  is instantiated as  $⟨ρ_0, ρ_1, ρ_2, ρ_3 \mid φ⟩ List⟨ρ_1, ρ_2⟩ \xrightarrow{π} Object⟨ρ_3⟩$ . Note that the newly introduced region identifiers are generalized as formal region parameters of the function type. This is a (heuristic) choice made in the case of higher order functions. Consider a higher order function  $f$  with a function typed parameter  $g$ . The fresh region identifiers introduced while instantiating the type of  $g$  could be alternatively generalized as formal region parameters of  $f$ , but we choose to generalize them as formal region parameters of  $g$ . We will discuss this aspect again later.

## 4.3 Constraint Generation

The constraint generation algorithm mimics the static type checker, but accumulates constraints that must hold for the type checking to succeed.

**Syntax of Constraints.** (See Fig. 9.) The constraints are expressed using a set  $R$  of region constants, a set  $V$  of region variables, and a set  $P$  of predicate variables.

Recall that a *region constant* may be either (a) a *formal region parameter* of a class or method, or (b) a *static region identifier* introduced by a `letregion` construct, or (c) an *open transferable region identifier* introduced by an `open` construct. A *region variable* is introduced to represent an unknown *actual* region parameter of a method invocation or



$$\begin{array}{l}
\pi \in R \text{ (Region constants)} \quad \nu \in V \text{ (Region vars)} \quad \varphi \in P \text{ (Predicate vars)} \quad \Delta \subseteq R \\
\rho \in \text{Region Identifiers} \quad ::= \quad \pi \mid \nu \\
\phi \in \text{Region Constraint} \quad ::= \quad \text{true} \mid \rho \succeq \rho \mid \phi \wedge \phi \\
F \in \text{Substitution} \quad ::= \quad \cdot \mid [\rho/\rho]F \\
\ell \in \text{Antecedent} \quad ::= \quad \phi \mid \varphi \mid \varphi \wedge \phi \\
r \in \text{Consequent} \quad ::= \quad \phi \mid F(\varphi) \\
\text{Constraint} \quad ::= \quad \ell \vdash r \mid \nu \in \Delta \mid \Delta \vdash \varphi \text{ ok}
\end{array}$$

■ **Figure 9** Syntax of constraints

object allocation, and the constraint-solver, if successful, will bind each region variable to a region constant.

A *region-constraint*  $\phi$  consists of a conjunction of *outlives-constraints* of the form  $\rho_1 \succeq \rho_2$ . A predicate variable  $\varphi$  is introduced to represent, *e.g.*, the unknown precondition of a method. The constraint-solver will end up binding it to a *region-constraint*  $\phi$  over a set of fixed formal region parameters. Our constraints also make use of *pending substitutions*  $F$ : A pending substitution serves to bind formal region parameters in  $\varphi$  to the actual region parameters used in a particular context: *E.g.*, in the validity constraint  $\pi_1 \succeq \pi_2 \vdash [\pi_1/\rho_1][\pi_2/\rho_2]\varphi$ , the pending substitution is  $[\pi_1/\rho_1][\pi_2/\rho_2]$ .

The constraints are primarily of the form  $\varphi_i \wedge \phi_{cx} \vdash \phi_{cs}$  or  $\varphi_i \wedge \phi_{cx} \vdash F_j(\varphi_j)$ . Here,  $\varphi_i$  is a predicate variable (representing the precondition of a method to be determined),  $\phi_{cs}$  is a region-constraint that is *required* to hold at a particular program point (within the method), and  $\phi_{cx}$  is a region-constraint that is *known* to hold at that program point. Constraints of the form  $\varphi_i \wedge \phi_{cx} \vdash F_j(\varphi_j)$  are generated by an invocation of a method with precondition  $\varphi_j$ .

We use *well-formedness constraints* of the form  $\rho \in \Delta$  to restrict the domain of unification for a region variable ( $\rho$ ) to a constant set  $\Delta = \{\pi_1, \dots, \pi_n\}$  of regions in scope, and well-formedness constraints of form  $\Delta \vdash \varphi \text{ ok}$  to restrict the domain of a predicate variable ( $\varphi$ ) to the set of all possible region-constraint formulas over a fixed set of regions ( $\Delta = \{\pi_1, \dots, \pi_n\}$ ) in scope.

**Constraint Solution.** We define an *assignment*  $\eta$  to be a pair of functions  $(\eta_R, \eta_P)$ , where  $\eta_R$  is a map from  $V$  to  $R$  and  $\eta_P$  is a map from  $P$  to a region-constraint formula. Such an assignment is said to *satisfy* a set of constraints  $C$  if every sequent in  $C$  is valid after the substitutions  $\eta_P$  and  $\eta_R$ . We say that  $C$  is *satisfiable* if it has a satisfying assignment.

**Constraint Generation.** The constraint generation algorithm is a direct adaption of the type checker: each type checking judgment is modified to produce a set of constraints that must hold for the type checker to succeed.

Fig. 10 illustrates this for selected language constructs. (The appendix contains the remaining constraint generation rules.) These rules use the same context as the corresponding typing judgment in Fig. 6, except that this is generalized to permit the use of region variables and predicate variables. Symbols  $\mathcal{A}$  and  $\Gamma$  retain their meaning, modulo this extension. Thus, the component  $\Phi$  of  $\mathcal{A}$  will now be a *symbolic expression* of the form  $\varphi \wedge \phi$ , where  $\varphi$  is the predicate variable representing the undetermined precondition of the method being analyzed. The algorithm proceeds top-down, analyzes an expression  $e$  in a context  $\mathcal{A}, \Gamma, r$ , and returns a type  $\tau$  and a set of constraints  $C$  (expressed in the rules as  $\mathcal{A}, \Gamma, r \vdash e : \tau \triangleleft C$ ), indicating

Expression Typing Constraint Generation  $\boxed{\mathcal{A}, \Gamma, r \vdash e : \tau \triangleleft C}$ 

$$\begin{array}{c}
\text{[NEW]} \quad \frac{\mathcal{A}, \Gamma, r \vdash \bar{e} : \bar{\tau} \triangleleft C_1 \quad \mathcal{A} \vdash N \text{ ok} \triangleleft C_2 \quad \text{fields}(N) = \bar{f} : \bar{\tau}}{\mathcal{A}, \Gamma, r \vdash \text{new } N(\bar{e}) : N \triangleleft C_1 \cup C_2} \\
\text{[NEWREGION]} \quad \frac{\mathcal{A}, \Gamma, r \vdash e : \langle \rho \rangle \text{unit} \xrightarrow{\tau} T @ \rho \triangleleft C}{\mathcal{A}, \Gamma, r \vdash \text{new Region} \langle T \rangle \langle \pi_{\top} \rangle (e) : \text{Region} \langle T \rangle \langle \pi_{\top} \rangle \triangleleft C} \\
\text{[LETREGION]} \quad \frac{\mathcal{A} = (\Delta, \Theta, \Phi) \quad \pi \notin \Delta \quad \mathcal{A}' = (\Delta \cup \{\pi\}, \Theta, \Phi \wedge (\Delta \succeq \pi))}{\mathcal{A}', \Gamma, \pi \vdash e : \tau \triangleleft C_1 \quad \mathcal{A} \vdash \tau \text{ ok} \triangleleft C_2} \\
\text{[FNAPPLY]} \quad \frac{\mathcal{A} = (\Delta, \Theta, \Phi) \quad C_1 = \{r\bar{r} \in \Delta\} \quad \mathcal{A}, \Gamma, r \vdash e : \langle \rho\bar{\rho} \mid \phi \rangle \bar{\tau}^1 \xrightarrow{\tau} \tau^2 \triangleleft C_2}{C_3 = \{\Phi \vdash [r\bar{r}/\rho\bar{\rho}]\phi\} \quad \mathcal{A}, \Gamma, r \vdash \bar{e} : [r\bar{r}/\rho\bar{\rho}]\bar{\tau}^1 \triangleleft C_4} \\
\text{[OPEN]} \quad \frac{\mathcal{A} = (\Delta, \Theta, \Phi) \quad \pi \notin \Delta \quad \mathcal{A}, \Gamma, r \vdash e_a : \text{Region} \langle T \rangle \langle \pi_{\top} \rangle \triangleleft C_1}{(\Delta \cup \{\pi\}, \Theta, \Phi), \Gamma[y \mapsto T @ \pi] \vdash e_b : \tau \triangleleft C_2 \quad \mathcal{A} \vdash \tau \text{ ok} \triangleleft C_3} \\
\mathcal{A}, \Gamma, r \vdash \text{open } e_a \text{ as } y @ \pi \text{ in } e_b : \tau \triangleleft (C_1 \cup C_2 \cup C_3)
\end{array}$$

■ **Figure 10** Constraint generation rules part 1

that the expression will have a type  $\tau$  provided the constraints  $C$  hold.

When we adapt a type-checking judgement rule to produce a corresponding constraint-generation rule, we treat the antecedent conditions of the type-checking rule in one of two ways. Many of these antecedent conditions are converted into constraints which are accumulated in the set  $C$ . However, some of the antecedent conditions are expected to be trivially satisfied and are expressed as antecedent conditions in the constraint-generation rule as well. E.g., our frontend ensures that the region constants introduced by `open` and `letregion` constructs are unique (and the elaboration phase ensures this by alpha-renaming them as needed). Thus, the precondition  $\rho \notin \Delta$  in rules `LETREGION` and `OPEN` is expected to hold true at constraint-generation time (and does not produce any constraints).

The rules for generating constraints from a method and class definition first build a context ( $\mathcal{A}$ ) containing a set ( $\Delta$ ) denoting regions that are currently live, a map ( $\Theta$ ) mapping type variables to their bounds, and a constraint formula ( $\Phi$ ) capturing constraints over live region variables. We use predicate variables ( $\varphi$  and  $\varphi_m$ ) to capture constraints over variables in  $\Delta$  that are yet to be inferred.

Let  $\text{GENCONSTRAINT}(q)$  denote the set of constraints generated from an elaborated program  $q$ . The following theorem states that constraint-generation is sound and complete:

► **Theorem 2.** *Let  $C = \text{GENCONSTRAINT}(q)$ . An assignment  $\eta$  (for the region and predicate variables in  $q$ ) satisfies  $C$  iff  $q[\eta]$  is well-typed.*

#### 4.4 The Constraint Solver

We refer to any validity constraint whose antecedent contains a predicate variable (*i.e.*, is of the form  $\varphi \wedge \phi_{cx}$ ) as an *abduction constraint*. Here,  $\phi_{cx}$  is either *true* (in which case, we call the

constraint a trivial abduction constraint) or a conjunction of one or more outlives-constraints (in which case, we call the constraint a non-trivial abduction constraint).

Non-trivial abduction constraints are a key challenge in solving the constraints. We now describe some special properties of the set of constraints  $C$  generated by our algorithm, which allow us to handle abduction constraints efficiently.

For any predicate variable  $\varphi$  used in  $C$ ,  $C$  has exactly one constraint of the form  $\Delta \vdash \varphi \text{ ok}$ . We will refer to this  $\Delta$  as  $\Delta_P^C(\varphi)$ . (Our constraint-generation algorithm guarantees the preceding property. Even otherwise, a set of constraints  $\Delta_i \vdash \varphi \text{ ok}$  can be replaced by the single equivalent constraint  $(\bigcap_i \Delta_i) \vdash \varphi \text{ ok}$ .) We say that an abduction constraint is  $C$ -decomposable if its antecedent is of the form  $\varphi \wedge \phi_{cx}$  where  $\varphi$  is a predicate variable and  $\phi_{cx}$  is a conjunction of zero or more outlives-constraints of the form  $\pi_1 \succeq \pi_2$  satisfying the following conditions: (1)  $\pi_2 \notin \Delta_P^C(\varphi)$ . (2) if  $\pi_1 \in \Delta_P^C(\varphi)$ , then for every  $\pi_f \in \Delta_P^C(\varphi)$ ,  $\pi_f \succeq \pi_2$  is a conjunct in  $\phi_{cx}$ . We will omit the reference to  $C$  in the above notation if no confusion is likely.

► **Lemma 3.** *Every abduction constraint in  $C$ , where  $C = \text{GENCONSTRAINT}(q)$ , is  $C$ -decomposable.*

► **Lemma 4.** *Consider any  $C$ -decomposable constraint  $\varphi \wedge \phi \vdash \pi_i \succeq \pi_j$  where both  $\pi_i$  and  $\pi_j$  are region constants. Let  $\eta$  satisfy  $C$ .*

- (a) *If  $\{\pi_i, \pi_j\} \subseteq \Delta_P(\varphi)$ :  $\eta$  satisfies  $\varphi \wedge \phi \vdash \pi_i \succeq \pi_j$  iff  $\eta$  satisfies  $\varphi \vdash \pi_i \succeq \pi_j$ .*  
 (b) *If  $\{\pi_i, \pi_j\} \not\subseteq \Delta_P(\varphi)$ :  $\eta$  satisfies  $\varphi \wedge \phi \vdash \pi_i \succeq \pi_j$  iff  $\eta$  satisfies  $\phi \vdash \pi_i \succeq \pi_j$*

The above lemma shows how we can reduce a non-trivial abduction constraint to either a trivial abduction constraint or a non-abduction constraint, provided that the consequent is an outlives-constraint without region variables.

**Constraint Solver.** The first step in our algorithm for solving a set of constraints  $C$  computes a set  $C^* \supseteq C$  of constraints by iteratively applying the following rules until a fixed point is reached:

1. (Initialization)  $\ell \vdash r \in C \Rightarrow \ell \vdash r \in C^*$
2. (Transitivity)
  - a.  $\ell \vdash \rho_1 \succeq \rho_2 \in C^*, \ell \wedge \phi \vdash \rho_2 \succeq \rho_3 \in C^* \Rightarrow \ell \wedge \phi \vdash \rho_1 \succeq \rho_3 \in C^*$
  - b.  $\ell \wedge \phi \vdash \rho_1 \succeq \rho_2 \in C^*, \ell \vdash \rho_2 \succeq \rho_3 \in C^* \Rightarrow \ell \wedge \phi \vdash \rho_1 \succeq \rho_3 \in C^*$
3. (Substitution)  $\ell \vdash F(\varphi) \in C^*, \varphi \vdash \phi \in C^* \Rightarrow \ell \vdash F(\phi) \in C^*$
4. (Abduction Decomposition)
  - a.  $\varphi \wedge \phi_{cx} \vdash \pi_i \succeq \pi_j \in C^*, \{\pi_i, \pi_j\} \subseteq \Delta_P(\varphi) \Rightarrow \varphi \vdash \pi_i \succeq \pi_j \in C^*$
  - b.  $\varphi \wedge \phi_{cx} \vdash \pi_i \succeq \pi_j \in C^*, \{\pi_i, \pi_j\} \not\subseteq \Delta_P(\varphi), \{\pi_i, \pi_j\} \subseteq R \Rightarrow \phi_{cx} \vdash \pi_i \succeq \pi_j \in C^*$

We can show that every new constraint added by the above rules is implied by existing constraints:

► **Theorem 5.** *Let  $C = \text{GENCONSTRAINT}(q)$ . An assignment  $\eta$  satisfies  $C$  iff  $\eta$  satisfies  $C^*$ .*

A key goal of this step is to identify the value every region variable must have in any solution of the set of constraints, as explained below. Consider a set of constraints  $D$ . We say that a region variable  $\rho$  occurs in a context  $\ell$  (in  $D$ ) if  $D$  contains some constraint  $\ell \vdash r$  where  $\rho$  occurs in  $r$ . We say that a region variable  $\rho$  is bound to a region constant  $\pi$  in a context  $\ell$  if  $\{\ell \vdash \rho \succeq \pi, \ell \vdash \pi \succeq \rho\} \subseteq D$ . We say that a region variable  $\rho$  is bound to a region constant  $\pi$  (in  $D$ ) if  $\rho$  is bound to  $\pi$  in every context  $\ell$  in which it occurs. We say

that a region variable  $\rho$  is *somewhere-bound* to a region constant  $\pi$  (in  $D$ ) if  $\rho$  is bound to  $\pi$  in some context  $\ell$  in which it occurs.

The set  $C^*$  makes it easy to identify a solution  $\hat{\eta}$  (if one exists), as below. For any predicate variable  $\varphi$ ,  $\hat{\eta}_P(\varphi)$  is defined to be  $\bigwedge\{\pi_1 \succeq \pi_2 \mid \pi_1, \pi_2 \in R, \varphi \vdash \pi_1 \succeq \pi_2 \in C^*\}$ . For any region variable  $\rho$ , we define  $\hat{\eta}_R(\rho)$  to be any element of the set  $\{\pi \in R \mid \rho \text{ is somewhere-bound to } \pi \text{ in } C^*\}$ , if this set is non-empty. If this set is empty for any  $\rho$ ,  $\hat{\eta}$  is undefined.

The set  $C^*$  also makes it easy to check if  $C$  is satisfiable. Let  $C_g^*$  denote the subset of all ground constraints (*i.e.*, constraints without any region variable or predicate variable) in  $C^*$ . Let  $WF_R$  denote the subset of all well-formedness constraints for region variables in  $C^*$ . Define  $SOLVE(C)$  as below.

$$SOLVE(C) = \begin{cases} \text{SOME}(\hat{\eta}) & \text{if } C_g^* \text{ is valid and } \hat{\eta} \text{ is defined and satisfies } WF_R \\ \text{NONE} & \text{otherwise} \end{cases}$$

► **Theorem 6.** *Let  $C = \text{GENCONSTRAINT}(q)$ . (Soundness) If  $SOLVE(C) = \text{SOME}(\eta)$ , then  $\eta$  satisfies  $C$ . (Completeness) If  $SOLVE(C) = \text{NONE}$ , then  $C$  is unsatisfiable.*

Checking the validity of a ground constraint is straightforward. A ground constraint is of the form  $\bigwedge_{i \in I} \ell_i \vdash r$ , where each  $\ell_i$  and  $r$  is an outlives-constraint. Let  $D$  denote  $\{\vdash \ell_i \mid i \in I\}$ . The given constraint is valid iff  $r$  belongs to  $D^*$ .

**Algorithmic Aspects.** Note that checking the validity of a ground constraint can be realized using a simple graph reachability algorithm. Given a set  $S$  of outlives constraints, define the directed graph  $G(S) = (V(S), E(S))$  as follows. Every distinct region identifier  $\rho$  in  $S$  is represented by a vertex, which we will also refer to as  $\rho$ . Every outlives constraint  $\rho_1 \succeq \rho_2$  is represented by an edge from  $\rho_1$  to  $\rho_2$ . It is easy to see that  $\bigwedge S \vdash \rho_1 \succeq \rho_2$  iff there exists a path from  $\rho_1$  to  $\rho_2$  in  $G(S)$ . Thus, a simple graph reachability algorithm can be used to check the validity of ground constraints.

This idea generalizes. Extending the simple graph reachability algorithm to incorporate the Substitution rule (in computing  $C^*$ ) turns the problem into a context-free reachability problem in graphs [16] (as usual for context-sensitive interprocedural analysis).

Algorithms for context-free reachability can be adapted to incorporate the Abduction Decomposition step. Alternatively, the iterative process described above is a standard fixed point computation and can be encoded using a set of Datalog rules, allowing us to compute the closure using any Datalog engine.

## 4.5 Soundness and Completeness: Discussion

Theorems 2 and 6 show that the second and third steps of the type-inference algorithm are sound. It is easy to verify the soundness of the first step (the elaboration phase): For any program  $p \in \llbracket FB \rrbracket$ , the elaboration phase produces a  $q$  such that for any assignment  $\eta$ , we have  $\llbracket q[\eta] \rrbracket = p$ . The soundness of the type inference algorithm follows.

Theorem 2 and 6 also establish the completeness of the constraint-generation and constraint-solving steps. The only source of incompleteness in the type inference algorithm is the set of heuristic choices made during the first step, as explained below.

(1) We determine the set of region parameters for a recursive class using the heuristic that a recursive occurrence of the class has the same parameters, in the same order, as the class itself. This heuristic fails, for example, if the program uses a recursive list type whose

elements alternatively come from two different regions. Such a program would require the following elaboration, which is beyond the scope of our approach:

```

class List < Object
{
  Object data;
  List next;
}
⇒
class List < (ρ0, ρ1, ρ2 | φ) < Object < (ρ0)
{
  Object < (ρ1) data;
  List < (ρ0, ρ2, ρ1) next;
}

```

(2) Our technique for region parameterization also uses a heuristic in the case of higher order programs. The following examples illustrates that principal types may not exist for higher order functions.

```

unit apply ( T → unit f, T x, T y ) { f(x); f(y); }

```

This method may be typed assuming either that  $f$  is polymorphic over the region that its parameter is allocated in (permitting  $x$  and  $y$  to be allocated in any regions), or by assuming that  $x$  and  $y$  are allocated in the same region  $\rho_1$  that  $f$  expects its parameters to be allocated in. Neither type subsumes the other. Our algorithm heuristically chooses the first option, as it appears to be the more likely and useful candidate.

If users provide partial region annotations, especially in situations (such as above) where elaboration makes a heuristic choice, the elaboration procedure can use the user-provided choices instead. This can help the type-inference overcome these limitations.

## 4.6 Modularity Aspects of Type Inference

The type inference algorithm, as presented, traverses the entire program to generate the set of constraints, which are solved en masse, using an iterative fixed point computation. However, the type inference can be realized in a modular and compositional fashion, subject only to the restrictions imposed by recursion.

In the elaboration phase, we can process a class  $C$  only after any class  $B$  that  $C$  depends on has been processed: class  $C$  depends on class  $B$  if  $B$  is either  $C$ 's base class or the type of any field of  $C$  depends on  $B$ . In effect, this means that any collection of mutually recursive classes must be processed together. Non-recursive dependences can be handled in a compositional fashion: if class  $C$  depends on  $B$  non-recursively, then the elaboration can be done for  $B$  first, and then  $C$  can be processed.

The same idea applies to the constraint-solving phase as well. Given a set of constraints, we say that a predicate variable  $\varphi_1$  *directly-depends* on another predicate variable  $\varphi_2$  if the set of constraints includes a constraint  $\varphi_1 \wedge \phi_{cx} \vdash F(\varphi_2)$ . We say that  $\varphi_1$  *depends* on  $\varphi_2$  if  $\varphi_1$  transitively depends on  $\varphi_2$ . The constraint solver needs to process any collection of mutually dependent predicate variables together. In effect, this requires the type inference to process any collection of mutually recursive methods together. However, methods that are not mutually recursive can be processed separately.

## 5 Implementation and Evaluation

As mentioned earlier, our work is a continuation of the work reported in [7], which provides users with a C#-based implementation of transferable regions (the features described in Section 2). This system has no type system and provides users with no safety guarantees. [7] presents evidence that realistic programs can be implemented using transferable regions and that this can yield significant performance gains. In particular, it reports speedups up to 34% for typical big-data analytics jobs. We now describe our implementation and experience with the type system and type inference algorithm presented in the current paper.

```

class LinkedList<T><R5,R4 | R4>R5> {
  ListNode<T><R5,R4> head; ...
  List<T><R17,R4> rev<R17,R4 | R4>R17>(unit u) {
    List<T><R17,R4> xs = new List<T><R17,R4>(this.head.val);
    ListNode<T><R5,R4> cur = this.head.next;
    while (!cur == Null) {
      xs.add<R17>(cur.val)
      cur = cur.next; }
    return xs;
  }
}

```

■ **Figure 11** Region-annotated definition of `rev` computed by BROOMC

We have implemented BROOMC, a prototype of BROOM compiler frontend, including its region type system and type inference, in 3k+ lines of OCaml. The input to BROOMC is a program in  $\llbracket FB^+ \rrbracket$ , an extended version of  $\llbracket FB \rrbracket$  that includes assignments, conditionals, loops, more primitive datatypes (*e.g.*, integers), and a null value. Our implementation of region type inference and constraint solving closely follows the description given in Sec. 4.

We performed two kinds of experiments to evaluate our region type system and type inference. First, we implemented some microbenchmarks ( $\leq 100$  LOC) consisting of standard classes such as pairs, lists, list iterators, etc., in  $\llbracket FB^+ \rrbracket$ , and used our inference engine to infer their region types. These classes are region-oblivious. Hence, as long as they are well-typed as per the core type system, BROOMC must be able to automatically construct its region-type-annotated definition without fail. BROOMC was able to infer the expected region types for all these classes under 10ms. Fig 11 shows the region-type-annotated definition computed for the list reverse method. Observe that BROOMC was able to infer that the list and its data (of type  $\tau$ ) can be allocated in different regions, as long as the latter outlives the former. This allows, for instance, a `preOrder` method to traverse a tree in a transferable region, and return a list of its nodes, where the list itself is allocated in the stack region.

Next, we translated 4 out of the 6 Naiad streaming query operator benchmarks (Naiad vertices) used in [7] to  $\llbracket FB^+ \rrbracket$ , and used BROOMC to verify their safety. The 2 remaining benchmarks were left out because their region behavior (from the perspective of the type system) is subsumed by the included benchmarks. The number of LOC performing operations on `Region` objects relative to the total LOC is 8% or under in the Naiad benchmarks. During the process, we found multiple instances of potential memory safety violations in the  $\llbracket FB^+ \rrbracket$  translation of all the 4 Naiad vertices, which we verified to be present in the original C# implementation as well. The cause of all safety violations is the creation of a reference from the outgoing message (a transferable region) to the payload of the incoming message. For example, the implementation of `SelectVertex` contains the following:

```

if (this.selector(inMsg.payload[i])) {
  outMsg.set(outputOffset, inMsg.payload[i]);
  ...
}

```

The `outMsg` is later transferred to a downstream actor, where the reference to `inMsg`'s payload becomes unsafe<sup>6</sup>. We eliminated such unsafe references by creating a clone of `inMsg.payload[i]` in `outMsg`, and our compiler was subsequently able to certify the safety of all references.

<sup>6</sup>This unsafe reference could have gone unnoticed during experiments in [7] because their experimental setup included only one actor.

Our experience with Naiad benchmarks shows the utility of our type inference/checking tool, particularly because it comes at no additional cost to the developer.

## 6 Related Work

Following Tofte and Talpin’s seminal work in [15, 20, 21], static type systems for safe region-based memory management have been extensively studied in the context of various languages and problem settings [8, 10, 24, 2, 1, 9, 23, 17, 11]. Our work differs from the existing proposals in one or more of the following respects.

1. Our design choice focuses on ensuring memory-safety while giving programmers control over region management and allocation of objects in regions. In contrast, some systems automate all aspects of memory management. This is a convenience-performance trade-off.
2. We support both lexically scoped (stack) regions and dynamic transferable regions (both programmer-managed).
3. We exploit a combination of a simple static type discipline and lightweight runtime checks to ensure memory safety. In particular, our approach circumvents the need for restrictive static mechanisms (e.g., linear types and unique pointers) or expensive runtime mechanisms (e.g., garbage collection and reference counting) in order to guarantee safety.
4. We present a full (interprocedural) type inference algorithm that eliminates the need to write region annotations on types.
5. Our underlying language is an object-oriented programming language, equipped with higher-order functions and parameterized (generic) types. These language features necessitate some non-trivial choices in the design of the region-parametricity aspect of the language, which also have an impact on aspects such type inference.

Tofte and Talpin’s approach [21] uses compiler-managed lexically scoped (stack) regions (as a replacement for GC). Our type inference is analogous to theirs in some respects, while differing in others. Their inference algorithm only generates equality constraints, solvable via unification. Our type inference algorithm generates partial order outlives constraints. Consequently, our constraint solving algorithm is more sophisticated, and is capable of inferring unknown outlives constraints over region arguments of polymorphic recursive functions.

Walker and Watkins [23] extend lambda calculus with first-class regions with dynamic lifetimes, and impose linear typing to control accesses to regions. Our open/close lexical block for transferable regions traces its origins to the `let!` expression in [23] and [22], which safely relaxes linear typing restrictions, allowing variables to be temporarily aliased. We don’t use linear typing (for references to regions), thus admit unrestricted aliasing, but use lightweight runtime checks for safety. Moreover, [23]’s linear type system is insufficient to enforce the invariants needed to ensure safety under region transfers, such as the absence of references that escape a transferable region.

Cyclone [8] equips C with programmer-managed stack regions, and a typing discipline that statically guarantees the safety of all pointer dereferences. Later proposals [10, 19] extends Cyclone with dynamic regions. BROOM differs from Cyclone in its non-intrusiveness design principle, which requires its safety mechanisms to not intrude on the programming practices of C#. BROOM programmers, for example, shouldn’t be forced to abandon iterators in favor of for-loops, annotate region types, or rewrite C#’s standard libraries to use in BROOM. Cyclone requires C programmers to use new language constructs and abandon some standard programming idioms in the interest of preserving safety. For instance, Cyclone

programmers are required to write region types for functions; the type inference is only intraprocedural. Ensuring safety in presence of dynamic regions requires using either unique pointers or reference-counted objects. Both approaches are intrusive. For example, unique pointers constrain, or in some cases forbid, the use of the familiar iterator pattern, which requires creation of aliases to objects in a collection. Some standard library functions, for example, those that use caching, may need to be rewritten. Moreover, even with unique pointers, safety cannot be guaranteed statically; checks against `NULL` are needed at run-time to enforce safety. For ref-counted objects, Cyclone requires programmers to use special functions (`alias_refptr` and `drop_refptr`) to create and destroy aliases. Reference count is affected only by these functions. An alias going out of scope, for instance, does not decrement the ref-count. The requirement to use additional constructs to manage aliases makes reference counting more-or-less as intrusive as unique pointers.

Our work differs from Cyclone also in terms of its technical contributions. While Cyclone equips C with a range of region constructs [19], the semantics of (a significant subset of) such constructs, and the safety guarantees of the language are not formalized. In contrast, the (static and dynamic) semantics of Broom has been rigorously defined with respect to a well-understood formal system (FGJ). The safety guarantees have been formalized and proved. Similar contrast can be made of region type inference in both the languages. Cyclone’s type inference was only ever described as being similar to Tofte and Talpin’s, and its effectiveness in presence of tracked pointers is not clear. In contrast, a detailed type inference algorithm is one of our core contributions.

Our region type system can also be thought of as a specialized ownership type system [5], where each region is the owner of all objects allocated in the region. An ownership type system for safe region-based memory management in real-time Java has been proposed by Boyapati *et al.* [2]. Their language permits only lexically-scoped (stack) regions. In contrast, we permit regions with dynamically determined lifetimes. Our language also admits generics and higher-order functions. We also establish type safety and transfer safety results that formalize the guarantees provided by our system. While Boyapati *et al.*’s language is explicitly typed, our language comes equipped with full type inference. However, several inference algorithms have been proposed in the context of other ownership type systems. Our type inference algorithm is also novel compared to these existing ownership inference algorithms, which are based on, e.g., pointer analysis [12] or boolean satisfiability [6]. (See Section 5.2 of Clark *et al.*’s survey of ownership type systems [5] for a more comprehensive discussion of ownership inference algorithms.) Some distinguishing characteristics of our algorithm is that it is customized to our problem, does not use any pointer analysis algorithm (which can be a source of imprecision) or SAT solvers (which can be a source of inefficiency), and comes with relative completeness guarantees.

Henglein *et al.* [9] propose a flow-sensitive approach for first-order programs to generalize Tofte and Talpin’s approach to dynamic regions. Cherem and Rugina [4] describe a flow-insensitive and context-sensitive analysis that transforms Java programs to use (dynamic) regions. However, neither of them supports dynamic regions as first-class objects; they cannot be stored in data structures or passed to methods. Furthermore, while Henglein *et al.* [9] require reference counting to ensure memory safety, Cherem and Rugina’s analysis [4] comes with no formal safety guarantees. Holk *et al.* [11] use regions to safely transfer data between the CPU and GPU in the context of Scheme. However, their setting only includes lexically-scoped regions for which Tofte and Talpin-style analysis suffices. In contrast, we provide first-class support for transferable regions with dynamic lifetimes.



---

**References**

---

- 1 Robert L. Bocchino, Jr., Vikram S. Adve, Danny Dig, Sarita V. Adve, Stephen Heumann, Rakesh Komuravelli, Jeffrey Overbey, Patrick Simmons, Hyojin Sung, and Mohsen Vakilian. A type and effect system for deterministic parallel java. In *Proceedings of the 24th ACM SIGPLAN Conference on Object Oriented Programming Systems Languages and Applications*, OOPSLA '09, pages 97–116, New York, NY, USA, 2009. ACM. URL: <http://doi.acm.org/10.1145/1640089.1640097>, doi:10.1145/1640089.1640097.
- 2 Chandrasekhar Boyapati, Alexandru Salcianu, William Beebe, Jr., and Martin Rinard. Ownership types for safe region-based memory management in real-time java. In *Proceedings of the ACM SIGPLAN 2003 Conference on Programming Language Design and Implementation*, PLDI '03, pages 324–337, New York, NY, USA, 2003. ACM. URL: <http://doi.acm.org/10.1145/781131.781168>, doi:10.1145/781131.781168.
- 3 Badrish Chandramouli, Jonathan Goldstein, Mike Barnett, Robert DeLine, John C. Platt, James F. Terwilliger, and John Wernsing. Trill: A high-performance incremental query processor for diverse analytics. *PVLDB*, 8(4):401–412, 2014. URL: <http://www.vldb.org/pvldb/vol8/p401-chandramouli.pdf>.
- 4 Sigmund Cherech and Radu Rugina. Region analysis and transformation for java programs. In *Proceedings of the 4th International Symposium on Memory Management*, ISMM '04, pages 85–96, New York, NY, USA, 2004. ACM. URL: <http://doi.acm.org/10.1145/1029873.1029884>, doi:10.1145/1029873.1029884.
- 5 Dave Clarke, Johan Östlund, Ilya Sergey, and Tobias Wrigstad. *Ownership Types: A Survey*, pages 15–58. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013. URL: [http://dx.doi.org/10.1007/978-3-642-36946-9\\_3](http://dx.doi.org/10.1007/978-3-642-36946-9_3), doi:10.1007/978-3-642-36946-9\_3.
- 6 Werner Dietl, Michael D. Ernst, and Peter Müller. Tunable static inference for generic universe types. In *ECOOP 2011 - Object-Oriented Programming - 25th European Conference, Lancaster, UK, July 25-29, 2011 Proceedings*, pages 333–357, 2011. URL: [http://dx.doi.org/10.1007/978-3-642-22655-7\\_16](http://dx.doi.org/10.1007/978-3-642-22655-7_16), doi:10.1007/978-3-642-22655-7\_16.
- 7 Ionel Gog, Jana Giceva, Malte Schwarzkopf, Kapil Vaswani, Dimitrios Vytiniotis, Ganesan Ramalingam, Manuel Costa, Derek Gordon Murray, Steven Hand, and Michael Isard. Broom: Sweeping out garbage collection from big data systems. In *15th Workshop on Hot Topics in Operating Systems, HotOS XV, Kartause Ittingen, Switzerland, May 18-20, 2015*, 2015. URL: <https://www.usenix.org/conference/hotos15/workshop-program/presentation/gog>.
- 8 Dan Grossman, Greg Morrisett, Trevor Jim, Michael Hicks, Yanling Wang, and James Cheney. Region-based memory management in cyclone. In *Proceedings of the ACM SIGPLAN 2002 Conference on Programming Language Design and Implementation*, PLDI '02, pages 282–293, New York, NY, USA, 2002. ACM. URL: <http://doi.acm.org/10.1145/512529.512563>, doi:10.1145/512529.512563.
- 9 Fritz Henglein, Henning Makholm, and Henning Niss. A direct approach to control-flow sensitive region-based memory management. In *Proceedings of the 3rd ACM SIGPLAN International Conference on Principles and Practice of Declarative Programming*, PPDP '01, pages 175–186, New York, NY, USA, 2001. ACM. URL: <http://doi.acm.org/10.1145/773184.773203>, doi:10.1145/773184.773203.
- 10 Michael Hicks, Greg Morrisett, Dan Grossman, and Trevor Jim. Experience with safe manual memory-management in cyclone. In *Proceedings of the 4th International Symposium on Memory Management*, ISMM '04, pages 73–84, New York, NY, USA, 2004. ACM. URL: <http://doi.acm.org/10.1145/1029873.1029883>, doi:10.1145/1029873.1029883.

- 11 Eric Holk, Ryan Newton, Jeremy Siek, and Andrew Lumsdaine. Region-based memory management for gpu programming languages: Enabling rich data structures on a spartan host. In *Proceedings of the 2014 ACM International Conference on Object Oriented Programming Systems Languages & Applications*, OOPSLA '14, pages 141–155, New York, NY, USA, 2014. ACM. URL: <http://doi.acm.org/10.1145/2660193.2660244>, doi:10.1145/2660193.2660244.
- 12 Wei Huang, Werner Dietl, Ana Milanova, and Michael D. Ernst. Inference and checking of object ownership. In *ECOOP 2012 - Object-Oriented Programming - 26th European Conference, Beijing, China, June 11-16, 2012. Proceedings*, pages 181–206, 2012. URL: [http://dx.doi.org/10.1007/978-3-642-31057-7\\_9](http://dx.doi.org/10.1007/978-3-642-31057-7_9), doi:10.1007/978-3-642-31057-7\_9.
- 13 Atsushi Igarashi, Benjamin C. Pierce, and Philip Wadler. Featherweight java: A minimal core calculus for java and gj. *ACM Trans. Program. Lang. Syst.*, 23(3):396–450, May 2001. URL: <http://doi.acm.org/10.1145/503502.503505>, doi:10.1145/503502.503505.
- 14 Martin Maas, Tim Harris, Krste Asanovic, and John Kubiawicz. Trash day: Coordinating garbage collection in distributed systems. In *Proceedings of the 15th USENIX Conference on Hot Topics in Operating Systems*, HOTOS'15, pages 1–1, Berkeley, CA, USA, 2015. USENIX Association. URL: <http://dl.acm.org/citation.cfm?id=2831090.2831091>.
- 15 Mads Tofte and Jean-Pierre Talpin. A Theory of Stack Allocation in Polymorphically Typed Languages. Technical Report DIKU-report 93/15, University of Copenhagen, 1993. URL: <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.51.6564>.
- 16 Thomas W. Reps. Program analysis via graph reachability. *Information & Software Technology*, 40(11-12):701–726, 1998. URL: [http://dx.doi.org/10.1016/S0950-5849\(98\)00093-7](http://dx.doi.org/10.1016/S0950-5849(98)00093-7), doi:10.1016/S0950-5849(98)00093-7.
- 17 The Rust Programming Language, 2015. Accessed: 2015-11-7 13:21:00. URL: <https://doc.rust-lang.org/book>.
- 18 Type-safe off-heap memory, 2016. Accessed: 2016-06-6 13:21:00. URL: <https://github.com/densh/scala-offheap>.
- 19 Nikhil Swamy, Michael Hicks, Greg Morrisett, Dan Grossman, and Trevor Jim. Safe manual memory management in cyclone. *Science of Computer Programming*, 62(2):122–144, 2006. Special Issue: Five perspectives on modern memory management - Systems, hardware and theory. URL: <http://www.sciencedirect.com/science/article/pii/S0167642306000785>, doi:<http://dx.doi.org/10.1016/j.scico.2006.02.003>.
- 20 Mads Tofte and Jean-Pierre Talpin. Implementation of the typed call-by-value  $\lambda$ -calculus using a stack of regions. In *Proceedings of the 21st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '94, pages 188–201, New York, NY, USA, 1994. ACM. URL: <http://doi.acm.org/10.1145/174675.177855>, doi:10.1145/174675.177855.
- 21 Mads Tofte and Jean-Pierre Talpin. Region-based memory management. *Inf. Comput.*, 132(2):109–176, February 1997. URL: <http://dx.doi.org/10.1006/inco.1996.2613>, doi:10.1006/inco.1996.2613.
- 22 Philip Wadler. Linear Types Can Change the World! In M. Broy and C. B. Jones, editors, *IFIP TC 2 Working Conference on Programming Concepts and Methods*, pages 561–581, Sea of Gallilee, Israel, April 1990. North-Holland.
- 23 David Walker and Kevin Watkins. On regions and linear types (extended abstract). In *Proceedings of the Sixth ACM SIGPLAN International Conference on Functional Pro-*

- gramming*, ICFP '01, pages 181–192, New York, NY, USA, 2001. ACM. URL: <http://doi.acm.org/10.1145/507635.507658>, doi:10.1145/507635.507658.
- 24 Bennett Norton Yates. A type-and-effect system for encapsulating memory in java. Master's thesis, Department of Computer Science and Information Science, University of Oregon, 1999.
  - 25 Matei Zaharia. New developments in spark, 2015. URL: <http://www.slideshare.net/databricks/new-developments-in-spark>.

## A Appendix

### A.1 Full Static Semantics

$$\begin{array}{ll}
\text{allocRgn}(A\langle r\bar{r}\rangle\langle\bar{T}\rangle) & = r & \text{bound}_{\Theta}(a@r) & = \Theta(a)@r \\
\text{allocRgn}(\langle\bar{\rho}|\phi\rangle\bar{\tau}^1 \xrightarrow{r} \tau^2) & = r & \text{bound}_{\Theta}(N) & = N \\
\text{shape}(A\langle\bar{r}\rangle\langle\bar{T}\rangle) & = A\langle\bar{T}\rangle & \text{fields}(\text{Object}\langle r\rangle) & = \bullet
\end{array}$$

$$\frac{CT(B) = \text{class } B\langle\bar{a}\triangleleft\bar{K}\rangle\langle\bar{\rho}|\phi\rangle\triangleleft N\{\bar{\tau}^f \bar{f}; \dots\} \quad \mathcal{S} = [\bar{r}/\bar{\rho}, \bar{T}/\bar{a}] \quad \text{fields}(\mathcal{S}(N)) = \bar{g} : \bar{\tau}^g}{\text{fields}(B\langle\bar{T}\rangle\langle\bar{r}\rangle) = \bar{g} : \bar{\tau}^g, \bar{f} : \mathcal{S}(\bar{\tau}^f)}$$

$$\frac{CT(B) = \text{class } B\langle\bar{a}\triangleleft\bar{K}\rangle\langle\bar{\rho}|\phi\rangle\triangleleft N\{\bar{\tau}^f \bar{f}; \bar{d}\} \quad m \notin \bar{d} \quad \mathcal{S} = [\bar{r}/\bar{\rho}, \bar{T}/\bar{a}]}{\text{mtype}(m, B\langle\bar{T}\rangle\langle\bar{r}\rangle) = \text{mtype}(m, \mathcal{S}(N))}$$

$$\frac{CT(B) = \text{class } B\langle\bar{a}\triangleleft\bar{K}\rangle\langle\bar{\rho}|\phi\rangle\triangleleft N\{\bar{\tau}^f \bar{f}; \bar{d}\} \quad \tau^2 m\langle\bar{\rho}_m|\phi_m\rangle\langle\bar{\tau}^1 \bar{x}\rangle\{\dots\} \in \bar{d} \quad \mathcal{S} = [\bar{r}/\bar{\rho}, \bar{T}/\bar{a}]}{\text{mtype}(m, B\langle\bar{T}\rangle\langle\bar{r}\rangle) = \mathcal{S}(\langle\bar{\rho}_m|\phi_m\rangle\bar{\tau}^1 \rightarrow \tau^2)}$$

$$\frac{\text{mtype}(m, N) = \langle\bar{\rho}_1|\phi_1\rangle\bar{\tau}^{11} \rightarrow \tau^{12} \quad \text{implies} \quad \mathcal{A}.\Phi \vdash \phi_2 \Leftrightarrow [\bar{\rho}_2/\bar{\rho}_1](\phi_1) \quad \text{and} \quad \bar{\tau}^{21} = [\bar{\rho}_2/\bar{\rho}_1](\bar{\tau}^{11}) \quad \text{and} \quad \mathcal{A} \vdash \tau^{22} <: [\bar{\rho}_2/\bar{\rho}_1](\tau^{12})}{\mathcal{A} \vdash \text{override}(m, N, \langle\bar{\rho}_2|\phi_1\rangle\bar{\tau}^{21} \rightarrow \tau^{22})}$$

■ **Figure 12** FEATHERWEIGHT BROOM: auxiliary definitions

Subtyping  $\boxed{\mathcal{A} \vdash \tau_1 <: \tau_2}$

$$\frac{\mathcal{A} \vdash \tau <: \tau \quad (\Delta, \Theta, \Phi) \vdash a@p <: \Theta(a)@p \quad CT(B) = \text{class } B\langle\bar{a}\triangleleft\bar{K}\rangle\langle\bar{\rho}|\phi\rangle\triangleleft N\{\dots\}}{\mathcal{A} \vdash B\langle\bar{T}\rangle\langle\bar{r}\rangle <: [\bar{r}/\bar{\rho}, \bar{T}/\bar{a}](N)}$$

$$\frac{\mathcal{A} \vdash \tau_1 <: \tau_2 \quad \mathcal{A} \vdash \tau_2 <: \tau_3 \quad \mathcal{A} \vdash \tau_1 <: \tau_3 \quad \mathcal{A}.\Phi \vdash \phi_1 \Rightarrow \phi_2 \quad \mathcal{A} \vdash \bar{\tau}^{11} <: \bar{\tau}^{21} \quad \mathcal{A} \vdash \tau^{22} <: \tau^{12}}{\mathcal{A} \vdash \langle\bar{\rho}|\phi_2\rangle\bar{\tau}^{21} \xrightarrow{r} \tau^{22} <: \langle\bar{\rho}|\phi_1\rangle\bar{\tau}^{11} \xrightarrow{r} \tau^{12}}$$

Type, and Type Constraint Well-formedness  $\boxed{\mathcal{A} \vdash \tau \text{ ok}, \Delta \vdash \phi \text{ ok}}$

$$\frac{\frac{r \in \Delta}{(\Delta, \Theta, \Phi) \vdash \text{Object}\langle r\rangle \text{ ok}} \quad \frac{r_0, r_1 \in \Delta}{\Delta \vdash r_0 \succeq r_1 \text{ ok}} \quad \frac{\Delta \vdash \phi_0 \text{ ok} \quad \Delta \vdash \phi_1 \text{ ok}}{\Delta \vdash \phi_0 \wedge \phi_1 \text{ ok}}}{r \in \Delta \quad \bar{\rho} \notin \Delta \quad \Delta' = \Delta \cup \{\bar{\rho}\} \quad \mathcal{A}' = (\Delta', \Theta, \Phi \wedge \phi) \quad \Delta' \vdash \phi \text{ ok} \quad \mathcal{A}' \vdash \bar{\tau}^1 \text{ ok} \quad \mathcal{A}' \vdash \tau^2 \text{ ok}}{(\Delta, \Theta, \Phi) \vdash \langle\bar{\rho}|\phi\rangle\bar{\tau}^1 \xrightarrow{r} \tau^2 \text{ ok}}$$

Class Well-formedness  $\boxed{B \text{ ok}}$

$$\frac{\Delta = \{\rho, \bar{\rho}\} \quad \Theta = [\bar{a} \mapsto \bar{K}] \quad \Phi = \phi \quad \mathcal{A} = (\Delta, \Theta, \Phi) \quad \Delta \vdash \phi \text{ ok} \quad \Theta \Vdash \bar{K} \text{ ok} \quad \bar{d} \text{ ok in } B \quad \mathcal{A} \vdash N, \bar{\tau}^f \text{ ok} \quad \text{shape}(N) \neq \text{Region}\langle T \rangle \quad \Phi \vdash \text{allocRgn}(\bar{\tau}^f) \succeq \rho \quad \text{allocRgn}(N) = \rho}{\text{class } B\langle\bar{a}\triangleleft\bar{K}\rangle\langle\bar{\rho}|\phi\rangle\triangleleft N\{\bar{\tau}^f \bar{f}; \bar{d}\} \text{ ok}}$$

■ **Figure 13** FEATHERWEIGHT BROOM: subtyping and well-formedness rules (continuation of Fig. 6)

Expression Typing  $\boxed{\mathcal{A}, \Gamma, r \vdash e : \tau}$

$$\begin{array}{c}
\frac{\mathcal{A}, \Gamma, r \vdash e_1 : \tau_1 \quad \mathcal{A}, \Gamma[x \mapsto \tau_1], r \vdash e_2 : \tau_2}{\mathcal{A}, \Gamma, r \vdash \text{let } x = e_1 \text{ in } e_2 : \tau_2} \quad \frac{\mathcal{A} = (\Delta, \Theta, \Phi) \quad 1 \notin \Delta \quad \mathcal{A}' = (\Delta \cup \{1\}, \Theta, \Phi \wedge \Delta \geq 1) \quad \mathcal{A}', \Gamma, 1 \vdash e : \tau \quad \mathcal{A} \vdash \tau \text{ ok}}{\mathcal{A}, \Gamma, r \vdash \text{letd } 1 \text{ in } e : \tau} \\
\\
\frac{\mathcal{A} = (\Delta, \Theta, \Phi) \quad \mathcal{A}' = (\Delta \cup \{1\}, \Theta, \Phi) \quad \mathcal{A}' \vdash T@1 \text{ ok} \quad \mathcal{A}', \Gamma, 1 \vdash e : T@1}{\mathcal{A}, \Gamma, r \vdash \text{new Region } \langle T \rangle \langle 1\tau 1 \rangle (e) : \text{Region } \langle T \rangle \langle \pi\tau \rangle} \\
\\
\frac{\mathcal{A} = (\Delta, \Theta, \Phi) \quad \mathcal{A}' = (\Delta \cup \{1\}, \Theta, \Phi) \quad 1 \in \Delta \text{ implies } s = \blacksquare \quad \mathcal{A}', \Gamma, 1 \vdash e : \tau \quad \mathcal{A} \vdash \tau \text{ ok}}{\mathcal{A}, \Gamma, r \vdash \text{opened } 1(s) \text{ in } e : \tau} \quad \frac{\mathcal{A}, \Gamma, r \vdash e : \tau \quad \mathcal{A} \vdash \tau <: \tau'}{\mathcal{A}, \Gamma, r \vdash e : \tau'}
\end{array}$$

■ **Figure 14** FEATHERWEIGHT BROOM: expression typing (continuation of Fig. 6)

## A.2 Full Operational Semantics

$\boxed{(e, \Sigma) \longrightarrow (e', \Sigma')}$

$$\begin{array}{c}
\text{[EVALORDER]} \quad \frac{(e, \Sigma) \longrightarrow (e', \Sigma')}{(E[e], \Sigma) \longrightarrow (E[e'], \Sigma')} \\
\\
\text{[EXCEPTION]} \quad \frac{(e, \Sigma) \longrightarrow \perp}{(E[e], \Sigma) \longrightarrow \perp} \\
\\
\text{[FIELDACCESS]} \quad \frac{\Sigma(1) = \blacksquare \quad \text{fields}(A\langle \bar{T} \rangle \langle 1\bar{1} \rangle) = \bar{f} : \bar{\tau}}{((\text{new } A\langle \bar{T} \rangle \langle 1\bar{1} \rangle(\bar{v})).f_i, \Sigma) \longrightarrow (v_i, \Sigma)} \\
\\
\text{[METHODINV]} \quad \frac{\Sigma(1) = \blacksquare \quad \text{mbody}(m, A\langle \bar{T} \rangle \langle 1\bar{1} \rangle) = \rho\bar{\rho}.\bar{x}.e}{((\text{new } A\langle \bar{T} \rangle \langle 1\bar{1} \rangle(\bar{v})).m\langle 1\bar{1}' \rangle(\bar{v}'), \Sigma) \longrightarrow (([1'\bar{1}'/\rho\bar{\rho}][\bar{v}'/\bar{x}][\text{new } A\langle \bar{T} \rangle \langle 1\bar{1} \rangle(\bar{v})/\text{this}]e, \Sigma)} \\
\\
\text{[FNAPPLY]} \quad \frac{v_a = \lambda@1'(\rho\bar{\rho})(\bar{\tau} \bar{x}).e \quad \Sigma(1') = \blacksquare}{(v_a\langle 1\bar{1} \rangle(\bar{v}), \Sigma) \longrightarrow (([\bar{v}/\bar{x}][1\bar{1}/\rho\bar{\rho}]e, \Sigma)} \\
\\
\text{[EXCEPTION]} \quad \frac{\Sigma(1) = \square \quad (e, \Sigma[1 \mapsto \blacksquare]) \longrightarrow \perp}{(\text{new Region } \langle T \rangle \langle 1\tau 1 \rangle (e), \Sigma) \longrightarrow \perp} \\
\\
\text{[LETEXP]} \quad \frac{}{(\text{let } x = v \text{ in } e, \Sigma) \longrightarrow ((v/x)e, \Sigma)}
\end{array}$$

Evaluation Context  $\boxed{E}$

$$\begin{array}{l}
E ::= \bullet \mid (\bullet).f \mid \bullet.m\langle \bar{1} \rangle(\bar{e}) \mid v.m\langle \bar{1} \rangle(\dots, \bullet, \dots) \mid \text{new } N(\dots, \bullet, \dots) \\
\quad \mid \text{new Region } \langle T \rangle \langle \pi\tau \rangle(\bullet) \mid \bullet\langle \bar{1} \rangle(\bar{e}) \mid v\langle \bar{1} \rangle(\dots, \bullet, \dots) \\
\quad \mid \text{let } x = \bullet \text{ in } e \mid \text{open } \bullet \text{ as } y@r \text{ in } e
\end{array}$$

■ **Figure 15** FEATHERWEIGHT BROOM: operational semantics (continuation of Fig. 8)

## A.3 Full Constraint Generation Semantics

Expression Typing Constraint Generation  $\boxed{\mathcal{A}, \Gamma, r \vdash e : \tau \triangleleft C}$

[UNIT]  $\mathcal{A}, \Gamma, r \vdash () : \text{unit} \triangleleft \{\}$

[VAR]  $\mathcal{A}, \Gamma, r \vdash x : \Gamma(\tau) \triangleleft \{\}$

[FIELDACCESS] 
$$\frac{\mathcal{A}, \Gamma, r \vdash e : \tau' \triangleleft C \quad \bar{f} : \bar{\tau} = \text{fields}(\text{bound}_{\mathcal{A}, \Theta}(\tau'))}{\mathcal{A}, \Gamma, r \vdash e.f_i : \tau_i \triangleleft C}$$

[LET] 
$$\frac{\mathcal{A}, \Gamma, r \vdash e_1 : \tau_1 \triangleleft C_1 \quad \mathcal{A}, \Gamma[x \mapsto \tau_1], r \vdash e_2 : \tau_2 \triangleleft C_2}{\mathcal{A}, \Gamma, r \vdash \text{let } x = e_1 \text{ in } e_2 : \tau_2 \triangleleft C_1 \cup C_2}$$

[METHODINV] 
$$\frac{\begin{array}{l} \mathcal{A}, \Gamma, r \vdash e_0 : \tau \triangleleft C_1 \quad C_2 = \{r\bar{r} \in \mathcal{A}.\Delta\} \\ \text{mtype}(m, \text{bound}_{\mathcal{A}, \Theta}(\tau)) = \langle \rho\bar{\rho} \mid \phi \rangle \bar{\tau}^1 \rightarrow \tau^2 \\ \mathcal{A} \vdash \langle \rho\bar{\rho} \mid \phi \rangle \bar{\tau}^1 \rightarrow \tau^2 \text{ ok} \triangleleft C_3 \quad \mathcal{A}, \Gamma, r \vdash \bar{e} : [r\bar{r}/\rho\bar{\rho}](\bar{\tau}^1) \triangleleft C_4 \\ C_5 = \{\mathcal{A}.\Phi \vdash [r\bar{r}/\rho\bar{\rho}](\phi)\} \end{array}}{\mathcal{A}, \Gamma, r \vdash e_0.m(r\bar{r})(\bar{e}) : [r\bar{r}/\rho\bar{\rho}](\tau^2) \triangleleft C_1 \cup C_2 \cup C_3 \cup C_4 \cup C_5}$$

[LAMBDA] 
$$\frac{\begin{array}{l} r \in \mathcal{A}.\Delta \quad \rho\bar{\rho} \notin \mathcal{A}.\Delta \quad \mathcal{A}' = (\mathcal{A}.\Delta \cup \{\rho\bar{\rho}\}, \mathcal{A}.\Theta, \mathcal{A}.\Phi \wedge \phi) \\ \mathcal{A}'.\Delta \vdash \phi \text{ ok} \quad \mathcal{A}' \vdash \bar{\tau}^1 \text{ ok} \triangleleft C_1 \quad \mathcal{A}' \vdash \tau^2 \text{ ok} \triangleleft C_2 \\ \mathcal{A}', \Gamma[\bar{x} \mapsto \bar{\tau}^1], \rho \vdash e : \tau^2 \triangleleft C_3 \end{array}}{\mathcal{A}, \Gamma, r \vdash \lambda @ r(\rho\bar{\rho} \mid \phi)(\bar{x} : \bar{\tau}^1).e : \langle \rho\bar{\rho} \mid \phi \rangle \bar{\tau}^1 \xrightarrow{r} \tau^2 \triangleleft \bigcup_{i=1}^4 C_i}$$

[SUBTYPING] 
$$\frac{\mathcal{A}, \Gamma, r \vdash e : \tau \triangleleft C_1 \quad \mathcal{A} \vdash \tau <: \tau' \triangleleft C_2}{\mathcal{A}, \Gamma, r \vdash e : \tau' \triangleleft C_1 \cup C_2}$$

Method Well-formedness Constraint Generation  $\boxed{\vdash d \text{ ok in } B \triangleleft C}$

[METHOD] 
$$\frac{\begin{array}{l} CT(B) = \text{class } B(\bar{a} \triangleleft \bar{K}) \langle \bar{\rho} \mid \varphi \rangle \triangleleft N \{ \dots \} \\ \mathcal{A} = (\Delta, \Theta, \Phi) = (\{\bar{\rho}, \rho_m, \bar{\rho}_m\}, [\bar{a} \mapsto \bar{K}], \varphi_m) \quad C_1 = \{\Delta \vdash \varphi_m \text{ ok}\} \\ \Gamma = \cdot [\text{this} \mapsto B(\bar{a}) \langle \bar{\rho} \rangle] [\bar{x} \mapsto \bar{\tau}^1] \quad \text{mtype}(m, N) = \langle \bar{\rho}_m \mid \phi_m \rangle \bar{\tau}^1 \rightarrow \tau^2 \\ \mathcal{A}, \Gamma, \rho_m \vdash e : \tau^2 \triangleleft C_2 \quad \mathcal{A} \vdash \bar{\tau}^1 \text{ ok} \triangleleft C_3 \quad \mathcal{A} \vdash \tau^2 \text{ ok} \triangleleft C_4 \end{array}}{\vdash \tau^2 m \langle \rho_m \bar{\rho}_m \mid \varphi_m \rangle (\bar{\tau}^1 \bar{x}) \{ \text{return } e; \} \text{ ok in } B \triangleleft (C_1 \cup C_2 \cup C_3 \cup C_4)}$$

Class Well-formedness Constraint Generation  $\boxed{\vdash B \text{ ok} \triangleleft C}$

[CLASS] 
$$\frac{\begin{array}{l} \mathcal{A} = (\Delta, \Theta, \Phi) = (\{\rho, \bar{\rho}\}, [\bar{a} \mapsto \bar{K}], \varphi) \\ C_1 = \{\Delta \vdash \varphi \text{ ok}\} \quad \Theta \Vdash \bar{K} \text{ ok} \quad \mathcal{A} \vdash N \text{ ok} \triangleleft C_2 \quad \mathcal{A} \vdash \bar{\tau}^f \text{ ok} \triangleleft C_3 \\ C_4 = \{\Phi \vdash \text{allocRgn}(\bar{\tau}^f) \succeq \rho \wedge \text{allocRgn}(N) = \rho\} \\ \vdash \bar{d} \text{ ok in } B \triangleleft C_5 \end{array}}{\vdash \text{class } B(\bar{a} \triangleleft \bar{K}) \langle \rho\bar{\rho} \mid \varphi \rangle \triangleleft N \{ \bar{\tau}^f \bar{x}; \bar{d} \} \text{ ok} \triangleleft \bigcup_{i=1}^5 C_i}$$

■ **Figure 16** FEATHERWEIGHT BROOM: Constraint generation rules part 2 (Continuation of Fig. 10)

**Subtyping Constraint Generation**  $\boxed{\mathcal{A} \vdash \tau_1 <: \tau_2 \triangleleft C}$

[REFLEXIVITY] 
$$\mathcal{A} \vdash \tau <: \tau \triangleleft \{\}$$

[UNIFY] 
$$\mathcal{A} \vdash \tau <: [\pi/\rho](\tau) \triangleleft \{\pi \succeq \rho, \rho \succeq \pi\}$$

[TRANSITIVITY] 
$$\frac{\mathcal{A} \vdash \tau_1 <: \tau_2 \triangleleft C_1 \quad \mathcal{A} \vdash \tau_2 <: \tau_3 \triangleleft C_2}{\mathcal{A} \vdash \tau_1 <: \tau_3 \triangleleft C_1 \cup C_2}$$

[FNSUBTYPING] 
$$\frac{C_1 = \{\mathcal{A}.\Phi \vdash \phi_1 \Rightarrow \phi_2\} \quad \mathcal{A} \vdash \overline{\tau^{11}} <: \overline{\tau^{21}} \triangleleft C_2 \quad \mathcal{A} \vdash \tau^{22} <: \tau^{12} \triangleleft C_3}{\mathcal{A} \vdash \langle \overline{\rho} | \phi_2 \rangle \overline{\tau^{21}} \xrightarrow{r} \tau^{22} <: \langle \overline{\rho} | \phi_1 \rangle \overline{\tau^{11}} \xrightarrow{r} \tau^{12} \triangleleft C_1 \cup C_2 \cup C_3}$$

**Type Well-formedness Constraint Generation**  $\boxed{\mathcal{A} \vdash \tau \text{ ok} \triangleleft C}$

[OBJECTTYPE] 
$$\frac{C = \{r \in \Delta\}}{(\Delta, \Theta, \Phi) \vdash \text{Object} \langle r \rangle \text{ ok} \triangleleft C}$$

[CLASSTYPE] 
$$\frac{CT(B) = \text{class } B \langle \overline{a} \triangleleft \overline{K} \rangle \langle \overline{\rho} | \phi \rangle \triangleleft N \{ \dots \} \quad \Theta \Vdash B \langle \overline{T} \rangle \text{ ok} \quad C = \{ \overline{r} \in \Delta, \Phi \vdash [\overline{r}/\overline{\rho}](\phi) \}}{(\Delta, \Theta, \Phi) \vdash B \langle \overline{T} \rangle \langle \overline{r} \rangle \text{ ok} \triangleleft C}$$

[TYPEPARAM] 
$$\frac{\Theta \Vdash T \text{ ok} \quad \Theta \Vdash T <: \text{Object} \quad C = \{r \in \Delta\}}{(\Delta, \Theta, \Phi) \vdash T@r \text{ ok} \triangleleft C}$$

[FNSTYPE] 
$$\frac{C_1 = \{r \in \Delta\} \quad \overline{\rho} \notin \Delta \quad \Delta' = \Delta \cup \{ \overline{\rho} \} \quad \mathcal{A}' = (\Delta', \Theta, \Phi \wedge \phi) \quad \Delta' \vdash \phi \text{ ok} \triangleleft C_2 \quad \mathcal{A}' \vdash \overline{\tau^1} \text{ ok} \triangleleft C_3 \quad \mathcal{A}' \vdash \tau^2 \text{ ok} \triangleleft C_4}{(\Delta, \Theta, \Phi) \vdash \langle \overline{\rho} | \phi \rangle \overline{\tau^1} \xrightarrow{r} \tau^2 \text{ ok} \triangleleft C_1 \cup C_2 \cup C_3 \cup C_4}$$

[REGIONTYPE] 
$$\frac{\Theta \Vdash T \text{ ok}}{(\Delta, \Theta, \Phi) \vdash \text{Region} \langle T \rangle \langle \pi_T \rangle \text{ ok} \triangleleft \{\}}$$

■ **Figure 17** FEATHERWEIGHT BROOM: Constraint generation rules part 3 (Continuation of Fig. 16)