

A Grainless Semantics for Parallel Programs with Shared Mutable Data

Stephen Brookes

*Department of Computer Science
Carnegie Mellon University
Pittsburgh, USA*

Abstract

We provide a new denotational semantic model, based on “footstep traces”, for parallel programs which share mutable state. The model embodies a classic principle proposed by Dijkstra: processes should be treated independently, with interference occurring only at synchronization points. The result is a model which makes fewer distinctions between programs than traditional trace models, helping to mitigate the combinatorial explosion triggered by interleaving. Indeed, for a sequential or synchronization-free program the footstep trace semantics is equivalent to a non-deterministic state transformation. The new model can be used to validate the soundness of concurrent separation logic, replacing the action trace semantic model used in previous work for that purpose and yielding a conceptually simpler proof. We include some example programs to facilitate comparison with earlier models, and we discuss briefly the relationship with a recent model by John Reynolds in which actions have discernable starts and finishes.

Key words: concurrency, shared memory, granularity, partial correctness, race condition, denotational semantics, logic

1 Outline

It is notoriously difficult to reason about parallel programs, because of the potential for interference or race conditions between concurrent threads. In traditional semantic models a process or thread denotes a set of traces, a trace is a sequence of atomic actions, and parallel composition is interpreted as fair interleaving. As is well known, this can lead to a combinatorial explosion when attempting to prove correctness of a parallel program, because of the number of possible interleavings to be considered.

¹ Email: brookes@cs.cmu.edu

If concurrent threads write to the same variable, without any guarantee of atomicity or exclusivity, the result may be unpredictable. Traditional semantic models usually make a granularity assumption about the atomicity of certain primitives at the hardware level. Assuming that reads and writes to integer-valued variables are atomic leads to a trace semantics with larger atomic steps and fewer interleavings, but this assumption is not realistic in practice because an integer value may not fit inside a single machine word. On the other hand, the slightly more realistic assumption that word-level reads and writes are atomic will yield more interleavings, tending to exacerbate the combinatorial problem, and the results may depend on word size. Instead of relying on possibly unfounded expectations about hardware implementation we would prefer to be able safely to abstract away from granularity while simultaneously reducing or avoiding the interleaving problem.

To an extent this is possible for simple shared-memory programs (without pointers), provided programs are designed in a disciplined manner. For such programs the “critical” variables can be detected by a static analysis. In the resource-based approach proposed by Hoare[12], Brinch Hansen[2,5] and others, the programmer must partition the critical variables among named resources, and access to a shared variable is only allowed inside the body of a critical region. Since resources are assumed to be implemented as semaphores, this ensures mutually exclusive access to critical variables and disallows programs with “racy” behavior whose execution results might depend on granularity. Owicki and Gries introduced a Hoare-style logic that reflects this discipline, with rules for parallel composition, regions and resources[15]. This provides a methodology for safe reasoning about partial correctness without dependence on granularity.

Unfortunately this methodology is unsuitable for parallel programs that use pointers, because the possibility of aliasing renders purely static detection of race conditions impossible, and the Owicki-Gries proof system is unsound in the presence of aliasing. Recently, O’Hearn[13] has proposed an adaptation of the Owicki-Gries logic to handle pointers, incorporating ideas from separation logic[16]. The author has developed a semantic model (using *action traces*) and used it to prove soundness of this concurrent separation logic[7]. The action traces semantic model interprets parallel composition as a resource-sensitive form of fair interleaving, and treats a potential race condition as a catastrophe, as suggested by John Reynolds. The soundness proof shows that a provable program is race-free.

The above work (including both logics [15,13] and semantics [7]) was designed with an eye to one of the classic principles of concurrent program design, originally articulated by Dijkstra [10]:

... processes should be loosely connected; ... apart from the (rare) moments of explicit intercommunication, ... processes are to be regarded as completely independent of each other.

The action traces model supports the definition of a “local enabling” relation that embodies this principle, allowing a characterization of process behavior in an environment that interferes only via synchronization. However, this model is unnecessarily fine-grained, containing traces that represent executions in less well-behaved environments. As a result the model distinguishes between some pairs of commands which satisfy exactly the same concurrent separation logic formulas. In fact the model supports compositional analysis of safety and liveness properties, in addition to partial and total correctness, because it retains information about the sequences of atomic state changes that occur along a computation. This level of detail is irrelevant for reasoning about partial and total correctness.

In this paper we develop a more abstract semantics, tailored specifically for reasoning about race-free partial and total correctness. Instead of action traces built from resource actions and atomic state actions such as reads and writes to individual variables, we work with *footstep traces* built from resource actions and (compressed sequences of) state actions, representing the cumulative effect of a sequence of state changes executed without interference. We restrict the structure of traces to reflect the assumption that interference can only occur within a critical region. The typical trace will therefore consist of an alternating sequence of state actions and resource actions, with “external” state changes only occurring when a resource is acquired, so that the new semantic model truly embodies Dijkstra’s principle.

This new semantics can also be used instead of action traces to establish the soundness of concurrent separation logic, providing a more streamlined soundness proof based on a more succinct semantics. Moreover, the footstep trace set of a command is independent of any assumption about the granularity of atomic actions. For a well designed concurrent program whose synchronization points are few and far between, our model requires fewer interleavings, helping to mitigate the combinatorial explosion.

To facilitate comparison with the action trace semantics of [7], and to retain as much common ground with that model as is still relevant, our technical development of the footstep trace semantics will echo the main concepts and definitions from the action trace setting, reformulating them as needed in terms of footsteps. (We also re-examine some of the program examples from [7] to highlight the advantages of the new semantics.)

The footstep trace model has been influenced by ideas of John Reynolds, who recently presented a model with similar aims. Reynolds’ model[17] also incorporates some ideas from action trace semantics, so it is not surprising that there are points of similarity. In Reynolds’ approach state actions such as reads and writes have a separate start and finish, leading to a trace model which makes more distinctions between processes than footstep traces; Reynolds then introduces an abstraction function that collapses a trace set by ignoring unnecessary distinctions between traces, based for instance on the relative orderings of independent steps. In contrast, our approach abstracts away from

duration, and we coalesce adjacent footsteps to avoid such distinctions, so that the collapsing is built in implicitly. For a command with no free resource names, such as a sequential command, or a parallel command with no need for synchronization, the footprint trace semantics degenerates into a form clearly equivalent to a (non-deterministic) state transformation. Thus our semantics achieves one of the desirable properties mentioned by Reynolds [17].

2 Syntax

Definition 2.1

The syntax for *commands* (ranged over by c) is given by the following abstract grammar, in which r ranges over *resource names*, i over *identifiers*, e over *integer expressions*, b over *boolean expressions*, and E over *list expressions*:

$$\begin{aligned} c ::= & \text{skip} \mid i := e \mid i := [e] \mid [e] := e' \mid i := \mathbf{cons}(E) \mid \mathbf{dispose} e \mid \\ & c_1; c_2 \mid c_1 \parallel c_2 \mid \mathbf{if} b \mathbf{then} c_1 \mathbf{else} c_2 \mid \mathbf{while} b \mathbf{do} c \mid \\ & \mathbf{local} i = e \mathbf{in} c \mid \mathbf{resource} r \mathbf{in} c \mid \mathbf{with} r \mathbf{when} b \mathbf{do} c \end{aligned}$$

An *allocation* $i := \mathbf{cons}(E)$ allocates a series of fresh heap cells; a *lookup* $i := [e]$ reads a heap cell and assigns its contents to i ; an *update* $[e] := e'$ changes the contents of a heap cell; and a *disposal* $\mathbf{dispose}(e)$ de-allocates a heap cell. Storage management is assumed to be handled globally.

A command of form $\mathbf{local} i = e \mathbf{in} c$ introduces a local variable named i , initialized to the value of e , with scope c .

A *resource block* $\mathbf{resource} r \mathbf{in} c$ introduces a local resource name r , with scope c . A command of form $\mathbf{with} r \mathbf{when} b \mathbf{do} c$ is a *conditional critical region* for resource r . A process attempting to enter a region must wait until the resource is available, acquire the resource and evaluate b : if b is **true** the process executes c then releases the resource; if b is **false** the process releases the resource and waits to try again. A resource can only be held by one process at a time. We use the abbreviation $\mathbf{with} r \mathbf{do} c$ when b is **true**.

Let $\mathbf{free}(c)$ be the set of identifiers occurring free in c , with similar notation for expressions. Let $\mathbf{reads}(c)$ be the set of identifiers having a free read occurrence in c , $\mathbf{writes}(c)$ be the set of identifiers having a free write occurrence in c ,

Definition 2.2

Let $\mathbf{res}(c)$ be the set of resource names occurring free in c , defined by the obvious structural induction. In particular,

$$\begin{aligned} \mathbf{res}(c_1 \parallel c_2) &= \mathbf{res}(c_1) \cup \mathbf{res}(c_2) \\ \mathbf{res}(\mathbf{with} r \mathbf{when} b \mathbf{do} c) &= \mathbf{res}(c) \cup \{r\} \\ \mathbf{res}(\mathbf{resource} r \mathbf{in} c) &= \mathbf{res}(c) - \{r\}. \end{aligned}$$

We omit the syntax of expressions, which includes the usual arithmetic operations and boolean connectives. We assume that expressions are “pure”, i.e. the value of an expression does not depend on the heap. A list expression E is a list e_0, \dots, e_n of integer expressions.

3 States

A state is a partial function $\sigma : \mathbf{Var} \rightarrow V_{int}$, where \mathbf{Var} is the disjoint union of \mathbf{Ide} and V_{int} . We use V_{int} for the set of locations (or heap cells), and we make no distinction between locations and integers. \mathbf{Ide} is the set of identifiers.

Let \mathbf{St} be the set of states. We use σ as a meta-variable ranging over \mathbf{St} . The state σ is *finite* if $\mathbf{dom}(\sigma)$ is a finite set. Note that $\mathbf{dom}(\sigma)$ is a subset of \mathbf{Var} , possibly including both identifiers and values.²

We use conventional notation for states, such as $[v_1 : v'_1, \dots, v_k : v'_k]$, where each $v_i \in \mathbf{Var}$ and each $v'_i \in V_{int}$. In particular $[\]$ denotes the empty state. We write $[\sigma \mid v : v']$ for the state that agrees with σ except at v , which it maps to v' . For $X \subseteq \mathbf{Var}$ let $\sigma \upharpoonright X = \{(v, v') \in \sigma \mid v \in X\}$ and $\sigma \setminus X = \{(v, v') \in \sigma \mid v \notin X\}$. We write $\sigma \smile \sigma'$ when $\sigma \upharpoonright \mathbf{dom}(\sigma') = \sigma' \upharpoonright \mathbf{dom}(\sigma)$, i.e. when the union $\sigma \cup \sigma'$ is a valid state. Let $\sigma \perp \sigma'$ when $\mathbf{dom}(\sigma) \cap \mathbf{dom}(\sigma') = \{\}$. Obviously if $\sigma \perp \sigma'$ then also $\sigma \smile \sigma'$. In the special case when $\sigma \perp \sigma'$ we may write $\sigma \cdot \sigma'$ for $\sigma \cup \sigma'$.

4 Actions

Actions are the building blocks of our semantic model. An action is either a “footstep”, representing a piece of state change, or a resource action representing an attempt to acquire a resource, or the releasing of a resource. We include both “normal” footsteps and “abnormal” footsteps representing a runtime error or a non-terminating computation. We use the term “footstep” by analogy with the usage of “footprint” in related work such as [14,13].

Definition 4.1

Let λ range over actions, as described by the following abstract grammar, in which σ, σ' range over states, X ranges over $\mathcal{P}(\mathbf{Var})$, and r ranges over resource names:

$$\lambda ::= (\sigma, \sigma')_X \mid (\sigma, \perp)_X \mid (\sigma, abort) \mid try(r) \mid acq(r) \mid rel(r)$$

The actions dealing with resources are interpreted as in the earlier semantic models: $try(r)$ represents an unsuccessful attempt to acquire r , $acq(r)$ represents a successful acquisition, and $rel(r)$ releases the resource.

A (normal) footstep has the form $(\sigma, \sigma')_X$ where σ and σ' are finite states and X is a subset of $\{v \in \mathbf{dom}(\sigma) \cap \mathbf{dom}(\sigma') \mid \sigma(v) = \sigma'(v)\}$. The set X represents the variables deemed to be “read-only” in the step, and σ' is to be

² It would be equally reasonable to present the state as a pair (s, h) where $s : \mathbf{Ide} \rightarrow V_{int}$ is the “store” and $h : V_{int} \rightarrow V_{int}$ is the “heap”.

regarded as all that is left of σ after the step, together with any new state introduced by allocations performed in the step. (Note that $X \subseteq \mathbf{Var}$ so X may include heap cells as well as identifiers.) Such a step represents the footprint of a command that “reads” σ and “writes” $\sigma' \setminus X$. The state portion $\sigma \upharpoonright X$ is read-only, and survives the step unchanged; indeed, the side-condition on X can be rephrased equivalently as $\sigma' \upharpoonright X = \sigma \upharpoonright X$.

A footstep of the form (σ, abort) , where σ is a (finite) state, represents a runtime error such as a race condition, or an attempt to dispose a location that is not in use. Such steps will be treated as catastrophic, so there is no need to record a read-only set here.

Finally a footstep of the form $(\sigma, \perp)_X$, where σ is a state, represents a non-terminating computation in which the variables in X are read-only and the variables in $\text{dom}(\sigma)$ are read/written or written. We include a read-only set in order to be able to predict race conditions accurately.

Definition 4.2

For a footstep $\lambda = (\sigma, \sigma')_X$ let $\text{dom}(\lambda) = \text{dom}(\sigma) \cup \text{dom}(\sigma')$, and $\text{mod}(\lambda) = \text{dom}(\lambda) - X$; this is the set of variables which are modified by λ . We also define $\text{allocates}(\lambda) = \text{dom}(\sigma') - \text{dom}(\sigma)$ and $\text{disposals}(\lambda) = \text{dom}(\sigma) - \text{dom}(\sigma')$.

For a non-terminating footstep $\lambda = (\sigma, \perp)_X$ we let $\text{mod}(\lambda) = \text{dom}(\sigma) - X$. Again this represents the set of identifiers assigned to or heap cells updated by the step. We need this set in order to keep a proper account of race conditions. There is no need, however, to define $\text{mod}(\lambda)$ when λ is an abort step.

To illustrate the above concepts and definitions we provide the following examples.

Example 4.3

Let $v \in \mathbf{Var}$ and $v', v'' \in V_{int}$.

- $([v : v'], [v : v''])_{\{v\}}$ is only well-formed if $v' = v''$, in which case this footstep represents a “read” of v with result v' .
- $([v : v'], [v : v''])_{\emptyset}$ is always well-formed, and represents a “read-write” of v ; note that even if $v' = v''$ we distinguish this from a pure read.
- $([v' : v''], [])_{\emptyset}$ represents a disposal of v' .
- $([], [v' : v''])_{\emptyset}$ represents an allocation of v' that initializes the value stored in this cell to v'' .

Example 4.4

The footstep $([x : 0, y : 1, z : 2, 1 : 42], [x : 1, y : 1, z : 99, 99 : 0])_{\{y\}}$ represents a state change that updates the value of x , disposes the heap cell denoted by y , and sets z to a fresh heap cell initialized to contain 0. Notice that the variable y is not itself updated, as indicated by the read-only set annotation.

5 Traces

We build traces out of finite or infinite sequences of actions. Since each step specifies just enough state to enable it, we need to take care when determining which sequences of steps make sense, bearing in mind our assumptions about the interactions between a command and its environment. With respect to resource actions, this means that each resource is held by at most one process at all times, so the acquire and release actions for each resource name along a trace will alternate (exactly as in the action traces model). More crucially, in line with the Dijkstra principle, we will build traces in which adjacent footsteps are always combined into a single step, abstracting away from the sequence of atomic actions represented by the footsteps. This is a radical departure from the action traces model and we make this design choice so that we ignore intermediate states irrelevant for reasoning about race-freedom and partial correctness.³ The technical details concerning this issue will be made precise in the forthcoming characterization of “catenable” sequences of actions.

We will use α, β to range over traces. Let \mathbf{Tr} be the set of traces.

Catenable actions

We characterize the conditions under which one footstep can be executed after another: λ_1 can follow λ_0 if the state needed for λ_1 is consistent with the effect of λ_0 , and λ_1 does not allocate any location that was allocated by λ_0 . When this holds we also say that the sequence $\lambda_0\lambda_1$ is *catenable*. Accordingly, we will build traces by concatenating sequences of actions which obey this constraint.

Definition 5.1

Let $\lambda_0 = (\sigma_0, \sigma'_0)_{X_0}$ and $\lambda_1 = (\sigma_1, \sigma'_1)_{X_1}$ be footsteps. We say that λ_1 *follows* λ_0 , written $\lambda_0 \succ \lambda_1$, if

- $\sigma'_0 \smile \sigma_1$, i.e. $\sigma'_0 \upharpoonright \text{dom}(\sigma_1) = \sigma_1 \upharpoonright \text{dom}(\sigma'_0)$
- $(\sigma'_0 \setminus \text{dom}(\sigma_1)) \perp \sigma'_1$

The first requirement here is that the state resulting from λ_0 must be consistent with the state needed for λ_1 . The second requirement can be rewritten equivalently as $\text{dom}(\sigma'_0) \cap \text{dom}(\sigma'_1) \subseteq \text{dom}(\sigma_1)$, which is the same as requiring $\text{allocates}(\lambda_1) \cap \text{dom}(\sigma'_0) = \{\}$.

Obviously this is not a symmetric relationship: $\lambda_0 \succ \lambda_1$ does not generally imply that $\lambda_1 \succ \lambda_0$.

The read-only sets play no role in determining the follows relation. They do contribute, however, in defining the combined effect of catenable actions.

³ This coalescing of adjacent steps is reminiscent of *mumbling*[6], but we work exclusively with “maximally mumbled” traces rather than including all possible mumbblings of a given sequence, since this results in a more succinct model.

When λ_1 follows λ_0 we can construct a single footstep, denoted $\lambda_0; \lambda_1$, that represents the combined effects of the successive steps.

Definition 5.2

Let $\lambda_0 = (\sigma_0, \sigma'_0)_{X_0}$ and $\lambda_1 = (\sigma_1, \sigma'_1)_{X_1}$ be footsteps such that $\lambda_0 \asymp \lambda_1$. We define $\lambda_0; \lambda_1$ to be the following footstep:

- if $(\sigma_1 \setminus \text{dom}(\sigma'_0)) \perp \sigma_0$ then $\text{disposals}(\lambda_0) \cap \text{dom}(\sigma_1) = \{\}$ and we let

$$(\sigma_0, \sigma'_0)_{X_0}; (\sigma_1, \sigma'_1)_{X_1} = (\sigma_0 \cdot (\sigma_1 \setminus \text{dom}(\sigma'_0)), \sigma'_1 \cdot (\sigma'_0 \setminus \text{dom}(\sigma_1)))_X$$

where $X = (X_0 - \text{mod}(\lambda_1)) \cup (X_1 - \text{mod}(\lambda_0))$.

- if $(\sigma_1 \setminus \text{dom}(\sigma'_0)) \not\perp \sigma_0$ then λ_0 disposes some cell needed by λ_1 , so we let

$$(\sigma_0, \sigma'_0)_{X_0}; (\sigma_1, \sigma'_1)_{X_1} = (\sigma_0 \cdot (\sigma_1 \setminus \text{dom}(\sigma'_0)), \text{abort}).$$

We extend the catenability relation \asymp and the sequencing operator $\lambda_0; \lambda_1$ to more general pairs of actions in the obvious way. For abort steps we make the following definitions:

- $(\sigma_0, \sigma'_0)_{X_0} \asymp (\sigma_1, \text{abort})$ if $\sigma'_0 \smile \sigma_1$ and $(\sigma_1 \setminus \text{dom}(\sigma'_0)) \perp \sigma_0$
 $(\sigma_0, \sigma'_0)_{X_0}; (\sigma_1, \text{abort}) = (\sigma_0 \cdot (\sigma_1 \setminus \text{dom}(\sigma'_0)), \text{abort})$
- $(\sigma_0, \text{abort}) \asymp \lambda$ for all actions λ
 $(\sigma_0, \text{abort}); \lambda = (\sigma_0, \text{abort})$

The corresponding definitions involving a non-terminating step are similar:

- Let $\lambda_0 = (\sigma_0, \sigma'_0)_{X_0}$ and $\lambda_1 = (\sigma_1, \perp)_{X_1}$.
 $(\sigma_0, \sigma'_0)_{X_0} \asymp (\sigma_1, \perp)_{X_1}$ if $\sigma'_0 \smile \sigma_1$ and $(\sigma_1 \setminus \text{dom}(\sigma'_0)) \perp \sigma_0$
 $(\sigma_0, \sigma'_0)_{X_0}; (\sigma_1, \perp)_{X_1} = (\sigma_0 \cdot (\sigma_1 \setminus \text{dom}(\sigma'_0)), \perp)$,
 where $X = (X_0 - \text{mod}(\lambda_1)) \cup (X_1 - \text{mod}(\lambda_0))$.
- $(\sigma_0, \perp)_{X_0} \asymp \lambda$ for all actions λ
 $(\sigma_0, \perp)_{X_0}; \lambda = (\sigma_0, \perp)_{X_0}$.

We allow any action λ to be followed by (or to follow) any resource action, so that for all resource names r , $\lambda \asymp \text{try}(r)$, $\lambda \asymp \text{acq}(r)$, $\lambda \asymp \text{rel}(r)$, $\text{try}(r) \asymp \lambda$, and so on.

Remark 5.3

Let δ be the footstep $([\], [\])_{\{\}}$. It is easy to check that, for all footsteps λ_0 and λ_1 , we have $\lambda_0 \asymp \delta$, $\delta \asymp \lambda_1$, and $\lambda_0; \delta = \lambda_0$, $\delta; \lambda_1 = \lambda_1$.

The following examples illustrate the above definitions, showing how the read-only information of successive steps gets combined, and motivating the carefully chosen side conditions in the previous development.

Example 5.4

Two consecutive reads of the same variable make sense provided they yield the same value; the result is again a read:

$$([x : v], [x : v])_{\{x\}} \asymp ([x : v'], [x : v'])_{\{x\}} \text{ if and only if } v = v'$$

$$([x : v], [x : v])_{\{x\}}; ([x : v], [x : v])_{\{x\}} = ([x : v], [x : v])_{\{x\}}$$

Example 5.5

A read followed by a write to the same variable is sensible when the read yields the same value as is used to start the write; the result is a write:

$$([x : v], [x : v])_{\{x\}} \asymp ([x : v'], [x : v'])_{\{x\}} \text{ if and only if } v = v'$$

$$([x : v], [x : v])_{\{x\}}; ([x : v], [x : v'])_{\{x\}} = ([x : v], [x : v'])_{\{x\}}$$

Example 5.6

Reads of distinct variables can be concatenated; their effects commute, and the result is a single read-only step.

$$([x : v], [x : v])_{\{x\}} \asymp ([y : v'], [y : v'])_{\{y\}} \text{ when } x \neq y$$

$$([x : v], [x : v])_{\{x\}}; ([y : v'], [y : v'])_{\{y\}} = ([x : v, y : v'], [x : v, y : v'])_{\{x, y\}}$$

Example 5.7

Writes to distinct variables can be concatenated; their effects commute; the result is a single action representing the combined reads.

$$([x : v_0], [x : v'_0])_{\{x\}} \asymp ([y : v_1], [y : v'_1])_{\{y\}} \text{ when } x \neq y$$

$$([x : v_0], [x : v'_0])_{\{x\}}; ([y : v_1], [y : v'_1])_{\{y\}} = ([x : v_0, y : v_1], [x : v'_0, y : v'_1])_{\{x, y\}}$$

Example 5.8

A single disposal can always be followed by another action. If the second step is another disposal for the same variable an error occurs (representing a dangling pointer).

$$([v : v'], [])_{\{v\}} \asymp (\sigma, \sigma')_X \text{ always}$$

$$([v : v'], [])_{\{v\}}; (\sigma, \sigma')_X = (\sigma \cup [v : v'], \sigma')_{X \setminus v} \text{ if } v \notin \text{dom}(\sigma)$$

$$([v : v'], [])_{\{v\}}; (\sigma, \sigma')_X = ([\sigma \mid v : v'], \text{abort}) \text{ if } v \in \text{dom}(\sigma)$$

Example 5.9

Since storage allocation is assumed to be managed globally, so that successive calls always return distinct heap cells, we need not allow consecutive allocations of the heap cell. Indeed, according to the above definitions we have $([], [v : v'])_{\{v\}} \not\asymp ([], [v : v''])_{\{v\}}$.

Example 5.10

A step that allocates a heap cell v can follow any step in which v was not read, written, or allocated, again because the storage manager always allocates “fresh” cells. We have $(\sigma, \sigma')_X \asymp ([], [v : v'])_{\{v\}}$ if $v \notin \text{dom}(\sigma')$, in which case $(\sigma, \sigma')_X; ([], [v : v'])_{\{v\}} = (\sigma, \sigma' \cup [v : v'])_{X \setminus v}$.

Catenable sequences

If $\lambda_0 \asymp \lambda_1$ and $\lambda_1 \asymp \lambda_2$, it does not necessarily follow that $(\lambda_0; \lambda_1) \asymp \lambda_2$. For example, consider the following footsteps:

$$\lambda_0 = ([v : 0], [])_{\{v\}} \quad \lambda_1 = ([y : 0], [y : 0])_{\{y\}} \quad \lambda_2 = ([v : 1], [])_{\{v\}}.$$

We have $\lambda_0 \asymp \lambda_1$ and $\lambda_1 \asymp \lambda_2$, but $\lambda_0; \lambda_1 = ([v : 0, y : 0], [y : 0])_{\{v, y\}}$, so that $(\lambda_0; \lambda_1) \not\asymp \lambda_2$. Indeed, the sequence of actions $\lambda_0 \lambda_1 \lambda_2$ clearly cannot be executed without the help of interference, because of the discrepancy in the

value of v between the first and third steps. These considerations lead to the following characterization of *catenable* sequences of actions.

Definition 5.11

The sequence of steps $\lambda_0\lambda_1\dots\lambda_n\dots$ is *catenable* if and only if $\lambda_0 \asymp \lambda_1$, $(\lambda_0; \lambda_1) \asymp \lambda_2$, $(\lambda_0; \lambda_1; \lambda_2) \asymp \lambda_3$, and so on. For such a sequence we let $cat(\lambda_0\dots\lambda_n\dots) = \lambda_0; \lambda_1; \dots; \lambda_n\dots$ be the result of catenating its steps.

We want to streamline our semantics so that we only include traces that reflect our underlying assumptions about interference. In particular, we only need to allow for the possibility of interference inside critical regions. So we extend the catenability relation to finite traces and footsteps as follows, by induction on trace length.

Definition 5.12

For all finite traces α , resource names r , and footsteps λ, λ'

$$\begin{aligned} \alpha\lambda \asymp \lambda' & \quad \text{iff } \lambda \asymp \lambda' \\ \alpha \text{ try}(r) \asymp \lambda' & \quad \text{iff } \alpha \asymp \lambda' \\ \alpha \text{ acq}(r) \asymp \lambda' & \quad \text{always} \\ \alpha \text{ rel}(r) \asymp \lambda' & \quad \text{iff } \alpha \asymp \lambda' \end{aligned}$$

Note that the step following a resource acquisition is not constrained here, allowing for the possibility of interference at synchronization.

We extend the catenation operation analogously, so that for a finite trace β and an action λ such that $\beta \asymp \lambda$ we obtain a definition of the trace $\beta; \lambda$.

Definition 5.13

For all finite traces α , resource names r , and footsteps λ, λ'

$$\begin{aligned} (\alpha\lambda); \lambda' & = \alpha(\lambda; \lambda') & \quad \text{when } \lambda \asymp \lambda' \\ (\alpha \text{ try}(r)); \lambda & = \alpha \text{ try}(r) \lambda' & \quad \text{when } \alpha \asymp \lambda' \\ (\alpha \text{ acq}(r)); \lambda' & = \alpha \text{ acq}(r) \lambda' & \quad \text{always} \\ (\alpha \text{ rel}(r)); \lambda' & = \alpha \text{ rel}(r) \lambda' & \quad \text{when } \alpha \asymp \lambda' \end{aligned}$$

Using these definitions we can characterize those finite sequences of actions satisfying the catenability constraints. This notion of catenability extends to an infinite sequence of actions in the obvious way: an infinite sequence is catenable if and only if each of its finite prefixes is catenable. However, the result of combining an infinite catenable sequence $\lambda_n = (\sigma_n, \sigma'_n)_{X_n}$ needs to be defined carefully. There are two cases:

- if for each n we have $\lambda_0; \dots; \lambda_n = (\tau_n, \tau'_n)_{Y_n}$ then the infinite sequence represents a non-terminating computation that will be enabled by the state $\bigcup_{n=0}^{\infty} \tau_n$. The read-only variables for this entire computation can be taken

to be $Y = \{v \in \bigcup_{n=0}^{\infty} X_n \mid \forall n. (v \in X_n \vee v \notin \text{dom}(\lambda_n))\}$. We therefore define $\lambda_0; \dots; \lambda_n; \dots = (\bigcup_{n=0}^{\infty} \tau_n, \perp)_Y$.

- if for some n we have $\lambda_0; \dots; \lambda_n = (\tau_n, \text{abort})$ then this will also be the case for all larger n , and we define $\lambda_0; \dots; \lambda_n; \dots$ to be (τ_n, abort) also.

The state appearing in a footstep constructed in this manner may well be infinite, even if all of the individual steps λ_n involve finite pieces of state. This has a natural computational interpretation: the process executing the step needs to access an infinite portion of the heap during its run. Of course this can only happen in an infinite computation, since each finite prefix of the computation trace affects only a finite portion of the state. The following example exhibits this behavior.

Example 5.14

Consider the sequence of steps $\lambda_n = ([x : n, n : 0], [x : n + 1])_{\{\}}^{\cup}$, for $n \geq 0$; we have, using the above definition,

$$\lambda_0; \lambda_1; \dots; \lambda_n; \dots = ([x : 0, 0 : 0, 1 : 0, \dots, n : 0, \dots], \perp)_{\{\}}^{\cup}.$$

This footstep represents a possible behavior of the command

while true do (dispose(x); $x := x + 1$).

So far we have characterized all cases where a single action λ may follow a trace α . It is easy to extend the notion of catenability, and the concatenation operator, to pairs of traces. When α and β are traces we write $\alpha \succ \beta$ to mean that β follows α , and we write $\alpha; \beta$ for the trace obtained by concatenating. We then define concatenation and iteration for sets of traces in the usual way, with adjustments so as to include only catenable cases:

- $T_0; T_1 = \{\alpha_0; \alpha_1 \mid \alpha_0 \in T_0 \ \& \ \alpha_1 \in T_1 \ \& \ \alpha_0 \succ \alpha_1\}$.
- $T^* = \bigcup_{n=0}^{\infty} T^n$, where $T^0 = \{\delta\}$ and $T^{n+1} = T^n; T$
- $T^\omega = \{\alpha_0; \alpha_1; \dots \alpha_n; \dots \mid \forall n \geq 0. \alpha_n \in T \ \& \ \alpha_0; \dots; \alpha_n \succ \alpha_{n+1}\}$.

Recall that δ is the trivial footstep, $([\], [\])_{\{\}}^{\cup}$.

Parallel composition

The behavior of a process running in parallel with other threads depends on the resources held by the process and is constrained by the resources being held by its environment. These sets of resources start empty and will always stay disjoint, given our assumptions about the implementation of resources via semaphores. Accordingly we define for each action λ a *resource enabling* relation $(A_1, A_2) \xrightarrow{\lambda} (A'_1, A'_2)$ on disjoint pairs of resource sets, to specify when a process holding resources A_1 , in an environment that holds A_2 , can perform this action, and the action's effect on resources. Although the special case where A_1 and A_2 are both empty deserves special attention, since it corresponds to the typical initial assumptions, the more general case lends itself better to inductive characterization and is a key ingredient in enforcing mutual

exclusion constraints on resource acquisition in the forthcoming definition of parallel composition.

Definition 5.15

Let A_1 and A_2 be disjoint sets of resource names. The resource enabling relations $\xrightarrow{\lambda}$ are given by:

$$\begin{aligned} (A_1, A_2) &\xrightarrow{\text{try}(r)} (A_1, A_2) && \text{always} \\ (A_1, A_2) &\xrightarrow{\text{acq}(r)} (A_1 \cup \{r\}, A_2) && \text{iff } r \notin A_1 \cup A_2 \\ (A_1, A_2) &\xrightarrow{\text{rel}(r)} (A_1 - \{r\}, A_2) && \text{iff } r \in A_1 \\ (A_1, A_2) &\xrightarrow{\lambda} (A_1, A_2) && \text{always, if } \lambda \text{ is not a resource action} \end{aligned}$$

This generalizes in the obvious way to a sequence of actions, and we write $(A_1, A_2) \xrightarrow{\alpha} \cdot$ to indicate that a process holding resources A_1 in an environment holding A_2 can perform the trace α .

We can characterize those pairs of actions that interfere as follows.

Definition 5.16

Let $\lambda_0 = (\sigma_0, \sigma'_0)_{X_0}$ and $\lambda_1 = (\sigma_1, \sigma'_1)_{X_1}$. We say that λ_0 and λ_1 interfere, written $\lambda_0 \bowtie \lambda_1$, if

- $\sigma_0 \smile \sigma_1$, i.e. $\sigma_0 \upharpoonright \text{dom}(\sigma_1) = \sigma_1 \upharpoonright \text{dom}(\sigma_0)$
- $\text{dom}(\sigma_0) \cap \text{dom}(\sigma_1) \not\subseteq X_0 \cap X_1$

When $\lambda_0 \bowtie \lambda_1$ define $\lambda_0 \bullet \lambda_1 = (\sigma_0 \cup \sigma_1, \text{abort})$.

If the overlap is on variables which are read-only for both footsteps, there is no race so we do not count this as an interfering case.

For abnormal footsteps we characterize the interfering cases as follows:

Definition 5.17

When λ_0 is (σ_0, abort) and/or λ_1 is (σ_1, abort) we specify that $\lambda_0 \bowtie \lambda_1$ holds if $\sigma_0 \smile \sigma_1$. And if λ_0 is $(\sigma_0, \perp)_{X_0}$ and λ_1 is $(\sigma_1, \sigma'_1)_{X_1}$ we specify that $\lambda_0 \bowtie \lambda_1$ holds if $\sigma_0 \smile \sigma_1$ and $\text{dom}(\sigma_0) \cap \text{dom}(\sigma_1) \not\subseteq X_0 \cap X_1$. The cases for $(\sigma_0, \perp)_{X_0}$ and $(\sigma_1, \perp)_{X_1}$, and for $(\sigma_0, \perp)_{X_0}$ and (σ_1, abort) , are analogous.

We now define, for each pair (A_0, A_1) of disjoint resource sets and each pair (α_0, α_1) of action sequences, the set $\alpha_{0A_0} \parallel_{A_1} \alpha_1$ of all *mutex fairmerges* of α_0 using A_0 with α_1 using A_1 . This is the set of all fair interleavings of α_0 with α_1 in which each step respects the resource constraints and a race condition causes an error. The definition for finite traces is inductive in the product of the lengths of the traces, and we include the empty sequence ϵ , to allow a simpler formulation: the base case is when one of the traces is empty. Note that the set $\alpha_{0A_0} \parallel_{A_1} \alpha_1$ may be empty: for instance when α_0 is the empty sequence but α_1 is not permitted because it needs a resource belonging to A_0 . Again, although the case where both resource sets is empty plays a special role in the sequel, the more general case has a naturally inductive formulation.

Definition 5.18 For finite traces α and β , and disjoint resource sets A_0, A_1 , the set of traces $\alpha_{A_0} \parallel_{A_1} \beta$ is given inductively by:

$$\begin{aligned}
 \alpha_{A_0} \parallel_{A_1} \epsilon &= \{ \alpha \mid (A_0, A_1) \xrightarrow{\alpha} \cdot \} \\
 \epsilon_{A_0} \parallel_{A_1} \beta &= \{ \beta \mid (A_1, A_0) \xrightarrow{\beta} \cdot \} \\
 (\lambda_0 \alpha_0)_{A_0} \parallel_{A_1} (\lambda_1 \alpha_1) \\
 &= \{ \lambda_0 \bullet \lambda_1 \} \quad \text{if } \lambda_0 \bowtie \lambda_1 \\
 &= \{ \lambda_0; \beta \mid (A_0, A_1) \xrightarrow{\lambda_0} (A'_0, A_1) \ \& \ \beta \in \alpha_{0A'_0} \parallel_{A_1} (\lambda_1 \alpha_1) \ \& \ \lambda_0 \asymp \beta \} \\
 &\quad \cup \{ \lambda_1; \beta \mid (A_1, A_0) \xrightarrow{\lambda_1} (A'_1, A_0) \ \& \ \beta \in (\lambda_0 \alpha_0)_{A_0} \parallel_{A'_1} \alpha_1 \ \& \ \lambda_1 \asymp \beta \} \\
 &\quad \text{if } \neg(\lambda_0 \bowtie \lambda_1)
 \end{aligned}$$

Again we remark that the definition of parallel composition builds in the resource constraints, only includes catenable cases, and combines adjacent footsteps. The definition of the set of mutex fairmerges extends to infinite traces in the usual way.

The fairmerge definition for pairs of traces extends in the obvious pointwise way to sets of traces: when T_0 and T_1 are sets of traces we let $T_{0A_0} \parallel_{A_1} T_1$ be the union of the sets $\alpha_{0A_0} \parallel_{A_1} \alpha_1$ as α_0 ranges over T_0 and α_1 ranges over T_1 . We will be mainly concerned with the case where $A_0 = A_1 = \{\}$, representing what happens when two processes start up with no initial resources. Indeed, we may abbreviate $T_{0\{\}} \parallel_{\{\}} T_1$ as $T_0 \parallel T_1$, omitting the empty subscripts.

6 Semantics of expressions

Let $Tag = \mathcal{P}(\mathbf{Var})$ be the set of *read-only sets* used to annotate footsteps. Recall that \mathbf{Var} is the union of \mathbf{Ide} and V_{int} . For integer expressions e we define $\llbracket e \rrbracket \subseteq (\mathbf{St} \times V_{int})_{Tag}$. Similarly for boolean expressions we define $\llbracket b \rrbracket \subseteq (\mathbf{St} \times V_{bool})_{Tag}$. We can also give a semantics for list expressions E along similar lines, so that the value of a list expression is a list of integer values: $\llbracket E \rrbracket \subseteq (\mathbf{St} \times V_{int}^*)_{Tag}$. We omit the details, which are straightforward.⁴

Our semantic clauses for expressions can be obtained from the corresponding clauses in the Reynolds semantics[17] by inserting and propagating the relevant information about read-only variables.

Definition 6.1 The semantic functions

$$\begin{aligned}
 \llbracket - \rrbracket &: \mathbf{Exp}_{int} \rightarrow \mathcal{P}((\mathbf{St} \times V_{int})_{Tag}) \\
 \llbracket - \rrbracket &: \mathbf{Exp}_{bool} \rightarrow \mathcal{P}((\mathbf{St} \times V_{bool})_{Tag})
 \end{aligned}$$

⁴ It would also be easy to incorporate into this semantic model “impure” expressions such as $[e]$, whose value depends on the heap.

are defined inductively by the following clauses:

$$\begin{aligned}
 \llbracket n \rrbracket &= \{([\], n)_{\{\}}\} \\
 \llbracket i \rrbracket &= \{([i : v], v)_{\{i\}} \mid v \in V_{int}\} \\
 \llbracket e_1 + e_2 \rrbracket &= \{(\sigma_1 \cup \sigma_2, v_1 + v_2)_{X_1 \cup X_2} \mid \\
 &\quad (\sigma_1, v_1)_{X_1} \in \llbracket e_1 \rrbracket \ \& \ (\sigma_2, v_2)_{X_2} \in \llbracket e_2 \rrbracket \ \& \ \sigma_1 \smile \sigma_2\} \\
 \llbracket \mathbf{true} \rrbracket &= \{([\], true)_{\{\}}\} \\
 \llbracket e_1 = e_2 \rrbracket &= \{(\sigma_1 \cup \sigma_2, v_1 = v_2)_{X_1 \cup X_2} \mid \\
 &\quad (\sigma_1, v_1)_{X_1} \in \llbracket e_1 \rrbracket \ \& \ (\sigma_2, v_2)_{X_2} \in \llbracket e_2 \rrbracket \ \& \ \sigma_1 \smile \sigma_2\} \\
 \llbracket e_1 \ \mathbf{and} \ e_2 \rrbracket &= \\
 &\quad \{(\sigma_1, false)_{X_1} \mid (\sigma_1, false)_{X_1} \in \llbracket e_1 \rrbracket\} \\
 &\quad \cup \{(\sigma_1 \cup \sigma_2, t)_{X_1 \cup X_2} \mid (\sigma_1, true)_{X_1} \in \llbracket e_1 \rrbracket \ \& \ (\sigma_2, t)_{X_2} \in \llbracket e_2 \rrbracket \ \& \ \sigma_1 \smile \sigma_2\}
 \end{aligned}$$

Here we have specified that a conjunction b_1 **and** b_2 is evaluated using the usual short-circuit; this detail is not crucial in what follows, and the ensuing development can be modified as needed if we use another interpretation such as strict, left-to-right evaluation.

There is no need to use a more elaborate semantic structure allowing for the possibility of interference during expression evaluation, since the model is intended only to describe executions in an environment that obeys the Dijkstra principle: interference is assumed to occur only on synchronization. It is therefore safe to ignore intermediate states.

A boolean expression gives rise to a pair of sets of single-footstep traces, in the obvious way: we distinguish the traces describing “true” cases from the traces describing “false” cases.

Definition 6.2

For a boolean expression b , let $\llbracket b \rrbracket_{true}, \llbracket b \rrbracket_{false} \subseteq \mathbf{Tr}$ be given by

$$\begin{aligned}
 \llbracket b \rrbracket_{true} &= \{(\sigma, \sigma)_X \mid (\sigma, true)_X \in \llbracket b \rrbracket\} \\
 \llbracket b \rrbracket_{false} &= \{(\sigma, \sigma)_X \mid (\sigma, false)_X \in \llbracket b \rrbracket\}
 \end{aligned}$$

Example 6.3

$$\llbracket x + y \rrbracket = \{(x : v, y : v'), v + v'\}_{x,y} \mid v, v' \in V_{int}\}$$

Example 6.4

$$\begin{aligned}
 \llbracket \mathbf{true} \rrbracket_{true} &= \{([\], [\])\} = \{\delta\} \\
 \llbracket \mathbf{true} \rrbracket_{false} &= \{\}
 \end{aligned}$$

Example 6.5 Let b be the expression $(x = y)$ **and** $(y = z)$.

$$\begin{aligned} \llbracket b \rrbracket_{false} &= \{[x : v, y : v'], [x : v, y : v']\}_{\{x, y\}} \mid v \neq v'\} \\ &\cup \{([x : v, y : v, z : v''], [x : v, y : v, z : v''])_{\{x, y, z\}} \mid v \neq v''\} \\ \llbracket b \rrbracket_{true} &= \{([x : v, y : v, z : v], [x : v, y : v, z : v])_{\{x, y, z\}} \mid v \in V_{int}\} \end{aligned}$$

7 Semantics of commands

For each command c we define $\llbracket c \rrbracket \subseteq \mathbf{Tr}$, by structural induction.

Definition 7.1

The semantic function $\llbracket - \rrbracket : \mathbf{Com} \rightarrow \mathcal{P}(\mathbf{Tr})$ is defined inductively by the following clauses:

$$\begin{aligned} \llbracket \mathbf{skip} \rrbracket &= \{([\], [\])\} \\ \llbracket i := e \rrbracket &= \{(\sigma \cup [i : v], [\sigma \mid i : v'])_{X \setminus i} \mid (\sigma, v')_X \in \llbracket e \rrbracket \ \& \ \sigma \smile [i : v]\} \\ \llbracket [e] := e' \rrbracket &= \{(\sigma \cup \sigma' \cup [v : v_0], [\sigma \cup \sigma' \mid v : v'])_{(X \cup X') \setminus v} \mid \\ &\quad (\sigma, v)_X \in \llbracket e \rrbracket \ \& \ (\sigma', v')_{X'} \in \llbracket e' \rrbracket \ \& \ \sigma \smile \sigma' \ \& \ (\sigma \cup \sigma') \smile [v : v_0]\} \\ \llbracket i := [e] \rrbracket &= \{(\sigma \cup [i : v_0, v : v'], [\sigma \mid i : v', v : v'])_{(X \cup \{v\}) \setminus i} \mid \\ &\quad (\sigma, v)_X \in \llbracket e \rrbracket \ \& \ \sigma \smile [i : v_0, v : v']\} \\ \llbracket i := \mathbf{cons}(E) \rrbracket &= \{(\sigma \cup [i : v], \sigma \cup [i : l, l : v_0, \dots, l + n : v_n])_{X \setminus i} \mid \\ &\quad (\sigma, [v_0, \dots, v_n])_X \in \llbracket E \rrbracket \ \& \ \sigma \smile [i : v] \ \& \ l, l + 1, \dots, l + n \notin \mathbf{dom}(\sigma)\} \\ \llbracket \mathbf{dispose} \ e \rrbracket &= \{(\sigma \cup [v : v'], \sigma \setminus v)_{X \setminus v} \mid (\sigma, v)_X \in \llbracket e \rrbracket \ \& \ \sigma \smile [v : v']\} \\ \llbracket c_1; c_2 \rrbracket &= \{\alpha_1; \alpha_2 \mid \alpha_1 \in \llbracket c_1 \rrbracket \ \& \ \alpha_2 \in \llbracket c_2 \rrbracket \ \& \ \alpha_1 \succ \alpha_2\} = \llbracket c_1 \rrbracket; \llbracket c_2 \rrbracket \\ \llbracket \mathbf{if} \ b \ \mathbf{then} \ c_1 \ \mathbf{else} \ c_2 \rrbracket &= \llbracket b \rrbracket_{true}; \llbracket c_1 \rrbracket \cup \llbracket b \rrbracket_{false}; \llbracket c_2 \rrbracket \\ \llbracket \mathbf{while} \ b \ \mathbf{do} \ c \rrbracket &= (\llbracket b \rrbracket_{true}; \llbracket c \rrbracket)^*; \llbracket b \rrbracket_{false} \cup (\llbracket b \rrbracket_{true}; \llbracket c \rrbracket)^\omega \\ \llbracket c_1 \parallel c_2 \rrbracket &= \bigcup \{\alpha_1 \parallel \alpha_2 \mid \alpha_1 \in \llbracket c_1 \rrbracket \ \& \ \alpha_2 \in \llbracket c_2 \rrbracket\} \\ \llbracket \mathbf{with} \ r \ \mathbf{when} \ b \ \mathbf{do} \ c \rrbracket &= \mathit{wait}^*; \mathit{enter} \cup \mathit{wait}^\omega \\ &\quad \text{where } \mathit{wait} = \mathit{acq}(r) \llbracket b \rrbracket_{false} \mathit{rel}(r) \cup \{\mathit{try}(r)\} \\ &\quad \mathit{enter} = \mathit{acq}(r) \llbracket b \rrbracket_{true}; \llbracket c \rrbracket \mathit{rel}(r) \\ \llbracket \mathbf{resource} \ r \ \mathbf{in} \ c \rrbracket &= \{\mathit{cat}(\alpha \setminus r) \mid \alpha \in \llbracket c \rrbracket_r\} \\ \llbracket \mathbf{local} \ i = e \ \mathbf{in} \ c \rrbracket &= \{(\sigma, \sigma)_X; (\alpha \setminus i) \mid \\ &\quad (\sigma, v)_X \in \llbracket e \rrbracket \ \& \ \alpha \in \llbracket c \rrbracket_{i=v} \ \& \ (\sigma, \sigma)_X \succ (\alpha \setminus i)\} \end{aligned}$$

Note that primitive commands such as **skip**, assignment, update, lookup, allocation, and disposal have traces consisting of a single footstep. The traces

of a conditional critical region contain resource actions. Note the restriction, when building traces of $c_1; c_2$, to catenable combinations; similarly for conditional commands, loops, conditional critical regions, and parallel composition.

The semantic clause for **resource** r **in** c involves traces of c which assume that r is *local*, so that no concurrent process can access r , and that r is initially available. This reflects the scoping rules: the scope of the local resource name is c and excludes any processes running in parallel.

For a trace set T we let T_r be the set of traces in T which are *sequential for* r and consistent with r being initially available; more precisely, these are the traces α in T such that $\alpha \setminus r$ is catenable and $\{\} \xrightarrow{\alpha[r]} \cdot$. We write $\alpha \setminus r$ for the sequence obtained by deleting all actions in α that involve resource name r . The catenability requirement on $\alpha \setminus r$ means that for each subsequence of α of the form $\beta acq(r)\lambda$ we have $\beta \approx \lambda$. This constraint reflects the fact that r is a local resource, so that there can be no interference from outside: no concurrent process can synchronize using the local resource. For a set A of resource names we define

$$\begin{aligned} A &\xrightarrow{try(r)} A && \text{if } r \in A \\ A &\xrightarrow{acq(r)} A \cup \{r\} && \text{if } r \notin A \\ A &\xrightarrow{rel(r)} A - \{r\} && \text{if } r \in A \end{aligned}$$

Thus we have $\{\} \xrightarrow{\alpha \setminus r} \cdot$ if and only if $\alpha \setminus r$ is a prefix of a trace in the set $(acq(r) try(r)^\infty rel(r)^\infty)$. This captures precisely the cases when α is consistent with r being a local resource name which is available for acquisition at the start of the trace.

The semantic clause for **local** $i = e$ **in** c involves a related notion of sequentiality: $\llbracket c \rrbracket_{i=v}$ is the set of traces of c which are *sequential for* i and consistent with the initial value v for i . A trace with these properties represents an execution in which no concurrent process can affect the value of i , and this corresponds to the fact that i is a local variable. For a trace α we define $\alpha \setminus i$ to be the trace obtained by deleting i from all states and read-only tags. When α is catenable and sequential for i , the trace $\alpha \setminus i$ will also be catenable.

There is considerable overlap with the corresponding Reynolds semantic clauses [17], and with action trace semantics [7]. The main differences involve the collapsing of consecutive state changes, book-keeping for read-only tags, and the restriction to catenable traces.

8 Examples

- (i) $\llbracket x := x + 1 \rrbracket = \{([x : v], [x : v + 1])_{\{\}} \mid v \in V_{int}\}$
- (ii) $\llbracket x := x + 1; x := x + 1 \rrbracket = \{([x : v, x : v + 2])_{\{\}} \mid v \in V_{int}\} = \llbracket x := x + 2 \rrbracket$
- (iii) $\llbracket x := x + 1 \parallel x := x + 1 \rrbracket = \{([x : v], abort) \mid v \in V_{int}\}$
- (iv) $\llbracket \text{with } r \text{ do } x := x + 1 \rrbracket = try(r)^* acq(r) \llbracket x := x + 1 \rrbracket rel(r) \cup \{try(r)^\omega\}$

- (v) $\llbracket \mathbf{with} \ r \ \mathbf{do} \ x:=x+1 \parallel \mathbf{with} \ r \ \mathbf{do} \ x:=x+1 \rrbracket$ contains traces of the forms $acq(r) \alpha \ rel(r) \ acq(r) \ \beta \ rel(r)$, $acq(r) \ \alpha \ rel(r) \ try(r)^\omega$, and $try(r)^\omega$, where $\alpha, \beta \in \llbracket x:=x+1 \rrbracket$, as well as traces of similar form containing additional $try(r)$ steps. Only traces of the first kind in which $\alpha \succ \beta$ are sequential for r . It follows that

$$\begin{aligned} & \llbracket \mathbf{resource} \ r \ \mathbf{in} \ (\mathbf{with} \ r \ \mathbf{do} \ x:=x+1 \parallel \mathbf{with} \ r \ \mathbf{do} \ x:=x+1) \rrbracket \\ &= \{ \alpha; \beta \mid \alpha, \beta \in \llbracket x:=x+1 \rrbracket \ \& \ \alpha \succ \beta \} \\ &= \llbracket x:=x+1; x:=x+1 \rrbracket = \llbracket x:=x+2 \rrbracket. \end{aligned}$$

- (vi) The command $[x]:=x+1 \parallel [y]:=y+1$ has traces of the form $([x : v, y : v', v : n, v' : n'], [x : v, y : v', v : v+1, v' : v'+1])_{\{x,y\}}$ where $v \neq v'$, and traces of the form $([x : v, y : v, v : n], abort)$, for all $v, n, v', n' \in V_{int}$.
- (vii) $\llbracket x:=1; y:=2 \rrbracket = \{([x:v, y:v'], [x:1, y:2])_{\{x,y\}} \mid v, v' \in V_{int}\} = \llbracket y:=2; x:=1 \rrbracket$.
- (viii) Let c_1 and c_2 be the commands

$$\begin{aligned} & \mathbf{with} \ r_1 \ \mathbf{do} \ \mathbf{with} \ r_2 \ \mathbf{do} \ x:=1 \\ & \mathbf{with} \ r_2 \ \mathbf{do} \ \mathbf{with} \ r_1 \ \mathbf{do} \ y:=2 \end{aligned}$$

The command $\mathbf{resource} \ r_1, r_2 \ \mathbf{in} \ (c_1 \parallel c_2)$ has the trace set

$$\{([\], \perp)_{\{x,y\}}\} \cup \{([x : v, y : v'], [x : 1, y : 2])_{\{x,y\}} \mid v, v' \in V_{int}\}.$$

The first case represents the potential for deadlock, which occurs if the first thread acquires r_1 and the second acquires r_2 .

- (ix) It is easy to see that

$$\begin{aligned} \llbracket x:=\mathbf{cons}(1) \rrbracket &= \{([x : v_0], [x : v, v : 1])_{\{x,y\}} \mid v_0, v \in V_{int}\} \\ \llbracket y:=\mathbf{cons}(2) \rrbracket &= \{([y : v_1], [y : v', v' : 2])_{\{x,y\}} \mid v_1, v' \in V_{int}\} \end{aligned}$$

It then follows that

$$\begin{aligned} & \llbracket x:=\mathbf{cons}(1) \parallel y:=\mathbf{cons}(2) \rrbracket \\ &= \llbracket x:=\mathbf{cons}(1) \rrbracket \parallel \llbracket y:=\mathbf{cons}(2) \rrbracket \\ &= \{([x : v_0, y : v_1], [x : v, y : v', v : 1, v' : 2])_{\{x,y\}} \mid v_0, v_1, v, v' \in V_{int}\} \\ &= \llbracket x:=\mathbf{cons}(1); y:=\mathbf{cons}(2) \rrbracket \end{aligned}$$

- (x) $\mathbf{dispose}(x); \mathbf{dispose}(y)$ has trace set

$$\begin{aligned} & \{([x : v, y : v', v : v_0, v' : v'_0], [x : v, y : v'])_{\{x,y\}} \mid v, v', v_0, v'_0 \in V_{int} \ \& \ v \neq v'\} \\ & \cup \{([x : v, y : v, v : v'], abort) \mid v, v' \in V_{int}\} \end{aligned}$$

- (xi) $\mathbf{while} \ \mathbf{true} \ \mathbf{do} \ ([x]:=0; x:=x+1)$ has the trace set

$$\{(\sigma, \perp)_{\{x,y\}} \mid (x, v) \in \sigma \ \& \ \text{dom}(\sigma) = \{x\} \cup \{n \in V_{int} \mid v \leq n\}\}.$$

(xii) Let PUT and GET be the following commands:

$$PUT :: \mathbf{with} \textit{buf} \mathbf{when} \textit{full} = 0 \mathbf{do} (z:=x; \textit{full}:=1)$$

$$GET :: \mathbf{with} \textit{buf} \mathbf{when} \textit{full} = 1 \mathbf{do} (y:=z; \textit{full}:=0)$$

The trace set of PUT is given by:

$$\llbracket PUT \rrbracket = \textit{wait}_0^* \textit{enter}_0 \cup \textit{wait}_0^\omega$$

$$\textit{wait}_0 = \{\textit{try}(\textit{buf})\} \cup \{\textit{acq}(\textit{buf}) \textit{full} = v \textit{rel}(\textit{buf}) \mid v \neq 0\}$$

$$\textit{enter}_0 = \textit{acq}(\textit{buf}) \textit{full} = 0 \llbracket z:=x; \textit{full}:=1 \rrbracket \textit{rel}(\textit{buf})$$

Similarly, the trace set of GET is given by:

$$\llbracket GET \rrbracket = \textit{wait}_1^* \textit{enter}_1 \cup \textit{wait}_1^\omega$$

$$\textit{wait}_1 = \{\textit{try}(\textit{buf})\} \cup \{\textit{acq}(\textit{buf}) \textit{full} = v \textit{rel}(\textit{buf}) \mid v \neq 1\}$$

$$\textit{enter}_1 = \textit{acq}(\textit{buf}) \textit{full} = 1 \llbracket y:=z; \textit{full}:=0 \rrbracket \textit{rel}(\textit{buf})$$

Let $prog$ be the program

$$\textit{full}:=0; \mathbf{resource} \textit{buf} \mathbf{in} (x:=\mathbf{cons}(-); PUT) \parallel (GET; \mathbf{dispose}(y)).$$

The trace set of this program is

$$\begin{aligned} \llbracket prog \rrbracket = \{ & ([\textit{full} : v_0, x : v_1, y : v_2, z : v_3], [\textit{full} : 0, x : v, y : v, z : v]) \}_{\{ \\ & \mid v, v_0, v_1, v_2, v_3 \in V_{int}\}. \end{aligned}$$

Note that this trace set consists of single footsteps, one for each possible initial state and heap cell chosen for x . In fact this is the same as the trace set of the sequential program $\textit{full}:=0; x:=\mathbf{cons}(-); z:=x; y:=z; \mathbf{dispose}(y)$.

9 Semantic properties

As these examples suggest, the trace set of a resource-free command has a degenerate form, containing only single-step traces. (The program of example (xii) has no free resource names.) We can establish a more general property from which this result follows. We already introduced $\mathbf{res}(c)$, the set of resource names occurring free in c . Similarly, for a trace α , let $\mathbf{res}(\alpha)$ be the set of resource names appearing in actions along α . The following result is easy to establish by structural induction.

Lemma 9.1 *For all commands c , and all traces $\alpha \in \llbracket c \rrbracket$, $\mathbf{res}(\alpha) \subseteq \mathbf{res}(c)$.*

Given the structural constraints obeyed by all traces in $\llbracket c \rrbracket$, i.e. that successive footsteps get combined between resource actions, it follows that a command c with no free resource names will have only singleton traces.

Theorem 9.2 *For all commands c such that $\mathbf{res}(c) = \{\}$, every trace of c is a footprint, one of the three forms $(\sigma_0, \sigma'_0)_X$, $(\sigma_0, \textit{abort})$, or $(\sigma_0, \perp)_X$, where σ_0*

and σ'_0 range over states.

The trace set of such a command corresponds to a (non-deterministic) state transformation, represented as a subset of $\mathbf{St} \times (\mathbf{St} \cup \{abort, \perp\})$. We can make this correspondence explicit as follows.

A footprint specifies a “footprint”: just the pieces of state needed to enable the step and describe its result [14,13]. Intuitively, a step of the form $(\sigma_0, \sigma'_0)_X$ is executable (without error) from a “global” state σ if σ_0 is a subset of σ and the part of σ not needed for the step, which is assumed to remain unchanged, can be properly combined with σ'_0 to produce the resulting state. On the other hand, if σ is consistent with σ_0 but missing part of σ_0 the step is still executable but leads to a runtime error because of an uninitialized variable. With this rationale, and similar considerations for the other forms of footsteps, we therefore define the *enabling relations* $\xRightarrow{\lambda} \subseteq \mathbf{St} \times (\mathbf{St} \cup \{abort, \perp\})$ as follows.

Definition 9.3

For a footprint λ the enabling relation $\xRightarrow{\lambda}$ is the least relation satisfying the following conditions:

- For a normal footprint $\lambda = (\sigma_0, \sigma'_0)_X$
 - if $\sigma_0 \subseteq \sigma$ and $(\sigma \setminus \sigma_0) \perp \sigma'_0$ then $\sigma \xRightarrow{(\sigma_0, \sigma'_0)_X} \sigma'_0 \cdot (\sigma \setminus \sigma_0)$
 - if $\sigma \upharpoonright \text{dom}(\sigma_0) \subset \sigma_0$ then $\sigma \xRightarrow{(\sigma_0, \sigma'_0)_X} abort$
- For an error step $\lambda = (\sigma_0, abort)$,
 - if either $\sigma_0 \subseteq \sigma$ or $\sigma \upharpoonright \text{dom}(\sigma_0) \subset \sigma_0$ then $\sigma \xRightarrow{(\sigma_0, abort)} abort$
- For a non-terminating step $\lambda = (\sigma_0, \perp)_X$,
 - if $\sigma_0 \subseteq \sigma$ then $\sigma \xRightarrow{(\sigma_0, \perp)_X} \perp$
 - if $\sigma \upharpoonright \text{dom}(\sigma_0) \subset \sigma_0$ then $\sigma \xRightarrow{(\sigma_0, \perp)_X} abort$.

This definition extends in the obvious way to $\sigma \xRightarrow{\alpha} \sigma'$ when α is a finite sequence of actions, and $\sigma \xRightarrow{\alpha} \cdot$ when α is an infinite sequence of actions.

We then obtain, for a resource-free command c , a non-deterministic state transformation $|c| \subseteq \mathbf{St} \times (\mathbf{St} \cup \{abort, \perp\})$ in the obvious way.

Definition 9.4

For a program c such that $\text{res}(c) = \{\}$, the state transformation denoted by c is given by:

$$|c| = \bigcup \{(\sigma, \sigma') \mid \exists \lambda \in \llbracket c \rrbracket. \sigma \xRightarrow{\lambda} \sigma'\}.$$

Note that for sequential programs we can reason compositionally in the usual way about state transformations; for example, with the obvious notion of composition we have $|c_1; c_2| = |c_2| \circ |c_1|$, when c_1 and c_2 are resource-free.

We can derive a state transformation from a more general set of footprint traces in the obvious way, generalizing the above account to allow for programs with more complex trace structure. It is natural to interpret a resource action

as having a trivial effect on the state, since such actions only affect and depend upon the availability of resources. Thus we will let $\sigma \xrightarrow{\lambda} \sigma$, for all states σ , when λ is a resource action. We can then generalize further to obtain the effect of a trace, by composing the effects of its actions in the appropriate order. Since an error is terminal, we define $abort \xrightarrow{\lambda} abort$ for all actions λ ; similarly we define $\perp \xrightarrow{\lambda} \perp$ for all λ . We can then obtain a state transformation for an arbitrary command by abstracting away from the structure of the command's traces, retaining only information about the effect of its traces.

Definition 9.5

For an arbitrary program c , the state transformation denoted by c is given by:

$$|c| = \bigcup \{(\sigma, \sigma') \mid \exists \alpha \in \llbracket c \rrbracket. \sigma \xrightarrow{\alpha} \sigma'\}.$$

Of course the general formula coincides with the special case formula given above, when applied to a resource-free program.

With these definitions in hand we can now formulate precisely what it means for a command to be *error-free* from a given initial state.

Definition 9.6

The command c is *error-free* from state σ if no trace of c executable from σ leads to *abort*:

$$\forall \alpha \in \llbracket c \rrbracket. \neg(\sigma \xrightarrow{\alpha} abort).$$

Equivalently, c is error-free from σ if and only if $(\sigma, abort) \notin |c|$.

Note that error-freeness implies the absence of race conditions and dangling pointers, as well as the avoidance of uninitialized variables. For instance, example (xii) above shows that the program

$$\begin{aligned} &full:=0; \mathbf{resource} \textit{ buf} \mathbf{ in} \\ &\quad (x:=\mathbf{cons}(-); \mathit{PUT}) \parallel (\mathit{GET}; \mathbf{dispose}(y)) \end{aligned}$$

is error-free from any state σ with $dom(\sigma) \supseteq \{full, x, y, z\}$. Even though the first process allocates a heap cell and the second disposes it, there is no danger here. This is an example of *ownership transfer* [13].

In contrast, the program

$$\begin{aligned} &full:=0; \mathbf{resource} \textit{ buf} \mathbf{ in} \\ &\quad (x:=\mathbf{cons}(-); \mathit{PUT}; \mathbf{dispose}(x)) \parallel (\mathit{GET}; \mathbf{dispose}(y)) \end{aligned}$$

is not error-free from any state, because even if the program is executed from a state in which $full, x, y, z$ all have initial values both processes will attempt to dispose the same cell. If the program is started from a state in which one of the free identifiers is uninitialized there will also be a runtime error because the program will try to read or write this identifier. The footstep trace set for this program is

$$\{([full : v_0, x : v_1, y : v_2, z : v_3], abort) \mid v_0, v_1, v_2, v_3 \in V_{int}\}.$$

Note that for each state σ there is a way to choose v_0, v_1, v_2, v_3 such that

$$\sigma \xrightarrow{([full:v_0, x:v_1, y:v_2, z:v_3], abort)} abort.$$

10 Concurrent separation logic

Concurrent separation logic [7,13] is a Hoare-style formalism for specifying and proving resource-sensitive partial correctness formulas for parallel programs that share mutable state, modifying the inference rules of the Owicki-Gries logic to make judicious use of separating conjunction [16] in a manner proposed by Peter O’Hearn. A resource-sensitive partial correctness formula has the form $\Gamma \vdash \{p\}c\{q\}$, where Γ is a resource context specifying a partition of the shared variables of c among resources, together with a resource invariant for each resource. The pre- and post-conditions p and q , as well as the resource invariants, are separation logic formulas [16] subject to certain restrictions. O’Hearn [13] has shown that this logic allows correctness proofs for a number of interesting examples, involving a novel form of “ownership transfer” in which permission to manipulate a pointer may be deemed to transfer from one process to another. The soundness of this logic is far from obvious: the original Owicki-Gries rules are unsound in the presence of pointers, and John Reynolds has shown that O’Hearn’s rules are not sound unless certain constraints are imposed on the class of formulas allowed as resource invariants.

Our earlier paper [7] introduced a notion of validity for concurrent separation logic formulas, phrased in terms of a resource-sensitive “local enabling” relation $\xRightarrow{\Gamma}$ based on *action trace semantics*, and proved soundness of the logic with respect to this notion of validity, assuming that resource invariants are *precise* separation logic formulas [16]. This semantics involved actions for reading and writing to identifiers, lookup and update of heap locations, and allocation and disposal of heap locations, in addition to resource actions, an idle action, and an error action. The idea is that when $\sigma \xRightarrow{\Gamma}^{\alpha} \sigma'$ the action trace α describes an execution in which a process performs the steps of α in a parallel context which obeys the Dijkstra principle by respecting the resource invariants and only accessing shared state when holding the relevant resources.

Using footstep traces instead of action traces leads to an analogous notion of local enabling, again describing executions in an environment that respects resources. We can then formulate the analogous notion of validity for concurrent separation logic formulas, based on footsteps. We say that $\Gamma \vdash \{p\}c\{q\}$ is *footstep-valid* if for all σ satisfying p , all footstep traces $\alpha \in \llbracket c \rrbracket$, and all $\sigma' \neq \perp$ such that $\sigma \xRightarrow{\Gamma}^{\alpha} \sigma'$, σ' satisfies q . (As in [7] but using footstep traces rather than action traces to underpin the definition.) Although the difference in underlying model suggests that the notion of footstep-validity might differ from the notion of validity used in [7], in fact the two notions coincide. We will discuss some of the ramifications of this result after a brief summary of the main ideas behind the proof.

Theorem 10.1

$\Gamma \vdash \{p\}c\{q\}$ is action-trace valid iff it is footstep-trace valid.

Proof (Sketch)

This can be proven by establishing a correspondence between action traces and footstep traces, and between the corresponding notions of effect and interference of the two models. For each action μ (in the alphabet of the action trace semantic model) we define a set $\mathcal{F}(\mu)$ of footsteps, such that for all states σ and σ' ,

$$\sigma \xrightarrow{\mu} \sigma' \text{ iff } \exists \lambda \in \mathcal{F}(\mu). \sigma \xrightarrow{\lambda} \sigma'.$$

One can also show that for all actions μ , and all states σ, σ'' :

$$(\exists \sigma'. \sigma \xrightarrow{\mu} \sigma' \xrightarrow{\mu'} \sigma'') \text{ iff } \exists \lambda \in \mathcal{F}(\mu), \lambda' \in \mathcal{F}(\mu'). \lambda \succ \lambda' \ \& \ \sigma \xrightarrow{\lambda; \lambda'} \sigma''.$$

This property connects the notion of *sequential* trace from [7] with catenability.

The relationship between the notions of interference used in the two models is expressed as follows. For all actions μ, μ' :

$$\mu \bowtie \mu' \Leftrightarrow \exists \lambda \in \mathcal{F}(\mu), \lambda' \in \mathcal{F}(\mu'). \lambda \bowtie \lambda'.$$

The resource context Γ specifies, for each resource occurring in c , a set of identifiers protected by that resource, and a resource invariant. By analogy with [7], we define a “local enabling” relation $\xrightarrow[\Gamma]{\lambda}$ for footstep actions λ , modifying the definition of $\xrightarrow{\lambda}$ to take account of the resource invariants and protection lists, and treating any violation of the resource discipline as an error. We show that for all μ and μ' :

$$\sigma \xrightarrow[\Gamma]{\mu} \sigma' \text{ iff } \exists \lambda \in \mathcal{F}(\mu). \sigma \xrightarrow[\Gamma]{\lambda} \sigma'.$$

In addition, for all actions μ and μ' , and all states σ, σ'' ,

$$(\exists \sigma'. \sigma \xrightarrow[\Gamma]{\mu} \sigma' \xrightarrow[\Gamma]{\mu'} \sigma'') \text{ iff } \exists \lambda \in \mathcal{F}(\mu), \lambda' \in \mathcal{F}(\mu'). \lambda \succ \lambda' \ \& \ \sigma \xrightarrow[\Gamma]{\lambda; \lambda'} \sigma''.$$

With the obvious extension of \mathcal{F} to sequences of actions, it then follows that for all action traces α , and all σ and σ'' , $\sigma \xrightarrow[\Gamma]{\alpha} \sigma''$ if and only if there is a footstep trace $\beta \in \mathcal{F}(\alpha)$ such that $\sigma \xrightarrow[\Gamma]{\beta} \sigma''$. Hence the two notions of validity coincide. \square

The following theorem can then be proven directly, using footstep trace semantics instead of action traces and following the proof strategy of [7]. It can also be deduced from the soundness result of [7] and the previous theorem.

Theorem 10.2 (Soundness of concurrent separation logic)

Concurrent separation logic is sound with respect to footstep-trace validity: every resource-sensitive partial correctness formula provable from the inference rules is footstep-trace valid.

As a consequence one can use footstep traces in preference to action traces to formalize a correctness proof of a parallel program, an advantage because of

the larger step size, which may reduce the number of interleavings. Indeed, we have already shown earlier that footstep trace semantics supports *sequential reasoning* about synchronization-free code fragments: one can replace any synchronization-free code fragment by another which denotes the same state transformation, without affecting the footstep trace semantics of the overall program. One can use more elementary reasoning techniques to deal with synchronization-free code, such as the familiar Hoare-style inference rules for sequential programs.

11 Levels of abstraction

Although we do not provide here a complete description of the action trace semantic model of [7], we do want to emphasize that the footstep trace model represents an evolutionary improvement over the earlier model, streamlined specifically for reasoning about partial correctness and race-freedom.

The key technical result relating the two models is simple to state:

Theorem 11.1

- (i) *If two expressions denote the same action trace set, they also denote the same footstep trace set.*
- (ii) *If two commands denote the same action trace set, they also denote the same footstep trace set.*

Proof (Sketch)

The details for expressions are simpler, so we focus on commands. We will use μ to range over the action alphabet used in action trace semantics. Let us say that an action trace is *loosely coupled* if it can be executed from some initial state, allowing interference only at synchronizations. We can formalize this property as follows. Define a *loosely coupled enabling relation* $\xrightarrow{\mu}_{\bullet}$ for actions μ , so that:

- $\sigma \xrightarrow{\mu}_{\bullet} \sigma'$ iff $\sigma \xrightarrow{\mu} \sigma'$, when μ is not an acquire action
- $\sigma \xrightarrow{\mu}_{\bullet} \sigma'$ for all σ and σ' , if μ is *acq*(r) for some resource name r .

As before, this generalizes to action traces in the obvious way by composition. We say that the action trace α is loosely coupled if there is a state σ such that $\sigma \xrightarrow{\alpha}_{\bullet} \cdot$.

For a loosely coupled action trace α the set $\mathcal{F}(\alpha)$ represents the set of catenable footstep traces generated by α . Let $\llbracket c \rrbracket_{fs}$ be the footstep trace set of c , and $\llbracket c \rrbracket_{lc}$ be the set of loosely coupled action traces of c . The following property can be proven by structural induction on c , based on the semantic clauses:

$$\llbracket c \rrbracket_{fs} = \{ cat(\beta) \mid \beta \in \mathcal{F}(\alpha) \ \& \ \alpha \in \llbracket c \rrbracket_{lc} \}.$$

Hence action trace equivalence implies footstep trace equivalence. \square

Thus footstep trace semantics is more abstract than action trace semantics. The converse does not always hold.

The example programs discussed earlier can also be used to emphasize the improvements achieved by the footstep trace model, since these same examples appear in the action traces paper. Their action trace sets contain traces in which the ordering of individual read and write actions is recorded; these intermediate details are ignored in the construction of the new model, because of the “mumbling” effect built into our definition of concatenation. For example, the commands $x:=1; x:=x+1$ and $x:=2$ have the same footstep traces, but different action traces, since $x:=1; x:=x+1$ has action traces of the form $x:=1 \ x=v \ x:=v+1$, for $v \in V_{int}$. Similarly the commands $x:=1; y:=2$ and $y:=2; x:=1$ have the same footstep traces but not the same action traces. In addition, the action trace model distinguishes between the commands **while true do** $x:=x+1$ and **while true do** $x:=x+2$, whereas they have the same footstep traces, namely $\{([x : v], \perp)_{\perp} \mid v \in V_{int}\}$.

A more compelling illustration of the advantages of footstep traces over action traces is provided by some programs involving the one-place buffer resource and the code fragments *PUT* and *GET*. The footstep trace semantics of the programs

$$\begin{aligned} & full:=0; \text{resource } buf \text{ in} \\ & (x:=\mathbf{cons}(-); PUT) \parallel (GET; \mathbf{dispose}(y)) \end{aligned}$$

and

$$\begin{aligned} & full:=0; \text{resource } buf \text{ in} \\ & (x:=\mathbf{cons}(-); PUT; \mathbf{dispose}(x)) \parallel GET \end{aligned}$$

are identical, both equal to the set

$$\begin{aligned} & \{([full : v_0, x : v_1, y : v_2, z : v_3], [full : 0, x : v, y : v, z : v])_{\perp} \} \\ & \mid v, v_0, v_1, v_2, v_3 \in V_{int} \}. \end{aligned}$$

Their action trace sets differ, because the action traces retain information about the relative ordering of steps: in the second program the disposal action may occur concurrently with *GET* but in the first program it follows the *GET*.

Furthermore, consider what happens when we iterate the buffer processes, running a sequence of puts concurrently with a sequence of gets. For all non-negative integers N , the program

$$\begin{aligned} & full:=0; \text{resource } buf \text{ in} \\ & (x:=\mathbf{cons}(-); PUT)^N \parallel (GET; \mathbf{dispose}(y))^N \end{aligned}$$

has *exactly the same* footstep traces as the above programs: again the trace set degenerates to a set of single footsteps, and the effect is still to (re)set *full*

to 0 and to make x, y , and z all equal; in contrast, the action trace set of this program (using the semantics of [7]) suffers from a combinatorial explosion, and contains many traces that yield the same effect.

Similarly all programs of the form

$$\mathit{full}:=0; \mathbf{resource\ buf\ in} \\ (x:=\mathbf{cons}(-); \mathit{PUT}; \mathbf{dispose}(x))^N \parallel \mathit{GET}^N$$

have the same trace set again.

Although these two families of programs (parameterized by N) are rather contrived, they demonstrate the potential of our new semantics for greatly reducing the need to consider interleavings. It is also worth remarking that the fact that each of these programs has the same set of footprint traces implies that they all satisfy exactly the same formulas of concurrent separation logic. This is by no means obvious; each family of programs requires a different proof strategy in the logic [13] because one needs to choose an appropriate resource invariant that expresses the ownership transfer policy implicit in the program structure. For programs in the first family ownership of the heap cell is deemed to transfer after a put operation, whereas this is not the case for programs in the second family.

12 Granularity

We have worked so far with states whose structure assumes that an integer-valued variable is stored as a single integer: we abstracted away from machine-level details such as word size. If instead we assume that programs are executed on a machine with a fixed, finite word size, it would be natural to introduce a more concrete model of state in which we represent an integer-valued variable as a list of word-valued component variables. For purely sequential programs it is well known that the more concrete model is unnecessary if we care only about high-level properties: a state transformation semantics based on integers is consistent with the more concrete semantics based on words. The same is true for the footprint trace semantics. It is straightforward to define a low-level word-based version of footprint traces, with variables represented as lists of components, each holding an integer small enough to fit inside a single word. Since the semantic definitions ensure that consecutive (and catenable) footsteps get collapsed into a single step, a sequence of word-level operations that implements an integer-level operation (like writing an integer to a variable) will give rise to a single word-level footprint. It can then be shown that the high-level footprint semantics is equipollent to the low-level, in that two programs have the same high-level semantics if and only if they have the same low-level semantics. This result is analogous to a similar result presented by the author at MFPS 2004 [8], which discussed high- and low-level versions of action trace semantics [7] and the relationship between the two versions. The result can be seen as a justification for adopting a more abstract model with

no loss of generality.

13 Future work

If c_1 and c_2 have the same footstep traces, then for all program contexts $P[-]$, $P[c_1]$ and $P[c_2]$ satisfy the same concurrent separation logic formulas.⁵ The converse fails: $(x:=1; y:=2) \parallel x:=2$ and $x:=1 \parallel x:=2$ have different footstep traces, since the first command has traces of the form $(\sigma, abort)$ where $x, y \in \text{dom}(\sigma)$ whereas the second command does not need a value for y ; nevertheless these commands satisfy the same concurrent separation logic formulas (none!), because they both cause a race. It would be interesting to see if we can further fine-tune the structure of footstep traces to avoid making distinctions like this.

We have indicated the greater succinctness of the new semantic model by applying it to some examples, and by proving a theorem to the effect that for synchronization-free programs the semantics takes a simple degenerate form equivalent to a (non-deterministic) state transformation. The semantic definitions may be used directly to calculate footstep trace sets and analyze program behavior, and this kind of manual analysis is made less tedious by the streamlined structure of the semantic model. The semantics also validates a number of laws of program equivalence which can be used to simplify program analysis: in particular, all laws of equivalence valid for action trace semantics are also valid for footstep traces, as was shown above. Nevertheless, the manual approach is still error-prone and for large-scale programs with many synchronization points we still face severe combinatorial problems, so that mechanical assistance would be an advantage. We believe that our approach may suggest a way to improve the efficiency of *model checking* [11], since our semantics is designed to avoid unnecessary interleaving and hence may help to tame the combinatorial explosion.

Recent papers [1,3] have proposed methods for reasoning about parallel programs based on *permissions*, building on earlier work of Boyland [4] and extending concurrent separation logic to deal with classic but logically tricky examples such as readers/writers [9]. Initial investigations suggest that the semantics presented in this paper offers a highly promising way to analyze such programs: the model's succinctness and its careful design to embody Dijkstra's principle facilitate compositional reasoning. We plan to use footstep trace semantics to prove the soundness of Bornat's permission-based logic [1], as it seems likely that the proof method used in [7] may be adjustable in a natural manner to incorporate permission accounting, by analogy with the way the original soundness proof formalizes and justifies reasoning about ownership transfer.

⁵ A similar implication holds for action trace semantics, by the previous remarks.

References

- [1] R. Bornat. *Variables as Resource in Separation Logic*. MFPS 2005, Proceedings. Martin Escardó (Ed.), Birmingham, England, May 18-21, 2005. To appear, Elsevier ENTCS (2005).
- [2] P. Brinch Hansen. *Structured multiprogramming*. Comm. ACM, 15(7):574-578, July 1972.
- [3] R. Bornat, C. Calcagno, P. O’Hearn, and M. Parkinson. *Permission accounting in separation logic*. Proc. POPL 2005. Pages 259–270. SIGPLAN-SIGACT. January 2005.
- [4] J. Boyland. *Checking interference with fractional permissions*. In: Static Analysis, 10th International Symposium, R. Cousot, ed. Springer LNCS vol. 2694 (2003), pp. 55–72.
- [5] P. Brinch Hansen. *Concurrent programming concepts*. ACM Computing Surveys 5(4):223-245, December 1973.
- [6] S. D. Brookes. *Full abstraction for a shared-variable parallel language*. Proc. 8th IEEE Symposium on Logic in Computer Science, IEEE Computer Society Press (1993), 98–109. Journal version in: *Inf. Comp.*, vol 127(2):145-163, Academic Press, June 1996.
- [7] S. D. Brookes. *A Semantics for Concurrent Separation Logic*. Invited paper, CONCUR 2004, Philippa Gardner and Nobuko Yoshida (Eds.), Springer LNCS 3170, London, August/September 2004, pp. 16-34. Full version to appear in *Theoretical Computer Science* (2005).
- [8] S. D. Brookes. *A Race-detecting Semantics for Concurrent Programs*. Slides from talk presented at MFPS 2004. Pittsburgh, May 2004.
- [9] P. J. Courtois, F. Heymans, and D. Parnas. *Concurrent control with “readers” and “writers”*. Comm. ACM. 14 (1971), pp. 667–668.
- [10] E. W. Dijkstra. *Cooperating sequential processes*. In: **Programming Languages**, F. Genuys (editor), pp. 43-112. Academic Press, 1968.
- [11] E. A. Emerson and E. M. Clarke. *Characterizing properties of parallel programs as fixpoints*. Proc. 7th ICALP. Springer LNCS vol. 85 (1981).
- [12] C. A. R. Hoare, *Towards a Theory of Parallel Programming*. In **Operating Systems Techniques**, C. A. R. Hoare and R. H. Perrott, editors, pp. 61-71, Academic Press, 1972.
- [13] P. W. O’Hearn. *Resources, Concurrency, and Local Reasoning*. Invited paper, CONCUR 2004, P. Gardner and N. Yoshida (Eds.), Springer LNCS 3170, London, August/September 2004, pp. 49-67. Full version to appear in *Theoretical Computer Science* (2005).

- [14] P. W. O’Hearn, J. C. Reynolds, and H. Yang. *Local reasoning about programs that alter data structures*. Proc. CSL 2001, L. Fribourg, ed. Springer LNCS vol. 2142 (2001), pp. 1–19.
- [15] S. Owicki and D. Gries, *Verifying properties of parallel programs: An axiomatic approach*, Comm. ACM. 19(5):279-285, May 1976.
- [16] J. C. Reynolds, *Separation logic: a logic for shared mutable data structures*, Invited paper. Proc. 17th IEEE Conference on Logic in Computer Science, LICS 2002, pp. 55-74. IEEE Computer Society, 2002.
- [17] J. C. Reynolds, *Towards a Grainless Semantics for Shared-Variable Concurrency*. Proc. 24th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2004), Chennai, India. Springer-Verlag, December 2004.

Appendix

We include a brief summary of the key definitions and results used to establish the relationship between the action trace semantic model and the footprint traces model. We refer to [7] for the action trace semantic clauses, but we include here a brief summary of the kinds of action appearing in [7] and the corresponding notions of effect and interference.

Let μ range over the actions appearing in the action trace semantics, characterized by the following abstract grammar:

$$\begin{aligned} \mu ::= & \delta \mid i=v \mid i:=v \mid \\ & [v]=v' \mid [v]:=v' \mid \\ & alloc(v, [v_0, \dots, v_n]) \mid disp(v) \mid \\ & try(r) \mid acq(r) \mid rel(r) \mid \\ & abort \end{aligned}$$

Here i ranges over identifiers, v, v', v_0, \dots, v_n over integers, and r over resource names.

We define, for each action μ a set $\mathcal{F}(\mu)$ of footprint actions:

$$\begin{aligned} \mathcal{F}(\delta) &= \{([\], [\])\}_{\{\}} \\ \mathcal{F}(i=v) &= \{([i : v], [i : v])\}_{\{i\}} \\ \mathcal{F}(i:=v) &= \{([i : v_0], [i : v])\}_{\{i\}} \mid v_0 \in V_{int} \\ \mathcal{F}([v]=v') &= \{([v : v'], [v : v'])\}_{\{v\}} \\ \mathcal{F}([v]:=v') &= \{([v : v_0], [v : v'])\}_{\{v\}} \mid v_0 \in V_{int} \\ \mathcal{F}(alloc(v, [v_0, \dots, v_n])) &= \{([\], [v : v_0, v + 1 : v_1, \dots, v + n : v_n])\}_{\{\}} \\ \mathcal{F}(disp(v)) &= \{([v : v_0], [\])\}_{\{v\}} \mid v_0 \in V_{int} \\ \mathcal{F}(acq(r)) &= \{acq(r)\} \\ \mathcal{F}(rel(r)) &= \{rel(r)\} \\ \mathcal{F}(try(r)) &= \{try(r)\} \\ \mathcal{F}(abort) &= \{abort\} \end{aligned}$$

Note that a write action $i:=v$ determines a set of footprints, each of which writes the value v to i ; the set $\mathcal{F}(i:=v)$ contains a footprint for each possible initial value of i . For similar reasons update actions $[v]:=v'$, allocations, and disposals all generate non-trivial sets of footprints. In all other cases the action μ corresponds to a single footprint action.

Recall the notion of *effect* given for action traces [7]; just as with footprint traces we define for each action μ a partial function $\xRightarrow{\mu}$ from \mathbf{St} to $\mathbf{St} \cup \{abort\}$.

The definition is by cases on the structure of μ :

- $\sigma \xrightarrow{\delta} \sigma'$ if and only if $\sigma = \sigma'$
- $\sigma \xrightarrow{i=v} \sigma'$ if and only if $(i, v) \in \sigma$ and $\sigma = \sigma'$
- $\sigma \xrightarrow{i:=v} \sigma'$ if and only if
 - either $i \in \text{dom}(\sigma)$ and $\sigma' = [\sigma \mid i : v]$
 - or $i \notin \text{dom}(\sigma)$ and $\sigma' = \text{abort}$
- $\sigma \xrightarrow{[v]=v'} \sigma'$ if and only if $(v, v') \in \sigma$ and $\sigma = \sigma'$
- $\sigma \xrightarrow{[v]:=v'} \sigma'$ if and only if
 - either $v \in \text{dom}(\sigma)$ and $\sigma' = [\sigma \mid v : v']$
 - or $v \notin \text{dom}(\sigma)$ and $\sigma' = \text{abort}$
- $\sigma \xrightarrow{\text{alloc}(v, [v_0, \dots, v_n])} \sigma'$ if and only if $\{v, v+1, \dots, v+n\} \cap \text{dom}(\sigma) = \{\}$ and $\sigma' = [\sigma \mid v : v_0, v+1 : v_1, \dots, v+n : v_n]$
- $\sigma \xrightarrow{\text{disp}(v)} \sigma'$ if and only if
 - either $v \in \text{dom}(\sigma)$ and $\sigma' = \sigma \setminus v$
 - or $v \notin \text{dom}(\sigma)$ and $\sigma' = \text{abort}$
- $\sigma \xrightarrow{\mu} \sigma'$ if and only if $\sigma' = \sigma$, when μ is a resource action
- $\sigma \xrightarrow{\text{abort}} \sigma'$ if and only if $\sigma' = \text{abort}$

As with footstep traces, we obtain the effect of a trace α by composition. For a finite action trace α the effect will be a partial function $\xrightarrow{\alpha}$ from \mathbf{St} to $\mathbf{St} \cup \{\text{abort}\}$. For an infinite trace α we will obtain a partial function $\xrightarrow{\alpha}$ from \mathbf{St} to $\{\perp\}$.

The notion of interfering actions (written $\mu \bowtie \mu'$), as defined in [7], is characterized by:

$$\begin{aligned} \mu \bowtie \mu' &\Leftrightarrow \mu = \text{abort} \vee \mu' = \text{abort} \\ &\vee \text{writes}(\mu) \cap \text{free}(\mu') \neq \{\} \\ &\vee \text{free}(\mu) \cap \text{writes}(\mu') \neq \{\}. \end{aligned}$$

We define the sets $\text{writes}(\mu)$ to be the set of identifiers or heap cells *written* by μ , and $\text{free}(\mu)$ to be the set of identifiers or heap cells either *read* or *written* by μ , so that interference corresponds to a concurrent write by one process to data being used by another [7].