

Full Abstraction for a Shared-Variable Parallel Language

STEPHEN BROOKES

School of Computer Science, Carnegie-Mellon University, Schenley Park, Pittsburgh, Pennsylvania 15213

We give a new denotational semantics for a shared-variable parallel programming language and prove full abstraction: the semantics gives identical meanings to commands if and only if they induce the same behavior in all program contexts. The meaning of a command is a set of "transition traces," which record the ways in which a command may interact with and be affected by its environment. We show how to modify the semantics to incorporate new program constructs, to allow for different levels of granularity or atomicity, and to model fair infinite computation, in each case achieving full abstraction with respect to an appropriate notion of program behavior. © 1996 Academic Press, Inc.

1. INTRODUCTION

One of the fundamental purposes of semantics is to provide rigorous means of proving the correctness of programs with respect to behavioral specifications. For any particular language different semantic models may be suitable for reasoning about different behavioral notions, such as *partial correctness*, *total correctness*, *deadlock-freedom*, and more general *liveness* and *safety* properties [14]. Ideally one would like a semantics in which the meaning of one term coincides with the meaning of another term if and only if the terms induce the same behavior in each program context; this guarantees that one term may be replaced by the other in any context without affecting the behavior of the overall program, thus supporting compositional or modular reasoning about program behavior. Such a semantics is said to be *equationally fully abstract* with respect to the given notion of behavior [17, 20, 22]. When the set of program behaviors is equipped with an approximation ordering and the semantic model has a partial order such that the meaning of one term is less than the meaning of another if and only if the behavior of the first term in each program context approximates the behavior of the second term in the same context, the semantics is said to be *inequationally fully abstract* with respect to the given notion of program behavior and approximation. Clearly an inequationally fully abstract semantics is also equationally fully abstract. Intuitively, a fully abstract semantics is at exactly the right level of abstraction to support compositional reasoning.

The difficulty of finding fully abstract semantics is well known [4, 17, 20, 22]. Many standard semantic models are

correct, in that whenever two terms induce different behavior in some context they denote different meanings, but *too concrete* since the converse may fail. Sometimes one can show that by adding extra syntactic constructs to the programming language the model becomes fully abstract. However, unless the extra constructs are computationally natural and the original language was clearly deficient because of their omission, the full abstraction problem for the original language is still important.

The standard state-transformation semantics for sequential while-programs is fully abstract with respect to partial correctness behavior. However, for a parallel version of this language [9, 18], in which parallel commands can interact by updating and reading shared variables, the full abstraction problem is more difficult. Parallel programs may exhibit non-deterministic behavior, depending on the scheduling of atomic actions, so the partial correctness behavior of a parallel command is naturally modelled as a non-deterministic state transformation, usually represented as a function from states to *sets* of states. However, the state transformation denoted by a parallel combination of commands cannot be determined solely from the state transformations denoted by the component commands; thus the state-transformation semantics for a parallel language is not even compositional, and is certainly not fully abstract. It is not even sufficient to model a program as a set of *sequences* of states, each sequence recording a possible execution history of the program, since this semantics still fails to be compositional. One needs a semantic model with more detailed structure, so that the possible interactions between commands executing in parallel may be modelled appropriately.

Hennessy and Plotkin [9] described a denotational semantics for this language, based on a recursively defined domain of *resumptions*, built with a powerdomain operator. However, the resumptions semantics is too concrete: **skip** and **skip** denote different resumptions even though they induce the same partial correctness behavior in all contexts. They showed that with the addition of extra features to the programming language, the resumptions model becomes fully abstract. However, one of the extra constructs is a rather peculiar form of coroutine execution which allows counting of the number of atomic steps taken by a

command executing in parallel. The problem remained of finding a fully abstract model for the original parallel language.

It is often reasonable to assume that program execution is *fair*, in that commands executing in parallel are not artificially prevented from making progress. In particular, execution is *weakly fair* if whenever a parallel command is continuously enabled it eventually gets to move [16, 8]. Even when parallel execution is being simulated on a shared uni-processor, using a scheduler that interleaves the activities of parallel commands, it is easy to implement weak fairness using a round-robin scheduling strategy. The assumption of (weak) fairness allows us to abstract away from the details of any particular scheduling strategy and from the relative speeds of parallel processes. It is well known that fair execution gives rise to the phenomenon of unbounded non-determinism—there are parallel programs guaranteed to terminate under a fair execution strategy but which may terminate in any of a (countably) infinite set of final states. The interaction of unbounded non-determinism with power-domain constructions is problematic. For instance it is difficult to see how to adapt the Hennessy–Plotkin resumptions model to incorporate fairness; moreover in Apt and Plotkin’s semantics [4], unbounded non-determinism causes lack of continuity of certain key functions. Park [19] provided a denotational semantics for shared-variable programs, including a formal treatment of fairness, but failed to achieve full abstraction for a familiar reason: Park’s semantics distinguished between `skip` and `skip; skip`. This left open the problem of finding a fully abstract model for shared-variable programs with respect to a suitable notion of fair behavior.

In this paper we solve these problems. We first describe a new denotational semantics for shared-variable programs, and we show that it is fully abstract with respect to partial correctness behavior. We do not need to add a coroutine construct to the language in order to achieve full abstraction; instead, our proof makes essential use of a synchronization primitive (the “conditional critical region” construct) already present in the language. We model the meaning of a command as a set of *transition traces*. A transition trace is a finite sequence of pairs of states recording a possible interaction sequence of the command with its environment; each pair of states represents the effect of a finite, possibly empty, sequence of atomic actions. The set of traces of a command is closed under two natural operations: “stuttering” (cf. Lamport [15]) and “mumbling.” This model is conceptually simpler than the resumptions model, since it does not require the use of powerdomains or recursively defined domains. The model also validates a number of intuitively natural equations and inequations between programs which fail in the resumptions model.

We also show how to incorporate fairness into our semantic framework in a natural way, by including infinite

transition traces representing fair computations. In doing this we build on the foundations laid by Park [19], with the key extra ingredient supplied by our use of *closed* trace sets. All operations on trace sets used in our semantic definitions are continuous, with respect to set inclusion. We prove full abstraction with respect to the appropriate notion of behavior, termed “strong correctness,” in which both termination and divergence (non-termination) are regarded as observable. Our semantics may therefore be used to reason about total correctness, partial correctness, safety and liveness properties of parallel programs executing fairly.

We also show that our semantic models are adaptable to a variety of settings: one may easily accommodate the addition of certain extra features to the programming language, and the results do not depend crucially on assumptions about the level of atomicity or granularity of execution.

2. SYNTAX

We discuss a standard shared-variable parallel language, as in [9, 18]. There are four syntactic sets: **Id**, the set of identifiers, ranged over by I ; **Exp**, the set of expressions, ranged over by E ; **BExp**, the set of boolean expressions, ranged over by B ; and **Com**, the set of commands, ranged over by C . Identifiers and expressions denote integer values, boolean expressions denote truth values, and the language contains the usual arithmetic and boolean operators and constants. For commands we specify the following grammar:

$$C ::= \text{skip} \mid I := E \mid C_1; C_2 \mid \text{if } B \text{ then } C_1 \text{ else } C_2 \mid \\ \text{while } B \text{ do } C \mid C_1 \parallel C_2 \mid \text{await } B \text{ then } C.$$

A command of the form `await B then C` is a *conditional critical region*,¹ converting C into an atomic action that is enabled only in states satisfying B ; we impose the (reasonable) syntactic restriction that C must be a finite sequence of assignments (or `skip`).²

We write $\text{free}[E]$ to denote the set of identifiers occurring free in E , and similarly for boolean expressions. For commands we define

$$\begin{aligned} \text{free}[\text{skip}] &= \{ \} \\ \text{free}[I := E] &= \{I\} \cup \text{free}[E] \\ \text{free}[C_1; C_2] &= \text{free}[C_1] \cup \text{free}[C_2] \end{aligned}$$

¹ Also known as a “conditional atomic action” [3].

² This restriction is motivated by pragmatic concerns: it is easy and inexpensive for a scheduler to suspend execution of other parallel commands while a finite sequence of assignments is performed, but not practical to suspend while an arbitrary command executes, since a loop may take an unbounded amount of time. Moreover, the restriction does not hamper the utility of the programming language and suffices for establishing full abstraction.

$$\begin{aligned} \text{free}[\text{if } B \text{ then } C_1 \text{ else } C_2] &= \text{free}[B] \cup \text{free}[C_1] \cup \text{free}[C_2] \\ \text{free}[\text{while } B \text{ do } C] &= \text{free}[B] \cup \text{free}[C] \\ \text{free}[\text{await } B \text{ then } C] &= \text{free}[B] \cup \text{free}[C] \\ \text{free}[C_1 \parallel C_2] &= \text{free}[C_1] \cup \text{free}[C_2]. \end{aligned}$$

3. OPERATIONAL SEMANTICS

We present a structural operational semantics similar to the semantics given in [9].

We use N for the set of (non-negative) integers, ranged over by n ; and $V = \{\text{tt}, \text{ff}\}$ for the set of truth values, ranged over by v . A state is a finite partial function from identifiers to integer values. Let $S = \text{Ide} \rightarrow_p N$ denote the set of states, ranged over by s . We write $\text{dom}(s)$ for the domain of s , and $[s|I=n]$ for the state which agrees with s except that it gives identifier I the value n . We use notation like $[I_1=n_1, \dots, I_k=n_k]$ for states.

When s is a state defined on (at least) the free identifiers of E , we write $\langle E, s \rangle \rightarrow^* n$ to indicate that E evaluates to n in state s . Similarly for boolean expressions. We assume that the semantics of expressions and boolean expressions are given by semantic functions \mathcal{E} and \mathcal{B} , characterized operationally by

$$\begin{aligned} \mathcal{E}[E] &= \{(s, n) \mid \langle E, s \rangle \rightarrow^* n\} \\ \mathcal{B}[B] &= \{(s, v) \mid \langle B, s \rangle \rightarrow^* v\}. \end{aligned}$$

For command execution we specify a set of configurations

$$\text{Conf} = \{\langle C, s \rangle \in \text{Com} \times S \mid \text{free}[C] \subseteq \text{dom}(s)\},$$

a transition relation $\rightarrow \subseteq \text{Conf} \times \text{Conf}$, and a subset of *successfully terminated* configurations. A configuration of the form $\langle C, s \rangle$ will represent a stage in a computation at which the remaining command to be executed is C , and the current state is s . A transition of the form $\langle C, s \rangle \rightarrow \langle C', s' \rangle$ represents a computation step in which an atomic step of C enabled in state s is performed, causing the state to change to s' and after which the remaining command to be executed is C' . The transition relation for commands is defined to be the smallest relation satisfying the syntax-directed transition rules given in Fig. 1. This means that a transition is possible if and only if it can be deduced from these rules. The successfully terminated configurations are those for which $\langle C, s \rangle \text{ term}$ is provable from the rules in Fig. 1.

Note how the rules specify the atomicity of boolean expression evaluation, assignment, and conditional critical regions, by means of \rightarrow^* , the reflexive transitive closure of the transition relation. The transition rules for $C_1 \parallel C_2$ clearly specify the interleaving of atomic actions performed by C_1 with those performed by C_2 , and we have specified that a parallel composition terminates only when both components have terminated.

$$\begin{array}{c} \langle \text{skip}, s \rangle \text{ term} \\ \langle E, s \rangle \rightarrow^* n \\ \hline \langle I := E, s \rangle \rightarrow \langle \text{skip}, [s|I=n] \rangle \\ \hline \frac{\langle C_1, s \rangle \rightarrow \langle C'_1, s' \rangle \quad \langle C_1, s \rangle \text{ term}}{\langle C_1; C_2, s \rangle \rightarrow \langle C'_1; C_2, s' \rangle} \quad \frac{\langle C_1, s \rangle \text{ term}}{\langle C_1; C_2, s \rangle \rightarrow \langle C_2, s \rangle} \\ \hline \frac{\langle B, s \rangle \rightarrow^* \text{tt}}{\langle \text{if } B \text{ then } C_1 \text{ else } C_2, s \rangle \rightarrow \langle C_1, s \rangle} \\ \hline \frac{\langle B, s \rangle \rightarrow^* \text{ff}}{\langle \text{if } B \text{ then } C_1 \text{ else } C_2, s \rangle \rightarrow \langle C_2, s \rangle} \\ \langle \text{while } B \text{ do } C, s \rangle \rightarrow \langle \text{if } B \text{ then } C; \text{while } B \text{ do } C \text{ else skip}, s \rangle \\ \hline \frac{\langle C_1, s \rangle \rightarrow \langle C'_1, s' \rangle \quad \langle C_2, s \rangle \rightarrow \langle C'_2, s' \rangle}{\langle C_1 \parallel C_2, s \rangle \rightarrow \langle C'_1 \parallel C_2, s' \rangle} \quad \frac{\langle C_2, s \rangle \rightarrow \langle C'_2, s' \rangle}{\langle C_1 \parallel C_2, s \rangle \rightarrow \langle C_1 \parallel C'_2, s' \rangle} \\ \hline \frac{\langle C_1, s \rangle \text{ term} \quad \langle C_2, s \rangle \text{ term}}{\langle C_1 \parallel C_2, s \rangle \text{ term}} \\ \hline \frac{\langle B, s \rangle \rightarrow^* \text{tt} \quad \langle C, s \rangle \rightarrow^* \langle C', s' \rangle \text{ term}}{\langle \text{await } B \text{ then } C, s \rangle \rightarrow \langle C', s' \rangle} \\ \hline \frac{\langle B, s \rangle \rightarrow^* \text{ff}}{\langle \text{await } B \text{ then } C, s \rangle \rightarrow \langle \text{await } B \text{ then } C, s \rangle} \end{array}$$

FIG. 1. Transition rules for commands.

A configuration like $\langle \text{await } B \text{ then } C, s \rangle$ is said to be “blocked” if B is false in s ; we model blocking operationally by means of a “busy-wait” or “idle” transition, as expressed in the final rule of Fig. 1.³ More generally a parallel command is said to be blocked in a given state if each of its component commands is blocked in this manner; thus C is blocked in state s iff the only possible transition of $\langle C, s \rangle$ is an “idle” transition back to itself.

A *computation* of a command C from state s is a maximal sequence of consecutive transitions starting from $\langle C, s \rangle$. A computation is either finite, ending in a successfully terminated configuration, or infinite.

4. PROGRAM BEHAVIOR

We will focus initially on two closely related notions of behavior, each determined by the finite computations of a program.

³ An alternative interpretation of blocking is obtained if we omit this transition rule: a blocked configuration would then have no enabled transition. We would then need to distinguish between two kinds of finite computation (successful and blocking). This would greatly complicate the presentation of later definitions and results, without gaining any generality.

DEFINITION 4.1. The *partial correctness* behavior function $\mathcal{M}: \mathbf{Com} \rightarrow \mathcal{P}(S \times S)$ is defined by

$$\mathcal{M}[C] = \{(s, s') \mid \langle C, s \rangle \rightarrow^* \langle C', s' \rangle \text{ term}\}.$$

A Hoare-style partial correctness assertion $\{P\} C \{Q\}$, where P and Q are boolean-valued expressions involving identifiers, is valid if every terminating computation of C from a state satisfying P ends in a state satisfying Q . As usual, P and Q are “conditions” drawn from some specification language. Assuming that the condition language contains at least expressions of the form

$$x_1 = n_1 \ \& \ \dots \ \& \ x_k = n_k,$$

where the x_i are identifiers and the n_i are numerals, it is easy to see that two programs have the same partial correctness behavior if and only if they satisfy the same Hoare-style assertions.

The partial correctness behavior \mathcal{M} ignores all intermediate states that arise during a computation, focussing only on the first and last states of successful computations. We now present a simple generalization. A (finite) *state trace* of C is a sequence of states occurring during a successful computation of C . Let S^+ denote the set of non-empty finite sequences of states.

DEFINITION 4.2. The (finite) *state trace* behavior function $\mathcal{S}: \mathbf{Com} \rightarrow \mathcal{P}(S^+)$ is given by

$$\begin{aligned} \mathcal{S}[C] = \{ & s_0 s_1 \dots s_k \mid \langle C, s_0 \rangle \rightarrow^* \langle C_1, s_1 \rangle \\ & \rightarrow^* \dots \rightarrow^* \langle C_k, s_k \rangle \text{ term} \}. \end{aligned}$$

Partial correctness behavior induces a preorder $\sqsubseteq_{\mathcal{M}}$ and an equivalence relation $\equiv_{\mathcal{M}}$ on commands,

$$\begin{aligned} C \sqsubseteq_{\mathcal{M}} C' &\Leftrightarrow \forall s. (\text{free}[C] \cup \text{free}[C']) \subseteq \text{dom}(s) \\ &\Rightarrow \mathcal{M}[C]s \subseteq \mathcal{M}[C']s \\ C \equiv_{\mathcal{M}} C' &\Leftrightarrow C \sqsubseteq_{\mathcal{M}} C' \ \& \ C' \sqsubseteq_{\mathcal{M}} C, \end{aligned}$$

where we write $\mathcal{M}[C]s$ for $\{s' \mid (s, s') \in \mathcal{M}[C]\}$.

Similarly we define the state trace preorder and equivalence relation on commands:

$$\begin{aligned} C \sqsubseteq_{\mathcal{S}} C' &\Leftrightarrow \forall s. (\text{free}[C] \cup \text{free}[C']) \subseteq \text{dom}(s) \\ &\Rightarrow \mathcal{S}[C]s \subseteq \mathcal{S}[C']s \\ C \equiv_{\mathcal{S}} C' &\Leftrightarrow C \sqsubseteq_{\mathcal{S}} C' \ \& \ C' \sqsubseteq_{\mathcal{S}} C, \end{aligned}$$

where $\mathcal{S}[C]s = \{s_1 \dots s_k \mid s s_1 \dots s_k \in \mathcal{S}[C]\}$.

None of these relations is *substitutive*, since we have:

$$x := 1; x := x + 1 \equiv_{\mathcal{M}} x := 2$$

$$(x := 1; x := x + 1) \parallel x := 2 \not\equiv_{\mathcal{M}} x := 2 \parallel x := 2$$

$$x := 1; x := x + 1 \equiv_{\mathcal{S}} x := 1; x := 2$$

$$(x := 1; x := x + 1) \parallel x := 2 \not\equiv_{\mathcal{S}} (x := 1; x := 2) \parallel x := 2.$$

We therefore define the substitutive preorders $\leq_{\mathcal{M}}$ and $\leq_{\mathcal{S}}$, and the substitutive equivalence relations $\equiv_{\mathcal{M}}$ and $\equiv_{\mathcal{S}}$ by

$$C \leq_{\mathcal{M}} C' \Leftrightarrow \forall P[-]. (P[C] \sqsubseteq_{\mathcal{M}} P[C'])$$

$$C \equiv_{\mathcal{M}} C' \Leftrightarrow C \leq_{\mathcal{M}} C' \ \& \ C' \leq_{\mathcal{M}} C$$

$$C \leq_{\mathcal{S}} C' \Leftrightarrow \forall P[-]. (P[C] \sqsubseteq_{\mathcal{S}} P[C'])$$

$$C \equiv_{\mathcal{S}} C' \Leftrightarrow C \leq_{\mathcal{S}} C' \ \& \ C' \leq_{\mathcal{S}} C,$$

where $P[-]$ ranges over program contexts, that is, programs with a hole (denoted $[-]$) into which a command may be substituted; and $P[C]$ denotes the program obtained by substituting C into the hole. Thus $C \equiv_{\mathcal{M}} C'$ if and only if C and C' are interchangeable in all program contexts without affecting partial correctness. By definition, $\leq_{\mathcal{M}}$ and $\equiv_{\mathcal{M}}$ are *congruences*, since they are preserved by all programming language constructs.

Since $\mathcal{M}[C] = \{(s, s') \mid s s' \in \mathcal{S}[C]\}$ it is easy to see that $C \sqsubseteq_{\mathcal{S}} C' \Rightarrow C \sqsubseteq_{\mathcal{M}} C'$, and similarly state trace equivalence implies partial correctness equivalence. The converse fails, as shown by

$$x := 1; x := x + 1 \not\equiv_{\mathcal{S}} x := 2$$

$$x := 1; x := x + 1 \equiv_{\mathcal{M}} x := 2.$$

Nevertheless the substitutive preorders $\leq_{\mathcal{M}}$ and $\leq_{\mathcal{S}}$ coincide, so that $\equiv_{\mathcal{M}}$ and $\equiv_{\mathcal{S}}$ also coincide; two commands induce the same state traces in every context if and only if they induce the same partial correctness behavior in every context.

PROPOSITION 4.3. *For all commands C and C' , $C \leq_{\mathcal{M}} C'$ if and only if $C \leq_{\mathcal{S}} C'$.*

Proof. Since $C \sqsubseteq_{\mathcal{S}} C'$ implies $C \sqsubseteq_{\mathcal{M}} C'$ it follows that $C \leq_{\mathcal{S}} C'$ implies $C \leq_{\mathcal{M}} C'$. For the converse we argue as follows.

Assume $C \not\leq_{\mathcal{S}} C'$. Then there is a context $P[-]$ and a state trace $\sigma = s_0 \dots s_k$ such that $\sigma \in \mathcal{S}[P[C]]$ but $\sigma \notin \mathcal{S}[P[C']]$. Since states are finite, for each state s there is a boolean expression IS_s such that for all states s' such that $\text{dom}(s') \supseteq \text{dom}(s)$, $\langle IS_s, s' \rangle \rightarrow^* \text{tt}$ if s' and s agree on

$\text{dom}(s)$, and $\langle \text{IS}_{s_0}, s' \rangle \rightarrow^* \text{ff}$ otherwise. Then the program context

$$Q[-] = P[-] \parallel (\text{await IS}_{s_0} \text{ then skip}; \dots; \\ \text{await IS}_{s_k} \text{ then skip})$$

has the property that $(s_0, s_k) \in \mathcal{M}[Q[C]]$ but $(s_0, s_k) \notin \mathcal{M}[Q[C']]$. Thus $C \not\ll_{\mathcal{M}} C'$. ■

As a corollary, a semantics will be fully abstract with respect to state traces if and only if it is fully abstract with respect to partial correctness.

5. RESUMPTION SEMANTICS

Hennessy and Plotkin [9] gave a denotational semantics based on a domain R of “resumptions,” defined recursively by the domain equation

$$R = S \rightarrow \mathcal{P}(S + (R \times S)),$$

where \mathcal{P} is a suitable powerdomain constructor, $+$ denotes the separated sum and \times denotes the cartesian product of domains. The resumption semantic function $\mathcal{R}: \mathbf{Com} \rightarrow R$ is characterized operationally by the following property:

$$\mathcal{R}[C]s = \{s \mid \langle C, s \rangle \text{ term}\} \\ \cup \{(\mathcal{R}[C'], s') \mid \langle C, s \rangle \rightarrow \langle C', s' \rangle\}.$$

This shows clearly how resumptions model intermediate state changes. The resumption semantic function can also be given a denotational definition, since $\mathcal{R}[C_1 \parallel C_2]$ can be determined from $\mathcal{R}[C_1]$ and $\mathcal{R}[C_2]$. Moreover, $\mathcal{M}[C]$ can be extracted from $\mathcal{R}[C]$ by an “unravelling” operation, as described in [9].

However, the resumption semantics makes many unnecessary distinctions between programs: for instance **skip** and **skip; skip** denote different resumptions even though they induce the same partial correctness properties in all contexts. To attain full abstraction Hennessy and Plotkin added a form of “coroutine” composition $C_1 \text{ co } C_2$ to the syntax of the programming language, together with a non-deterministic choice operation $C_1 \text{ or } C_2$. The operational behavior of $C_1 \text{ co } C_2$ is to perform single atomic steps alternately from C_1 and C_2 until one of them terminates, and $C_1 \text{ or } C_2$ can behave either like C_1 or like C_2 . The transition rules for $C_1 \text{ co } C_2$ are

$$\frac{\langle C_1, s \rangle \rightarrow \langle C'_1, s' \rangle}{\langle C_{11} \text{ co } C_2, s \rangle \rightarrow \langle C_2 \text{ co } C'_1, s' \rangle} \\ \frac{\langle C_1, s \rangle \text{ term}}{\langle C_1 \text{ co } C_2, s \rangle \text{ term}}.$$

These two extra constructs permit program contexts to be built which can count the number of atomic actions taken by a command, thus distinguishing between **skip** and **skip; skip**. The resumptions model then becomes fully abstract for this extended language. Nevertheless, this coroutine construct seems rather *ad hoc* and difficult to motivate. The full abstraction problem for the original language remained open.

6. TRANSITION TRACES

The main problem with the resumptions model is that it represents explicitly the one-step transition relation \rightarrow and is therefore forced to distinguish between too many commands. Instead we design a semantic model based on the reflexive, transitive closure of the transition relation (denoted \rightarrow^*).

Informally, a (finite) *transition trace* of a command C is defined to be a sequence of pairs of states $(s_0, s'_0)(s_1, s'_1) \dots (s_k, s'_k)$ such that it is possible for C to perform a computation from s_0 to s'_k if execution is interrupted k times, the i th interruption changing the state from s'_i to s_{i+1} ($0 \leq i < k$). A transition trace of this form is *interference-free* iff $s'_i = s_{i+1}$ for each i , so that the trace corresponds to a computation in which the state is never changed by the command’s environment. The degenerate case ($k = 0$) yields simply a pair (s, s') such that C has a computation from s terminating in s' . Formally, we write $\mathcal{T}[C]$ for the set of transition traces of C , characterized operationally by

$$\mathcal{T}[C] = \{(s_0, s'_0)(s_1, s'_1) \dots (s_k, s'_k) \mid \\ \langle C, s_0 \rangle \rightarrow^* \langle C_1, s'_0 \rangle \ \& \\ \langle C_1, s_1 \rangle \rightarrow^* \langle C_2, s'_1 \rangle \ \& \\ \dots \ \& \\ \langle C_k, s_k \rangle \rightarrow^* \langle C', s'_k \rangle \text{ term}\}.$$

Clearly \mathcal{M} and \mathcal{S} can be extracted from \mathcal{T} , in the following sense: for all commands C ,

$$\mathcal{M}[C] = \{(s, s') \mid (s, s') \in \mathcal{T}[C]\} \\ \mathcal{S}[C] = \{s_0 s_1 \dots s_k \mid (s_0, s_1)(s_1, s_2) \dots (s_{k-1}, s_k) \in \mathcal{T}[C]\}.$$

The above operational characterization of \mathcal{T} has obvious but important consequences.

PROPOSITION 6.1. *The set of transition traces of a command C is closed under “stuttering” and “mumbling”: for all $\alpha, \beta \in (S \times S)^*$ and all $s, s', s'' \in S$,*

$$\alpha\beta \in \mathcal{T}[C] \Rightarrow \alpha(s, s) \beta \in \mathcal{T}[C] \\ \alpha(s, s')(s', s'') \beta \in \mathcal{T}[C] \Rightarrow \alpha(s, s'') \beta \in \mathcal{T}[C].$$

The first closure condition corresponds to reflexivity of \rightarrow^* , the second to transitivity. Since a stutter introduces a repetition and a mumble absorbs an intermediate state, the terminology is closely related to everyday usage.⁴ Given a set T of transition traces, we let T^\dagger , the *closure* of T , be the smallest set containing T and closed under stuttering and mumbling. We say that T is *closed* if $T = T^\dagger$. Proposition 6.1 thus states that $\mathcal{F}[[C]]$ is closed.

Let $\Sigma = S \times S$, and let $\mathcal{P}^\dagger(\Sigma^+)$ denote the set of closed sets of (non-empty) traces, ordered by inclusion. It is easy to see that this forms a complete lattice, with least element the empty set and with least upper bounds given by unions. We will show that the trace semantic function $\mathcal{F}: \mathbf{Com} \rightarrow \mathcal{P}^\dagger(\Sigma^+)$, characterized operationally above, can be defined compositionally.

Since **skip** causes no state change, no matter how many times it is interrupted, it is easy to see that the traces of **skip** are finite sequences of stuttering steps:

$$\begin{aligned} \mathcal{F}[[\mathbf{skip}]] &= \{(s_0, s_0) \dots (s_k, s_k) \mid k \geq 0 \ \& \\ &\quad \forall i. (0 \leq i \leq k \Rightarrow s_i \in S)\} \\ &= \{(s, s) \mid s \in S\}^\dagger. \end{aligned}$$

Similarly, since we assume that assignments and conditional critical regions are atomic,

$$\begin{aligned} \mathcal{F}[[I := E]] &= \{(s, [s \mid I = n]) \mid (s, n) \in \mathcal{E}[[E]]\}^\dagger \\ \mathcal{F}[[\mathbf{await} \ B \ \mathbf{then} \ C]] \\ &= \{(s, s') \mid (s, \tau\tau) \in \mathcal{B}[[B]] \ \& \ (s, s') \in \mathcal{F}[[C]]\}^\dagger. \end{aligned}$$

The traces of $C_1; C_2$ are built by concatenation. When T_1 and T_2 are closed sets of traces we define

$$T_1; T_2 = \{\alpha\beta \mid \alpha \in T_1 \ \& \ \beta \in T_2\}^\dagger,$$

denoting the smallest closed set of traces containing all relevant concatenations; the closure operator accounts for the possibility of mumbling between the end of a trace in T_1 and the start of a trace in T_2 . It then follows that $\mathcal{F}[[C_1; C_2]] = \mathcal{F}[[C_1]]; \mathcal{F}[[C_2]]$.

⁴ We should note, however, that our use of the term “stuttering” differs from Lamport [15]. In Lamport’s setting a program denotes a set of sequences of states, and a set $X \subseteq \mathcal{P}(S^+)$ is said to be closed under stuttering if for all ρ, ρ' in S^* and all $s \in S$,

$$\rho s \rho' \in X \Leftrightarrow \rho s s \rho' \in X.$$

The implication from left to right corresponds to reflexivity of \rightarrow^* , and the reverse implication to transitivity. Thus in Lamport’s terminology “stuttering” combines repetition of a state and absorption of an intermediate state.

The traces of a while-loop **while** B **do** C are obtained by closure from traces of form

$$\beta_1 \gamma_1 \beta_2 \gamma_2 \dots \beta_k \gamma_k \alpha,$$

where $k \geq 0$, each β_i corresponds to an evaluation of B to $\tau\tau$, each γ_i is a trace of C , and α corresponds to an evaluation of B to ff . We therefore extend the Kleene-star operation to closed sets of traces in the obvious way: T^* denotes the smallest set containing T and the empty trace,⁵ closed under stuttering, mumbling and concatenation. It then follows that

$$\mathcal{F}[[\mathbf{while} \ B \ \mathbf{do} \ C]] = (\mathcal{F}[[B]]; \mathcal{F}[[C]])^*; \mathcal{F}[[\neg B]],$$

where we write $\mathcal{F}[[B]] = \{(s, s) \mid (s, \tau\tau) \in \mathcal{B}[[B]]\}^\dagger$. Similarly we have

$$\begin{aligned} \mathcal{F}[[\mathbf{if} \ B \ \mathbf{then} \ C_1 \ \mathbf{then} \ C_2]] \\ &= \mathcal{F}[[B]]; \mathcal{F}[[C_1]] \cup \mathcal{F}[[\neg B]]; \mathcal{F}[[C_2]], \end{aligned}$$

writing $\mathcal{F}[[\neg B]]$ for $\{(s, s) \mid (s, \text{ff}) \in \mathcal{B}[[B]]\}^\dagger$.

The transition traces of $C_1 \parallel C_2$ are built by interleaving. For α and β in Σ^* let $\alpha \parallel \beta$ be the set of all transition traces built by interleaving α with β . This may be defined inductively as follows:

$$\begin{aligned} \alpha \parallel \varepsilon &= \varepsilon \parallel \alpha = \{\alpha\} \\ \sigma\alpha \parallel \rho\beta &= \{\sigma\gamma \mid \gamma \in \alpha \parallel \rho\beta\} \cup \{\rho\gamma' \mid \gamma' \in \sigma\alpha \parallel \beta\}, \end{aligned}$$

where σ and ρ range over Σ , α and β range over Σ^* , and ε is the empty trace. When T_1 and T_2 are closed sets of traces we define

$$T_1 \parallel T_2 = \bigcup \{\alpha \parallel \beta \mid \alpha \in T_1 \ \& \ \beta \in T_2\}^\dagger.$$

It then follows that $\mathcal{F}[[C_1 \parallel C_2]] = \mathcal{F}[[C_1]] \parallel \mathcal{F}[[C_2]]$.

We summarize these results in the following denotational description of \mathcal{F} .

PROPOSITION 6.2. *The (finite) transition traces semantic function $\mathcal{F}: \mathbf{Com} \rightarrow \mathcal{P}^\dagger(\Sigma^+)$ is characterized uniquely by the following clauses:*

$$\begin{aligned} \mathcal{F}[[\mathbf{skip}]] &= \{(s, s) \mid s \in S\}^\dagger \\ \mathcal{F}[[I := E]] &= \{(s, [s \mid I = n]) \mid (s, n) \in \mathcal{E}[[E]]\}^\dagger \\ \mathcal{F}[[C_1; C_2]] &= \mathcal{F}[[C_1]]; \mathcal{F}[[C_2]] \end{aligned}$$

⁵ Although transition traces are always non-empty, some of our definitions are simpler if we include the empty trace.

$$\begin{aligned}
\mathcal{T}[\text{if } B \text{ then } C_1 \text{ else } C_2] &= \mathcal{T}[B]; \mathcal{T}[C_1] \cup \mathcal{T}[\neg B]; \mathcal{T}[C_2] \\
\mathcal{T}[\text{while } B \text{ do } C] &= (\mathcal{T}[B]; \mathcal{T}[C])^*; \mathcal{T}[\neg B] \\
\mathcal{T}[C_1 \parallel C_2] &= \mathcal{T}[C_1] \parallel \mathcal{T}[C_2] \\
\mathcal{T}[\text{await } B \text{ then } C] &= \{(s, s') \in \mathcal{T}[C] \mid (s, s') \in \mathcal{T}[B]\}^\dagger.
\end{aligned}$$

Note that all operations on closed sets of traces used in this semantic definition are monotone (even continuous) with respect to set inclusion. An alternative (and equivalent) definition of the trace semantics of loops can be given using least fixed points. For each boolean expression B and command C , the function F on $\mathcal{P}^+(\Sigma^+)$ given by $F(T) = (\mathcal{T}[B]; \mathcal{T}[C]; T \cup \mathcal{T}[\neg B])$ is continuous. Its least fixed point is therefore

$$\begin{aligned}
\mu T. F(T) &= \bigcup_{n=0}^{\infty} F^n\{\} \\
&= (\mathcal{T}[B]; \mathcal{T}[C])^*; \mathcal{T}[\neg B] \\
&= \mathcal{T}[\text{while } B \text{ do } C].
\end{aligned}$$

It is also worth noting that the loop semantics can be obtained by taking the least fixed point of the analogous operation on (not necessarily closed) sets of traces and then taking the closure. This can be made precise as follows. For (arbitrary) trace sets T_1 and T_2 let $T_1 \cdot T_2 = \{\alpha\beta \mid \alpha \in T_1 \ \& \ \beta \in T_2\}$, so that $T_1; T_2 = (T_1 \cdot T_2)^\dagger$. Clearly the collection of (arbitrary) trace sets, ordered by inclusion, also forms a complete lattice. The function $F'(T) = (\mathcal{T}[B] \cdot \mathcal{T}[C] \cdot T \cup \mathcal{T}[\neg B])$ is continuous, with least fixed point given by

$$\mu T. F'(T) = (\mathcal{T}[B] \cdot \mathcal{T}[C])^* \cdot \mathcal{T}[\neg B],$$

and the closure of this set is $(\mathcal{T}[B]; \mathcal{T}[C])^*; \mathcal{T}[\neg B] = \mathcal{T}[\text{while } B \text{ do } C]$.

7. FULL ABSTRACTION

Given the assumption that expression evaluation is atomic, the only important aspect of an expression's operational behavior in the transition rules for commands is its final value. It follows trivially that two expressions induce the same partial correctness behavior in all program contexts if and only if they evaluate to the same results in all states. Thus, \mathcal{E} is fully abstract for the expression sub-language, and \mathcal{B} is fully abstract for the boolean expression sub-language.

We now show that the transition trace semantics for commands is fully abstract. We define $\mathcal{T}[C]s = \{s'\alpha \mid (s, s')\alpha \in \mathcal{T}[C]\}$ and

$$\begin{aligned}
C \sqsubseteq_{\mathcal{T}} C' &\Leftrightarrow \forall s. (\text{free}[C] \cup \text{free}[C'] \subseteq \text{dom}(s) \\
&\Rightarrow \mathcal{T}[C]s \subseteq \mathcal{T}[C']s) \\
C \equiv_{\mathcal{T}} C' &\Leftrightarrow C \sqsubseteq_{\mathcal{T}} C' \ \& \ C' \sqsubseteq_{\mathcal{T}} C.
\end{aligned}$$

PROPOSITION 7.1. *The transition traces semantics \mathcal{T} is inequationally fully abstract: for all commands C and C' , $C \sqsubseteq_{\mathcal{T}} C' \Leftrightarrow C \leq_{\mathcal{M}} C'$.*

Proof. Suppose $C \sqsubseteq_{\mathcal{T}} C'$. Since \mathcal{T} is a denotational semantics, for each program context $P[-]$ the only relevant aspect of C in determining $\mathcal{T}[P[C]]$ is $\mathcal{T}[C]$. Moreover, all operations used in the semantic definitions are monotone with respect to set inclusion. Thus we get $\mathcal{T}[P[C]] \subseteq \mathcal{T}[P[C']]$. But then for all relevant states s ,

$$\begin{aligned}
\mathcal{M}[P[C]]s &= \{s' \mid (s, s') \in \mathcal{T}[P[C]]\} \\
&\subseteq \{s' \mid (s, s') \in \mathcal{T}[P[C']]\} \\
&= \mathcal{M}[P[C']]s.
\end{aligned}$$

This shows that $C \sqsubseteq_{\mathcal{T}} C' \Leftrightarrow C \leq_{\mathcal{M}} C'$.

Recall that for each state s there is a boolean expression IS_s that evaluates to tt from s' if s' agrees with s on $\text{dom}(s)$, and evaluates to ff otherwise. Similarly there is a command MAKE_s such that

$$\langle \text{MAKE}_s, s' \rangle \rightarrow^* \langle \text{skip}, s \rangle$$

for all states such that $\text{dom}(s') = \text{dom}(s)$.

Now suppose $C \not\sqsubseteq_{\mathcal{T}} C'$, so that there is some transition trace $\alpha = (s_0, s'_0)(s_1, s'_1) \cdots (s_k, s'_k)$ belonging to $\mathcal{T}[C]$ and not to $\mathcal{T}[C']$. Let DO_α be the command

$$\begin{aligned}
&\text{await IS}_{s'_0} \text{ then MAKE}_{s_1}; \\
&\text{await IS}_{s'_0} \text{ then MAKE}_{s_2}; \\
&\dots \\
&\text{await IS}_{s'_{k-1}} \text{ then MAKE}_{s_k}.
\end{aligned}$$

The traces of DO_α are obtained by closure from the trace $\bar{\alpha} = (s'_0, s_1)(s'_1, s_2) \cdots (s'_{k-1}, s_k)$. Let $P_\alpha[-]$ be the program context $[-] \parallel \text{DO}_\alpha$. Each trace of $P_\alpha[C]$ is obtained by closure from the result of interleaving a trace of C with $\bar{\alpha}$. In particular, since $\alpha \in \mathcal{T}[C]$ it follows that the interference-free trace

$$(s_0, s'_0)(s'_0, s_1)(s_1, s'_1)(s'_1, s_2) \cdots (s'_{k-1}, s_k)(s_k, s'_k)$$

is possible for $P_\alpha[C]$. Since trace sets are closed under mumbering, it follows that $(s_0, s'_k) \in \mathcal{M}[P[C]]$.

However, (s_0, s'_k) cannot belong to the set $\mathcal{M}[P[C']]$, since this would require C' to have a trace α' such that interleaving α' with $\bar{\alpha}$ can yield an interference-free trace of form

$$(s_0, t_1)(t_1, t_2)(t_2, t_3)\dots(t_{r-1}, t_r)(t_r, s'_k).$$

But it is easy to see that this is possible only if α itself is derivable from α' by stuttering and mumbling. It would then follow that $\alpha \in \mathcal{F}[C']$, contradicting our assumption.

Thus, $C \not\sqsubseteq_{\mathcal{F}} C'$ implies $C \not\sqsubseteq_{\mathcal{M}} C'$. That completes the proof. ■

For example, consider the commands $C = x := 1$; $x := x + 1$ and $C' = x := 1$; $x := 2$. They have the same partial correctness behavior but different transition traces: clearly

$$\alpha = ([x = 0], [x = 1])([x = 0], [x = 1])$$

is a trace of C but not of C' . The context $P_{\alpha}[-]$ built in the proof above is

$$[-] \parallel \mathbf{await} \ x = 1 \ \mathbf{then} \ x := 0$$

and it is clear that $P_{\alpha}[C]$ may terminate with $x = 1$ but that $P_{\alpha}[C']$ cannot.

Similarly, consider the commands $x := 0$ and $x := 0$; $x := 0$. It is easy to see that

$$\mathcal{F}[x := 0] \subseteq \mathcal{F}[x := 0; x := 0],$$

and this inclusion is proper. The trace $([x = 1], [x = 0])([x = 1], [x = 0])$ is possible for $x := 0$; $x := 0$ but not for $x := 0$. These two commands can be distinguished by running them in parallel with the command $\mathbf{await} \ x = 0 \ \mathbf{then} \ x := 1$.

8. LAWS OF PARALLEL PROGRAMMING

We can use trace semantics to prove equations and inequations between programs, with the guarantee that these laws may be used safely for reasoning about partial correctness, or state traces, in any program context. The following Proposition summarizes some laws, writing \equiv for $\equiv_{\mathcal{F}}$ (or, equivalently, for $\equiv_{\mathcal{M}}$ or $\equiv_{\mathcal{S}}$) and \sqsubseteq for $\sqsubseteq_{\mathcal{F}}$ (or, equivalently, for $\sqsubseteq_{\mathcal{M}}$ or $\sqsubseteq_{\mathcal{S}}$).

PROPOSITION 8.1. *The following laws are valid:*

$$\begin{aligned} \mathbf{skip}; C &\equiv C \equiv C; \mathbf{skip} \\ (C_1; C_2); C_3 &\equiv C_1; (C_2; C_3) \\ C \parallel \mathbf{skip} &\equiv C \\ C_1 \parallel C_2 &\equiv C_2 \parallel C_1 \end{aligned}$$

$$\begin{aligned} (C_1 \parallel C_2) \parallel C_3 &\equiv C_1 \parallel (C_2 \parallel C_3) \\ (\mathbf{if} \ B \ \mathbf{then} \ C_1 \ \mathbf{else} \ C_2); C &\equiv \mathbf{if} \ B \ \mathbf{then} \ C_1; C \ \mathbf{else} \ C_2; C \\ \mathbf{if} \ (B_1 \ \& \ B_2) \ \mathbf{then} \ C_1 \ \mathbf{else} \ C_2 \\ &\sqsubseteq \mathbf{if} \ B_1 \ \mathbf{then} \ (\mathbf{if} \ B_2 \ \mathbf{then} \ C_1 \ \mathbf{else} \ C_2) \ \mathbf{else} \ C_2 \\ \mathbf{while} \ B \ \mathbf{do} \ C &\equiv \mathbf{if} \ B \ \mathbf{then} \ C; \ \mathbf{while} \ B \ \mathbf{do} \ C \ \mathbf{else} \ \mathbf{skip}. \end{aligned}$$

Proof. The laws may be easily validated, taking advantage of natural algebraic identities involving $T_1; T_2$, $T_1 \parallel T_2$, and T^* . For example, \mathbf{skip} is a unit for sequential and parallel composition because trace sets are closed under stuttering and mumbling. It is obvious from the definitions that for all closed sets T_1 , T_2 , and T_3 , we have

$$\begin{aligned} (T_1; T_2); T_3 &= T_1; (T_2; T_3) \\ &= \{\alpha\beta\gamma \mid \alpha \in T_1 \ \& \ \beta \in T_2 \ \& \ \gamma \in T_3\}^{\dagger} \\ (T_1 \parallel T_2) \parallel T_3 &= T_1 \parallel (T_2 \parallel T_3) \\ &= \bigcup \{\alpha \parallel \beta \parallel \gamma \mid \alpha \in T_1 \ \& \ \beta \in T_2 \ \& \ \gamma \in T_3\}^{\dagger} \\ T_1 \parallel T_2 &= T_2 \parallel T_1, \end{aligned}$$

since concatenation and interleaving operations on transition traces are clearly associative, and interleaving of transition traces is obviously a symmetric operation. The rules for conditionals and loops are straightforward, since $(T_1 \cup T_2); T = (T_1; T) \cup (T_2; T)$ and $\mathcal{F}[B_1 \ \& \ B_2] \subseteq \mathcal{F}[B_1]; \mathcal{F}[B_2]$. ■

In addition, since for all $\alpha, \beta, \gamma \in \Sigma^*$ we have $\alpha(\beta \parallel \gamma) \sqsubseteq (\alpha\beta) \parallel \gamma$, we obtain the following law:

$$C_1; (C_2 \parallel C_3) \sqsubseteq (C_1; C_2) \parallel C_3,$$

from this law, we may derive the inequation $C_1; C_2 \sqsubseteq C_1 \parallel C_2$.

Since the semantics is tailored to correspond to partial correctness, and we model blocking as busy waiting, the semantics also validates the following laws:

$$\begin{aligned} \mathbf{while} \ \mathbf{true} \ \mathbf{do} \ \mathbf{skip} &\sqsubseteq C \\ \mathbf{while} \ \mathbf{true} \ \mathbf{do} \ \mathbf{skip} &\equiv \mathbf{await} \ \mathbf{false} \ \mathbf{then} \ C. \end{aligned}$$

This is reasonable, since a diverging program satisfies only vacuous partial correctness properties, and induces non-termination in every program context.

Since assignment is atomic, this semantics validates the law $I := I \equiv \mathbf{skip}$. Since boolean expression evaluation is atomic, we can also show that

$$\begin{aligned} \mathcal{F}[B]; \mathcal{F}[B] &= \mathcal{F}[B] \\ \mathcal{F}[B]; \mathcal{F}[\neg B] &\subseteq \mathcal{F}[B] \cap \mathcal{F}[\neg B], \end{aligned}$$

and the semantics validates the following laws:

$$\begin{aligned}
& \text{while } B \text{ do } C \\
& \equiv \text{if } B \text{ then while } B \text{ do } C \text{ else skip} \\
& \text{if } B \text{ then } C_1 \text{ else } C_2 \\
& \equiv \text{if } B \text{ then (if } B \text{ then } C_1 \text{ else } C_2) \text{ else } C_2.
\end{aligned}$$

If the expression language is deterministic, so that for all E and s the set $\mathcal{E}[E]s$ contains at most one value, we also obtain the inequation

$$I := [E_1/I]E_2 \sqsubseteq I := E_1; I := E_2,$$

where $[E_1/I]E_2$ denotes the expression obtained by substituting E_1 for each free occurrence of I in E_2 , with appropriate changes of bound variable to avoid capturing any free identifiers of E_1 .

By way of contrast, we discuss briefly some results of de Bakker [6] concerning the equivalence in all *sequential* contexts of sequences of simple assignments, and we compare what happens in the parallel case. A simple assignment has form $I := I'$, where I and I' are (not necessarily distinct) identifiers. The following list of axioms is given by de Bakker:

$$\begin{aligned}
x := y; y := x & \equiv x := y \\
x := y; x := z & \equiv x := z \quad (x \neq z) \\
x := y; z := x & \equiv x := y; z := y \\
x := y; z := y & \equiv z := y; x := y
\end{aligned}$$

Together with some simple rules, such as

$$\frac{C_1; x := x' \equiv C_2; x := x' \quad C_1; y := y' \equiv C_2; y := y' \quad x \neq y}{C_1 \equiv C_2}$$

this gives a complete axiomatization of equivalence in all sequential contexts. That is, two sequences of simple assignments induce the same behavior in all sequential contexts if and only if their equivalence can be proven from these axioms and rules.

However, each of these axioms fails for parallel contexts:

- The first of de Bakker's axioms is invalid in the context $[- \parallel x := 1]$.
- The second is invalid in the context $[- \parallel w := x]$.
- The third is invalid in the context $[- \parallel x := 0]$.
- The fourth is invalid in the context $[- \parallel y := 0]$.

Instead, the following inequational versions of the first two axioms hold in all parallel contexts:

$$\begin{aligned}
x := y & \sqsubseteq x := y; y := x \\
x := z & \sqsubseteq x := y; x := z.
\end{aligned}$$

The second of these is itself a special case of the law given above for $I := E_1; I := E_2$. The third and fourth de Bakker axioms cannot be weakened to an inequality in either direction. For instance,

$$x := y; z := x \sqsubseteq x := y; z := y$$

is invalid in the context $[- \parallel y := y + 1]$ and the reverse inequality fails in the context $[- \parallel x := x + 1]$.

9. FINER GRANULARITY

Our semantics can be adapted to deal with finer levels of granularity. For instance, we might allow interruption of an assignment $I := E$ during the evaluation of E , and interruption of a conditional during the evaluation of its test. To make the discussion precise, suppose that we have the following abstract syntax for boolean expressions and integer expressions:

$$\begin{aligned}
B & ::= \text{true} \mid \text{false} \mid \neg B \mid B_1 \& B_2 \mid E_1 \leq E_2 \\
E & ::= 0 \mid 1 \mid I \mid E_1 + E_2 \mid \text{if } B \text{ then } E_1 \text{ else } E_2.
\end{aligned}$$

To adapt the operational semantics we introduce the set \mathbf{BExp}' of extended boolean expressions, defined by adding the clauses $B ::= v$ ($v \in V$) to the grammar for \mathbf{BExp} , and the set \mathbf{Exp}' of extended integer expressions, defined by adding $E ::= n$ ($n \in \mathbb{N}$) to the grammar for \mathbf{Exp} . We use configurations of form $\langle E, s \rangle$ and $\langle B, s \rangle$, where E and B are extended expressions. A configuration of form $\langle n + E_2, s \rangle$ (with $n \in \mathbb{N}$) represents a stage in evaluation of a sum expression where the left-hand expression has been evaluated to the integer n and the right-hand expression remaining to be evaluated is E_2 ; a configuration of form $n \in \mathbb{N}$ represents the final result of evaluation.

A fine-grained transition system for expressions is described in Figs. 2 and 3. Note that the transition rules specify that a conjunction $B_1 \& B_2$ is evaluated from left-to-right with a short-circuit strategy, avoiding evaluation of B_2 if B_1 evaluates to ff . On the other hand we specify that in a sum expression $E_1 + E_2$ the two sub-expressions are evaluated in parallel. These choices were made solely for illustration, and the transition rules and subsequent semantic definitions may easily be modified to model different evaluation strategies.

Now that expression evaluation is no longer atomic, the semantic functions \mathcal{E} and \mathcal{B} are not fully abstract. Instead

$$\begin{array}{c}
\langle \text{true}, s \rangle \rightarrow \text{tt} \\
\langle \text{false}, s \rangle \rightarrow \text{ff} \\
\frac{\langle B, s \rangle \rightarrow \langle B', s \rangle}{\langle \neg B, s \rangle \rightarrow \langle \neg B', s \rangle} \\
\frac{\langle B, s \rangle \rightarrow \text{tt}}{\langle \neg B, s \rangle \rightarrow \text{ff}} \\
\frac{\langle B, s \rangle \rightarrow \text{ff}}{\langle \neg B, s \rangle \rightarrow \text{tt}} \\
\frac{\langle B_1, s \rangle \rightarrow \langle B'_1, s \rangle}{\langle B_1 \& B_2, s \rangle \rightarrow \langle B'_1 \& B_2, s \rangle} \\
\frac{\langle B_1, s \rangle \rightarrow \text{tt}}{\langle B_1 \& B_2, s \rangle \rightarrow \langle B_2, s \rangle} \\
\frac{\langle B_1, s \rangle \rightarrow \text{ff}}{\langle B_1 \& B_2, s \rangle \rightarrow \text{ff}} \\
\frac{\langle E_1, s \rangle \rightarrow \langle E'_1, s \rangle}{\langle E_1 \leq E_2, s \rangle \rightarrow \langle E'_1 \leq E_2, s \rangle} \\
\frac{\langle E_2, s \rangle \rightarrow \langle E'_2, s \rangle}{\langle E_1 \leq E_2, s \rangle \rightarrow \langle E_1 \leq E'_2, s \rangle} \\
\langle m \leq n, s \rangle \rightarrow \text{tt} \quad \text{if } m \leq n \\
\langle m \leq n, s \rangle \rightarrow \text{ff} \quad \text{if } m > n
\end{array}$$

FIG. 2. Fine-grained transition rules for boolean expressions.

we need to extend the transition traces semantics to cover expressions, to allow for the possibility that the state may change during evaluation. Since we assume that expression evaluation never causes any side-effects, we can use a slightly simpler trace structure than for commands:⁶

$$\begin{array}{c}
\mathcal{T}[[B]] = \{((s_0, s_0)(s_1, s_1)\dots(s_k, s_k), v) | \\
\langle B, s_0 \rangle \rightarrow^* \langle B_1, s_0 \rangle \& \\
\langle B_1, s_1 \rangle \rightarrow^* \langle B_2, s_1 \rangle \& \\
\dots \& \\
\langle B_k, s_k \rangle \rightarrow^* v\}
\end{array}$$

⁶ Actually, we could have used traces of form $(s_0 s_1 \dots s_k, v)$, with minor modifications in what follows. Our notation is deliberately chosen for uniformity, so as to simplify some of the details that follow.

$$\begin{array}{c}
\langle 0, s \rangle \rightarrow 0 \\
\langle 1, s \rangle \rightarrow 1 \\
\langle I, s \rangle \rightarrow s[I] \\
\frac{\langle B, s \rangle \rightarrow \langle B', s \rangle}{\langle \text{if } B \text{ then } E_1 \text{ else } E_2, s \rangle \rightarrow \langle \text{if } B' \text{ then } E_1 \text{ else } E_2, s \rangle} \\
\frac{\langle B, s \rangle \rightarrow \text{tt}}{\langle \text{if } B \text{ then } E_1 \text{ else } E_2, s \rangle \rightarrow \langle E_1, s \rangle} \\
\frac{\langle B, s \rangle \rightarrow \text{ff}}{\langle \text{if } B \text{ then } E_1 \text{ else } E_2, s \rangle \rightarrow \langle E_2, s \rangle} \\
\frac{\langle E_1, s \rangle \rightarrow \langle E'_1, s \rangle}{\langle E_1 + E_2, s \rangle \rightarrow \langle E'_1 + E_2, s \rangle} \\
\frac{\langle E_2, s \rangle \rightarrow \langle E'_2, s \rangle}{\langle E_1 + E_2, s \rangle \rightarrow \langle E_1 + E'_2, s \rangle} \\
\langle m + n, s \rangle \rightarrow k \quad \text{if } m + n = k
\end{array}$$

FIG. 3. Fine-grained transition rules for integer expressions.

$$\begin{array}{c}
\mathcal{T}[[E]] = \{((s_0, s_0)(s_1, s_1)\dots(s_k, s_k), n) | \\
\langle E, s_0 \rangle \rightarrow^* \langle E_1, s_0 \rangle \& \\
\langle E_1, s_1 \rangle \rightarrow^* \langle E_2, s_1 \rangle \& \\
\dots \& \\
\langle E_k, s_k \rangle \rightarrow^* n\}
\end{array}$$

Thus a trace $((s_0, s_0)(s_1, s_1)\dots(s_k, s_k), v) \in \mathcal{T}[[B]]$ means that there is an evaluation of B from initial state s_0 resulting in value v , during which the environment makes k interruptions, the i th interruption changing the state to s_i . In particular allowing no interruptions corresponds to the definition of \mathcal{B} , and $\mathcal{B}[[B]] = \{(s, n) \mid ((s, s), n) \in \mathcal{T}[[B]]\}$. Note that the traces of an expression are again closed under (the obvious analogues of) stuttering and mumbling. For boolean expressions this amounts to the following:

PROPOSITION 9.1. *For all boolean expressions B , all states s , all $\alpha, \beta \in \Sigma^*$, and all truth values v ,*

$$\begin{array}{c}
(\alpha\beta, v) \in \mathcal{T}[[B]] \Rightarrow (\alpha(s, s)\beta, v) \in \mathcal{T}[[B]] \\
(\alpha(s, s)(s, s)\beta, v) \in \mathcal{T}[[B]] \Rightarrow (\alpha(s, s)\beta, v) \in \mathcal{T}[[B]].
\end{array}$$

We write $\mathcal{P}^+(\Sigma^+ \times V)$ for the set of closed sets, ordered again by inclusion. Similar properties hold for integer expressions, so that $\mathcal{T}[[E]]$ is a closed subset of $\Sigma^+ \times N$.

So far we have characterized $\mathcal{T}[[B]]$ and $\mathcal{T}[[E]]$ operationally. As with commands, we can also give denotational definitions. We give the details only for boolean expressions.

PROPOSITION 9.2. *The fine-grained trace semantics $\mathcal{T}: \mathbf{BExp} \rightarrow \mathcal{P}^+(\Sigma^+ \times V)$ is uniquely characterized by the following clauses:*

$$\begin{aligned} \mathcal{T}[\mathbf{true}] &= \{(s, s), \text{tt}\} \mid s \in S\}^\dagger \\ \mathcal{T}[\mathbf{false}] &= \{(s, s), \text{ff}\} \mid s \in S\}^\dagger \\ \mathcal{T}[\neg B] &= \{(\alpha, \neg v) \mid (\alpha, v) \in \mathcal{T}[B]\}, \\ &\text{where } \neg \text{tt} = \text{ff}, \neg \text{ff} = \text{tt} \\ \mathcal{T}[B_1 \ \& \ B_2] &= \{(\alpha, \text{ff}) \mid (\alpha, \text{ff}) \in \mathcal{T}[B_1]\} \\ &\cup \{(\alpha\beta, v) \mid (\alpha, \text{tt}) \in \mathcal{T}[B_1] \ \& \ (\beta, v) \in \mathcal{T}[B_2]\}^\dagger \\ \mathcal{T}[E_1 \leq E_2] &= \{(\gamma, m \leq n) \mid (\alpha, m) \in \mathcal{T}[E_1] \ \& \\ &(\beta, n) \in \mathcal{T}[E_2] \ \& \ \gamma \in \alpha \parallel \beta\}^\dagger. \end{aligned}$$

An operational characterization of the fine-grained trace semantics of commands is given exactly as before, but using the fine-grained transition relation \rightarrow from Fig. 4:

$$\begin{aligned} \mathcal{T}[C] &= \{(s_0, s'_0)(s_1, s'_1) \dots (s_k, s'_k) \mid \\ &\langle C, s_0 \rangle \rightarrow^* \langle C_1, s'_0 \rangle \ \& \\ &\langle C_1, s_1 \rangle \rightarrow^* \langle C_2, s'_1 \rangle \ \& \\ &\dots \ \& \\ &\langle C_k, s_k \rangle \rightarrow^* \langle C', s'_k \rangle \text{ term}\}. \end{aligned}$$

PROPOSITION 9.3. *The fine-grained trace semantics of commands is uniquely characterized by the following clauses:*

$$\begin{aligned} \mathcal{T}[\mathbf{skip}] &= \{(s, s) \mid s \in S\}^\dagger \\ \mathcal{T}[I := E] &= \{\alpha(s, [s \mid I = n]) \mid (\alpha, n) \in \mathcal{T}[E]\}^\dagger \\ \mathcal{T}[C_1; C_2] &= \mathcal{T}[C_1]; \mathcal{T}[C_2] \\ \mathcal{T}[\mathbf{if} \ B \ \mathbf{then} \ C_1 \ \mathbf{else} \ C_2] \\ &= \mathcal{T}[B]; \mathcal{T}[C_1] \cup \mathcal{T}[\neg B]; \mathcal{T}[C_2] \\ \mathcal{T}[\mathbf{while} \ B \ \mathbf{do} \ C] &= (\mathcal{T}[B]; \mathcal{T}[C])^*; \mathcal{T}[\neg B] \\ \mathcal{T}[C_1 \parallel C_2] &= \mathcal{T}[C_1] \parallel \mathcal{T}[C_2] \\ \mathcal{T}[\mathbf{await} \ B \ \mathbf{then} \ C] &= \{(s, s') \in \mathcal{T}[C] \mid (s, s) \in \mathcal{T}[B]\}^\dagger. \end{aligned}$$

Here we write $\mathcal{T}[B]$ for the set $\{\alpha \mid (\alpha, \text{tt}) \in \mathcal{T}[B]\}$ and $\mathcal{T}[\neg B]$ for $\{\alpha \mid (\alpha, \text{ff}) \in \mathcal{T}[B]\}$.

Again all operations on trace sets used in this semantics are continuous with respect to set inclusion.

Of course, since the operational semantics of commands is now fine-grained, we are now interested in notions of behavior defined as before but based instead on the fine-grained transition relation of Fig. 4.

$$\begin{aligned} &\langle \mathbf{skip}, s \rangle \text{term} \\ \frac{\langle E, s \rangle \rightarrow \langle E', s \rangle}{\langle I := E, s \rangle \rightarrow \langle I := E', s \rangle} \quad \frac{\langle E, s \rangle \rightarrow n}{\langle I := E, s \rangle \rightarrow \langle I := n, s \rangle} \\ &\langle I := n, s \rangle \rightarrow \langle \mathbf{skip}, [s \mid I = n] \rangle \\ \frac{\langle C_1, s \rangle \rightarrow \langle C'_1, s' \rangle}{\langle C_1; C_2, s \rangle \rightarrow \langle C'_1; C_2, s' \rangle} \quad \frac{\langle C_1, s \rangle \text{term}}{\langle C_1; C_2, s \rangle \rightarrow \langle C_2, s \rangle} \\ &\frac{\langle B, s \rangle \rightarrow \langle B', s \rangle}{\langle \mathbf{if} \ B \ \mathbf{then} \ C_1 \ \mathbf{else} \ C_2, s \rangle \rightarrow \langle \mathbf{if} \ B' \ \mathbf{then} \ C_1 \ \mathbf{else} \ C_2, s \rangle} \\ &\frac{\langle B, s \rangle \rightarrow \text{tt}}{\langle \mathbf{if} \ B \ \mathbf{then} \ C_1 \ \mathbf{else} \ C_2, s \rangle \rightarrow \langle C_1, s \rangle} \\ &\frac{\langle B, s \rangle \rightarrow \text{ff}}{\langle \mathbf{if} \ B \ \mathbf{then} \ C_1 \ \mathbf{else} \ C_2, s \rangle \rightarrow \langle C_2, s \rangle} \\ &\langle \mathbf{while} \ B \ \mathbf{do} \ C, s \rangle \rightarrow \langle \mathbf{if} \ B \ \mathbf{then} \ C; \mathbf{while} \ B \ \mathbf{do} \ C \ \mathbf{else} \ \mathbf{skip}, s \rangle \\ \frac{\langle C_1, s \rangle \rightarrow \langle C'_1, s' \rangle}{\langle C_1 \parallel C_2, s \rangle \rightarrow \langle C'_1 \parallel C_2, s' \rangle} \quad \frac{\langle C_2, s \rangle \rightarrow \langle C'_2, s' \rangle}{\langle C_1 \parallel C_2, s \rangle \rightarrow \langle C_1 \parallel C'_2, s' \rangle} \\ &\frac{\langle C_1, s \rangle \text{term} \ \langle C_2, s \rangle \text{term}}{\langle C_1 \parallel C_2, s \rangle \text{term}} \\ &\frac{\langle B, s \rangle \rightarrow^* \text{tt} \ \langle C, s \rangle \rightarrow^* \langle C', s' \rangle \text{term}}{\langle \mathbf{await} \ B \ \mathbf{then} \ C, s \rangle \rightarrow \langle C', s' \rangle} \\ &\frac{\langle B, s \rangle \rightarrow^* \text{ff}}{\langle \mathbf{await} \ B \ \mathbf{then} \ C, s \rangle \rightarrow \langle \mathbf{await} \ B \ \mathbf{then} \ C, s \rangle} \end{aligned}$$

FIG. 4. Fine-grained transition rules for commands.

PROPOSITION 9.4. *The fine-grained semantics is fully abstract: for all terms t and t' of the same syntactic type, $t \sqsubseteq_{\mathcal{T}} t' \Leftrightarrow t \leq_{\mathcal{H}} t'$.*

Proof. For commands the proof is similar to the proof of Proposition 7.1.

For boolean expressions t and t' with different transition traces it is easy to construct a context of form $C \parallel \mathbf{if} \ [-] \ \mathbf{then} \ z := 0 \ \mathbf{else} \ z := 1$ (for a suitably chosen C and z) that distinguishes between them.

For integer expressions with different transition traces we can find a discriminating context of form $C \parallel z := [-]$. ■

For example, the boolean expressions $x \leq x$ and \mathbf{true} are not semantically equivalent: $([x = 0], [x = 0])([x = 1], [x = 1]), \text{ff} \in \mathcal{T}[x \leq x] - \mathcal{T}[\mathbf{true}]$. As a result they may induce different behavior in contexts such as

$$x := 1 \parallel \mathbf{if} \ [-] \ \mathbf{then} \ z := 0 \ \mathbf{else} \ z := 1.$$

The relationships given in Proposition 8.1 continue to hold for the fine-grained semantics. However, the identity $I := I \equiv \text{skip}$ fails because assignment is not atomic. For example,

$$x := 0; (x := x \parallel x := 1) \not\equiv_{\text{H}} x := 0; (\text{skip} \parallel x := 1),$$

because $([x=0], [x=0])([x=1], [x=0])$ is a trace of $x := x$ but not of skip . Instead we get the inequality $\text{skip} \sqsubseteq I := I$. Similarly, we have only an inclusion $\mathcal{F} \llbracket B \rrbracket; \mathcal{F} \llbracket B \rrbracket \supseteq \mathcal{F} \llbracket B \rrbracket$, the converse direction failing in general; and we obtain the following inequalities

$$\begin{aligned} & \text{while } B \text{ do } C \\ & \quad \sqsubseteq \text{if } B \text{ then while } B \text{ do } C \text{ else skip} \\ & \text{if } B \text{ then } C_1 \text{ else } C_2 \\ & \quad \sqsubseteq \text{if } B \text{ then (if } B \text{ then } C_1 \text{ else } C_2) \text{ else } C_2. \end{aligned}$$

10. FAIRNESS

So far we have ignored the possibility of infinite computation and non-termination. This was appropriate for reasoning about partial correctness, and in general for any property determined entirely by the finite traces of a program. However, many parallel programs are designed specifically not to terminate, and we would like a semantics suitable for reasoning about total correctness, and about safety and liveness properties, in addition to partial correctness.

When reasoning about parallel programs it is often natural to make a *fairness* assumption [19]: when running commands in parallel, no individual command is forever denied its turn for execution. This kind of assumption permits us to abstract away from scheduling details, such as the relative speeds of parallel processes, and to deduce program properties that hold in any “realistic” implementation of the programming language. At least for the form of fairness described informally above, it is easy to design schedulers that guarantee fair execution, so that the fairness assumption provides a convenient and reasonable abstraction. In this section we extend our semantics to incorporate finite and infinite traces corresponding to fair executions. We provide an operational characterization of the set of fair traces of a program, followed by a denotational characterization. Since we model blocking as busy waiting, there is no need to distinguish between “weak” and “strong” fairness: a non-terminated command is *always* enabled, since even a command stuck at an *await* with its test false can perform an idle step.

In order to characterize the fair infinite computations of a command operationally, care must be taken to keep track of which parallel component performs each atomic action

(see, for example, [8]). The crucial case is for a command of form $C_1 \parallel C_2$. An infinite transition sequence of $C_1 \parallel C_2$ is fair if and only if for each $i = 1, 2$ the sub-sequence of steps involving C_i is either finite and ends in a successfully terminated configuration, or is infinite.

For example, consider the program

$$C = (x := 0 \parallel \text{while } x = 1 \text{ do } y := y + 1)$$

started in initial state $[x=1, y=1]$. This program has infinitely many terminating computations, in which the loop body is executed a finite number of times and then the assignment to x occurs, causing the loop to terminate. For each positive integer k there is such a computation ending in the state $[x=0, y=k]$. There is also an infinite computation of C in which the loop body is repeated forever. However, this computation is *unfair*, since it can only occur if the assignment to x is forever denied a chance to execute. Under the fairness assumption this program is guaranteed to terminate. Since the final value of y is an arbitrary positive integer this program exhibits unbounded non-determinism.

For another example, the program

$$\begin{aligned} & x := 0; y := 0; (\text{await } x = 1 \text{ then } y := 1 \parallel \\ & \quad \text{while } y = 0 \text{ do } x := 1 - x) \end{aligned}$$

has a non-terminating weakly fair computation, in which the *await* command only gets to move after an even number of loop iterations, when x is equal to 0, so that it never changes the value of y . This shows that even under fairness assumptions it can be possible for blocking to persist.

We write

$$\langle C_0, s_0 \rangle \rightarrow \langle C_1, s_1 \rangle \rightarrow \dots \langle C_n, s_n \rangle \rightarrow \dots \text{fair}$$

to indicate a fair infinite computation. Similarly, we write

$$\langle C_0, s_0 \rangle \rightarrow^* \langle C_1, s_1 \rangle \rightarrow^* \dots \langle C_n, s_n \rangle \rightarrow^* \dots \text{fair}$$

to indicate an infinite sub-sequence of a fair infinite computation; in particular, this indicates that infinitely many of the transition sequences $\langle C_i, s_i \rangle \rightarrow^* \langle C_{i+1}, s_{i+1} \rangle$ involve a non-zero number of steps. Finally, we generalize further to non-consecutive transition sequences, writing

$$\begin{aligned} & \langle C_0, s_0 \rangle \rightarrow^* \langle C_1, s'_0 \rangle \& \\ & \langle C_1, s_1 \rangle \rightarrow^* \langle C_2, s'_1 \rangle \& \dots \& \\ & \langle C_n, s_n \rangle \rightarrow^* \langle C_{n+1}, s'_n \rangle \dots \text{fair} \end{aligned}$$

when there is a fair computation of C_0 from initial state s_0 during which execution is interrupted infinitely often, the i th interruption changing the state from s'_i to s_{i+1} (for each $i \geq 0$). Each (s_i, s'_i) represents a finite (possibly empty) sequence of atomic actions performed by the command, and infinitely many of these action sequences must be non-empty.

This provides a natural generalization of transition traces to encompass fairness. We therefore define $\Sigma^\infty = \Sigma^+ \cup \Sigma^\omega$, the set of finite or infinite transition traces, and we characterize the fair trace semantics operationally as follows:

$$\begin{aligned} \mathcal{T}[C] = & \{(s_0, s'_0)(s_1, s'_1)\dots(s_k, s'_k) \mid \\ & \langle C, s_0 \rangle \rightarrow^* \langle C_1, s'_0 \rangle \& \dots \& \\ & \langle C_k, s_k \rangle \rightarrow^* \langle C', s'_k \rangle \text{ term}\} \cup \\ & \{(s_0, s'_0)\dots(s_n, s'_n)\dots \mid \\ & \langle C, s_0 \rangle \rightarrow^* \langle C_1, s'_0 \rangle \& \dots \& \\ & \langle C_n, s_n \rangle \rightarrow^* \langle C_{n+1}, s'_n \rangle \& \dots \text{fair}\}. \end{aligned}$$

Note that C has an infinite interference-free trace beginning in state s iff $\langle C, s \rangle$ has a fair infinite computation from s .

For obvious reasons only finitely many interruptions can occur between successive atomic actions by C ; consequently, $\mathcal{T}[C]$ is again closed under stuttering and mumbling, where we allow finitely many stutters or mumbles between successive stages in a trace, and we allow use of a closure operation at infinitely many positions in an infinite trace. Formally, we extend the definition of closure to subsets of Σ^∞ as follows. A set T of finite and infinite traces is closed if and only if T satisfies the conditions of Proposition 6.1 and also

$$\begin{aligned} \alpha_0 \alpha_1 \dots \alpha_n \dots \in T \\ \Rightarrow \alpha_0(s_0, s_0) \alpha_1(s_1, s_1) \dots \alpha_n(s_n, s_n) \dots \in T \\ \alpha_0(s_0, s'_0)(s'_0, s''_0) \alpha_1(s_1, s'_1)(s'_1, s''_1) \alpha_2 \dots \in T \\ \Rightarrow \alpha_0(s_0, s''_0) \alpha_1(s_1, s'_1) \alpha_2 \dots \in T. \end{aligned}$$

We let $\mathcal{P}^+(\Sigma^\infty)$ denote the set of closed sets of finite or infinite traces. This again forms a complete lattice under set inclusion.

Again we show that \mathcal{T} can be defined compositionally. The clauses defining the traces of **skip** and assignment are unchanged.

We extend concatenation to infinite traces in the obvious way: if α is finite we define $\alpha\beta$ as before to be the result of concatenating β on the end of α ; if α is infinite we define $\alpha\beta$ to be α . Then we define $T_1; T_2 = \{\alpha\beta \mid \alpha \in T_1, \beta \in T_2\}^\dagger$ and again we have $\mathcal{T}[C_1; C_2] = \mathcal{T}[C_1]; \mathcal{T}[C_2]$.

We define T^* on closed sets of finite or infinite traces as before but using the extended notion of concatenation. We also define

$$T^\omega = \{\alpha_0 \alpha_1 \dots \alpha_n \dots \mid \forall n \geq 0. \alpha_n \in T\}^\dagger.$$

We then extend the trace set of a loop command to include infinite traces corresponding to divergence:

$$\begin{aligned} \mathcal{T}[\text{while } B \text{ do } C] \\ = (\mathcal{T}[B]; \mathcal{T}[C])^*; \mathcal{T}[\neg B] \cup (\mathcal{T}[B]; \mathcal{T}[C])^\omega. \end{aligned}$$

Similarly we extend the trace set of an await command to include infinite traces corresponding to blocking:

$$\begin{aligned} \mathcal{T}[\text{await } B \text{ then } C] \\ = \{(s, s') \mid (s, s) \in \mathcal{T}[B] \& (s, s') \in \mathcal{T}[C]\}^\dagger \cup (\mathcal{T}[\neg B])^\omega. \end{aligned}$$

For α and β in Σ^∞ let $\alpha \parallel \beta$ be the set of all traces built by fairly interleaving α with β . Perhaps the simplest way to define $\alpha \parallel \beta$ formally, following Park [19], is

$$\begin{aligned} \alpha \parallel \beta &= \{\gamma \mid (\alpha, \beta, \gamma) \in \text{fairmerge}\} \\ \text{fairmerge} &= (L^*RR^*L)^\omega \cup (L \cup R)^* A \\ L &= \{(\sigma, \epsilon, \sigma) \mid \sigma \in \Sigma\} \\ R &= \{(\epsilon, \sigma, \sigma) \mid \sigma \in \Sigma\} \\ A &= \{(\alpha, \epsilon, \alpha) \mid \alpha \in \Sigma^\infty\} \cup \{(\epsilon, \beta, \beta) \mid \beta \in \Sigma^\infty\}, \end{aligned}$$

where we extend concatenation to work on sets and on triples of traces in the obvious way. To be precise, for $X, Y \subseteq \Sigma^\infty$ let $XY = \{\alpha\beta \mid \alpha \in X, \beta \in Y\}$; for triples let $(\alpha_1, \alpha_2, \alpha_3)(\beta_1, \beta_2, \beta_3) = (\alpha_1\beta_1, \alpha_2\beta_2, \alpha_3\beta_3)$. Intuitively, when $(\alpha, \beta, \gamma) \in (L^*RR^*L)^\omega$ each of α and β is infinite and γ is an interleaving of all of α with all of β . When $(\alpha, \beta, \gamma) \in (L \cup R)^* A$ at least one of α and β is finite, and γ is again an interleaving of all of α with all of β ; as soon as all of one trace has been consumed there is no longer any fairness requirement.

Then we define a fair interleaving operator on closed sets of traces by

$$T_1 \parallel T_2 = \bigcup \{\alpha_1 \parallel \alpha_2 \mid \alpha_1 \in T_1 \& \alpha_2 \in T_2\}^\dagger,$$

with the result that $\mathcal{T}[C_1 \parallel C_2] = \mathcal{T}[C_1] \parallel \mathcal{T}[C_2]$. With these definitions in hand, we can define \mathcal{T} denotationally. We give details for the coarse-grained case; the corresponding fine-grained version is obtainable similarly.

DEFINITION 10.1. The fair transition traces semantic function $\mathcal{F}: \mathbf{Com} \rightarrow \mathcal{P}^\dagger(\Sigma^\infty)$ is defined by the following clauses:

$$\begin{aligned} \mathcal{F}[\mathbf{skip}] &= \{(s, s) \mid s \in S\}^\dagger \\ \mathcal{F}[I := E] &= \{(s, [s \mid I = n]) \mid (s, n) \in \mathcal{E}[E]\}^\dagger \\ \mathcal{F}[C_1; C_2] &= \mathcal{F}[C_1]; \mathcal{F}[C_2] \\ \mathcal{F}[\mathbf{if } B \mathbf{ then } C_1 \mathbf{ else } C_2] \\ &= \mathcal{F}[B]; \mathcal{F}[C_1] \cup \mathcal{F}[\neg B]; \mathcal{F}[C_2] \\ \mathcal{F}[\mathbf{while } B \mathbf{ do } C] \\ &= (\mathcal{F}[B]; \mathcal{F}[C])^*; \mathcal{F}[\neg B] \cup (\mathcal{F}[B]; \mathcal{F}[C])^\omega \\ \mathcal{F}[C_1 \parallel C_2] &= \mathcal{F}[C_1] \parallel \mathcal{F}[C_2] \\ \mathcal{F}[\mathbf{await } B \mathbf{ then } C] \\ &= \{(s, s') \in \mathcal{F}[C] \mid (s, s) \in \mathcal{F}[B]\}^\dagger \cup (\mathcal{F}[\neg B])^\omega. \end{aligned}$$

Yet again all operations on trace sets used in this semantics are continuous with respect to set inclusion. The semantics of while-loops again has an equivalent formulation using fixed-points, but now we must use a *greatest* fixed-point. As before, the function $F'(T) = (\mathcal{F}[B] \cdot \mathcal{F}[C] \cdot T \cup \mathcal{F}[\neg B])$ is monotone on the complete lattice of (not necessarily closed) sets of traces. Its greatest fixed point is

$$\begin{aligned} \nu T. F'(T) &= (\mathcal{F}[B] \cdot \mathcal{F}[C])^* \cdot \mathcal{F}[\neg B] \cup \\ &(\mathcal{F}[B] \cdot \mathcal{F}[C])^\omega, \end{aligned}$$

and the closure of this set is $\mathcal{F}[\mathbf{while } B \mathbf{ do } C]$.⁷

To show how the fairmerge definition works, we now return to the example programs discussed above. First, it is easy to show that all interference-free traces of the program

$$x := 0 \parallel \mathbf{while } x = 1 \mathbf{ do } y := y + 1$$

are finite, corresponding to the fact that according to the fair operational semantics this program always terminates. Second, consider the program

$$\mathbf{await } x = 1 \mathbf{ then } y := 1 \parallel \mathbf{while } y = 0 \mathbf{ do } x := 1 - x.$$

⁷ The reader can easily check that $\mathcal{F}[\mathbf{while } B \mathbf{ do } C]$ does *not* coincide with the greatest fixed point of the corresponding operator on closed sets, $\lambda T. (\mathcal{F}[B]; \mathcal{F}[C]; T \cup \mathcal{F}[\neg B])$. For instance, when $B = \mathbf{true}$ and $C = \mathbf{skip}$ this functional is simply $\lambda T. T$ and its greatest fixed point is $\mathcal{P}^\dagger(\Sigma^\infty)$, the set of all traces; but $\mathcal{F}[\mathbf{while } \mathbf{true} \mathbf{ do } \mathbf{skip}] = \{(s, s) \mid s \in S\}^\omega$, consisting only of the infinite stuttering sequences.

Introduce the following abbreviations:

$$\begin{aligned} \sigma_- &= ([x = 1, y = 0], [x = 0, y = 0]) \\ \sigma_+ &= ([x = 0, y = 0], [x = 1, y = 0]) \\ \tau &= ([x = 0, y = 0], [x = 0, y = 0]) \end{aligned}$$

Then $(\sigma_- \sigma_+)^\omega$ is a trace of the while-loop, and τ^ω is a trace of the await-command. And we have $((\sigma_- \sigma_+)^\omega, \tau^\omega, (\sigma_- \tau \sigma_+)^\omega) \in \mathit{fairmerge}$. The trace $(\sigma_- \tau \sigma_+)^\omega$ represents a fair infinite execution of the parallel program in which the await command is continually blocked. The program also has traces corresponding to executions in which the await command is scheduled in a state satisfying its test, so that y gets set to 1.

10.1. Full Abstraction

We now generalize the notion of state trace behavior to include infinite state traces. Let $S^\infty = S^+ \cup S^\omega$ be the set of non-empty finite and infinite sequences of states.

DEFINITION 10.2. The fair state trace behavior function $\mathcal{S}: \mathbf{Com} \rightarrow \mathcal{P}(S^\infty)$ is given by:

$$\begin{aligned} \mathcal{S}[C] &= \{s_0 s_1 \dots s_k \mid \langle C, s_0 \rangle \rightarrow^* \langle C_1, s_1 \rangle \rightarrow^* \\ &\dots \rightarrow^* \langle C_k, s_k \rangle \mathit{term}\} \cup \\ &\{s_0 \dots s_n \dots \mid \langle C, s_0 \rangle \rightarrow^* \langle C_1, s_1 \rangle \rightarrow^* \\ &\dots \langle C_n, s_n \rangle \rightarrow^* \dots \mathit{fair}\}. \end{aligned}$$

By definition, as before, $\mathcal{S}[C]$ is determined by the interference-free subset of $\mathcal{F}[C]$.

PROPOSITION 10.3. *The fair transition trace semantics is fully abstract with respect to fair state traces: for all commands C and C' , $C \sqsubseteq_{\mathcal{F}} C' \Leftrightarrow C \sqsubseteq_{\mathcal{S}} C'$.*

Proof. As before the forward implication is straightforward. For the converse, suppose C has a trace α that is impossible for C' . If α is finite we may argue as before. Otherwise, suppose α is infinite, say

$$\alpha = (s_0, s'_0)(s_1, s'_1) \dots (s_n, s'_n) \dots$$

Let x_1, \dots, x_k be the (finitely many) identifiers occurring free in either C or C' . Without loss of generality we can assume that the domain of each state in α is $\{x_1, \dots, x_k\}$. Let t, c, y_1, \dots, y_k and z_1, \dots, z_k be distinct fresh identifiers. Let $\bar{x} := \bar{x}$ stand for the command

$$x_1 := z_1; \dots; x_k := z_k$$

and $\bar{y} := \bar{0}$ stand for $y_1 := 0; \dots; y_k := 0$. Let $\bar{x} = \bar{y}$ be the boolean expression

$$(x_1 = y_1) \& \dots \& (x_k = y_k).$$

Let $\text{CHOOSE}(\bar{y})$ be the following command:

$$\begin{aligned} & \bar{y} := \bar{0}; t := 0; (\quad t := 1 \\ & \quad \parallel \text{ while } t = 0 \text{ do } y_1 := y_1 + 1 \\ & \quad \parallel \text{ while } t = 0 \text{ do } y_2 := y_2 + 1 \\ & \quad \parallel \dots\dots\dots \\ & \quad \parallel \text{ while } t = 0 \text{ do } y_k := y_k + 1 \\ &) \end{aligned}$$

with a similar form for $\text{CHOOSE}(\bar{z})$. Intuitively, $\text{CHOOSE}(\bar{y})$ "guesses" a state; fairness guarantees termination and that for each state s there is a computation of $\text{CHOOSE}(\bar{y})$ that guesses s , by setting each y_i to $s(x_i)$ for $i = 1, \dots, k$. We will write $\bar{y} = s(\bar{x})$ to indicate the state $[y_1 = s(x_1), \dots, y_k = s(x_k)]$.

Let $P[-]$ be the context

$$\begin{aligned} & [-] \parallel \text{ while true do } (\text{CHOOSE}(\bar{y}); \\ & \quad \text{CHOOSE}(\bar{z}); \\ & \quad c := c + 1; \\ & \quad \text{await } \bar{x} = \bar{y} \text{ then } \bar{x} := \bar{z}). \end{aligned}$$

Then $P[C]$ has an infinite state trace characteristic of α while $P[C']$ does not. Specifically, $P[C]$ has a fair execution in which at stage n the loop guesses the states s'_n and s_{n+1} then waits until C reaches s'_n before setting the state to s_{n+1} . This corresponds to a state trace of the form

$$\begin{aligned} & [\bar{x} = s_0(\bar{x}) \mid y = \bar{0} \mid \bar{z} = \bar{0} \mid c = 0] \\ & [\bar{x} = s_0(\bar{x}) \mid \bar{y} = s'_0(\bar{x}) \mid \bar{z} = s_1(\bar{x}) \mid c = 1] \\ & [\bar{x} = s'_0(\bar{x}) \mid \bar{y} = s'_0(\bar{x}) \mid \bar{z} = s_1(\bar{x}) \mid c = 1] \\ & [\bar{x} = s_1(\bar{x}) \mid \bar{y} = s'_1(\bar{x}) \mid \bar{z} = s_2(\bar{x}) \mid c = 2] \\ & [\bar{x} = s'_1(\bar{x}) \mid \bar{y} = s'_1(\bar{x}) \mid \bar{z} = s_2(\bar{x}) \mid c = 2] \\ & [\bar{x} = s_2(\bar{x}) \mid \bar{y} = s'_2(\bar{x}) \mid \bar{z} = s_3(\bar{x}) \mid c = 3] \\ & \dots\dots\dots, \end{aligned}$$

with alternating steps made by the context and by C . Since C' cannot manipulate the y_i, z_i or c , and does not have the trace α , $P[C']$ has no state trace of this form. ■

The laws given in Proposition 8.1 continue to hold for the fair trace semantics. However, the inequation

$$C_1; (C_2 \parallel C_3) \sqsubseteq (C_1; C_2) \parallel C_3$$

and its corollary $C_1; C_2 \sqsubseteq C_1 \parallel C_2$ may fail when C_1 has infinite traces. Nevertheless, these inequations still hold if C_1 has only finite traces, for instance when C_1 contains no loops or awaits.

Note that the busy-wait treatment of blocking means that this semantics still identifies **await false then skip** with **while true do skip**, since they both denote the set of all infinite stuttering sequences. However, the inequation **while true do skip** $\sqsubseteq C$ now fails.

11. ROBUSTNESS

The full abstraction results given above relied only on certain general properties: monotonicity of the semantic definitions, compositionality of \mathcal{S} , and the fact that the behavior of a program is embedded in its trace set. We can therefore extend these results to deal with any additional program constructs that do not violate these properties.

11.1 Non-deterministic Choice

For instance, we may add a non-deterministic choice construct C_1 or C_2 , with operational semantics given by

$$\begin{aligned} & \langle C_1 \text{ or } C_2, s \rangle \rightarrow \langle C_1, s \rangle \\ & \langle C_1 \text{ or } C_2, s \rangle \rightarrow \langle C_2, s \rangle. \end{aligned}$$

Then $\mathcal{S}[C_1 \text{ or } C_2] = \mathcal{S}[C_1] \cup \mathcal{S}[C_2]$, and all of the previous development goes through with minor modifications.⁸ For the language extended with non-deterministic choice this semantics is again fully abstract, and the laws of programming given earlier continue to hold. In addition, **or** is idempotent, commutative and associative, **or** distributes through sequential and parallel composition, and $C \sqsubseteq C'$ if and only if $(C \text{ or } C') = C'$. The following laws of equivalence hold:

$$\begin{aligned} & C \text{ or } C \equiv C \\ & C_1 \text{ or } C_2 \equiv C_2 \text{ or } C_1 \\ & (C_1 \text{ or } C_2) \text{ or } C_3 \equiv C_1 \text{ or } (C_2 \text{ or } C_3) \\ & (C_1 \text{ or } C_2); C \equiv (C_1; C) \text{ or } (C_2; C) \\ & C; (C_1 \text{ or } C_2) \equiv (C; C_1) \text{ or } (C; C_2) \\ & (C_1 \text{ or } C_2) \parallel C \equiv (C_1 \parallel C) \text{ or } (C_2 \parallel C) \\ & \text{if } B \text{ then } (C \text{ or } C') \text{ else } C_2 \\ & \equiv (\text{if } B \text{ then } C \text{ else } C_2) \text{ or } (\text{if } B \text{ then } C' \text{ else } C_2). \end{aligned}$$

⁸ Of course, the coroutine construct C_1 **co** C_2 from Hennessy-Plotkin cannot be handled by our semantics, since $\mathcal{S}[C_1 \text{ co } C_2]$ cannot be determined from $\mathcal{S}[C_1]$ and $\mathcal{S}[C_2]$.

The coarse-grained semantics satisfies the law

$$\begin{aligned} I_1 := E_1 \parallel I_2 := E_2 \\ \equiv (I_1 := E_1; I_2 := E_2) \text{ or } (I_2 := E_2; I_1 := E_1), \end{aligned}$$

but this fails in the fine-grained case since $x := x + 1 \parallel x := x + 1$ has the trace $([x = 0], [x = 1])$, but this is not a trace of $x := x + 1; x := x + 1$. In the fine-grained semantics we obtain instead the following inequation:

$$\begin{aligned} I_1 := E_1 \parallel I_2 := E_2 \\ \supseteq (I_1 := E_1; I_2 := E_2) \text{ or } (I_2 := E_2; I_1 := E_1). \end{aligned}$$

11.2. Local Variables

We can also expand the programming language to include *local variable declarations* by adding the syntactic clause

$$C ::= \text{new } I = E \text{ in } C'.$$

The intention is that I is a local variable initialized to the value of E and that C' may update and read I . Each occurrence of I in C' is bound by the declaration, so that

$$\text{free}[\text{new } I = E \text{ in } C'] = \text{free}[E] \cup (\text{free}[C'] - \{I\}).$$

Since the scope of the declaration *only* includes C' no other process executing in parallel may read or write to the local variable. An advantage provided by this extension is that it permits a natural form of reasoning about the behavior of programs under “locality” assumptions, such as the assumption that no parallel command will alter the value of certain shared-variables. It is possible to incorporate local variable declarations into both coarse- and fine-grained semantic models. Again we give details for the coarse-grained case.

To define the operational behavior of this construct we first introduce generalized commands of the form $\text{new } [I = n] \text{ in } C$, representing a block with body C using local variable I with current value n . The transition rules are

$$\begin{array}{c} \frac{\langle E, s \rangle \rightarrow^* n}{\langle \text{new } I = E \text{ in } C, s \rangle \rightarrow \langle \text{new } [I = n] \text{ in } C, s \rangle} \\ \frac{\langle C, [s \mid I = n] \rangle \rightarrow \langle C', [s' \mid I = n'] \rangle}{\left(\frac{\langle \text{new } [I = n] \text{ in } C, s \rangle}{\rightarrow \langle \text{new } [I = n'] \text{ in } C', [s' \mid I = s(I)] \rangle} \right)} \\ \frac{\langle C, [s \mid I = n] \rangle \text{ term}}{\langle \text{new } [I = n] \text{ in } C, s \rangle \text{ term}} \end{array}$$

where we extend the update operation on states so that $[s' \mid I = s(I)]$ denotes $s' \setminus I$ when $I \notin \text{dom}(s)$, where $s' \setminus I$ is the restriction of s' to the set $\text{dom}(s') - \{I\}$. Note the careful treatment of I , distinguishing between local and global versions of the identifier; it is easy to see that every computation of $\text{new } I = E \text{ in } C$ from a state s leaves the value of (the global) I unchanged.

To give a denotational description we first introduce some auxiliary definitions. For a trace α with n th step (s_n, s'_n) let $\alpha \setminus I$ be the trace with n th step $(s_n, [s'_n \mid I = s_n(I)])$. By convention, if $I \notin \text{dom}(s_n)$ this collapses to (s_n, s'_n) . Clearly $\alpha \setminus I$ is a trace in which the value of I is never changed “internally.” We also define $\langle I = k \rangle \alpha$ to be the trace beginning with the step $([s_0 \mid I = k], s'_0)$ and whose n th step for $n \geq 1$ is $([s_n \mid I = s'_{n-1}(I)], s'_n)$. This operation constructs a trace in which I is initialized to the value k and is never changed “externally.” Then we define

$$\begin{aligned} \mathcal{T}[\text{new } I = E \text{ in } C'] \\ = \{ \alpha(\beta \setminus I) \mid (\alpha, k) \in \mathcal{T}[E] \ \& \ \langle I = k \rangle \beta \in \mathcal{T}[C'] \}^\dagger. \end{aligned}$$

The definition captures the idea that C' treats I as a local variable, and even though its environment may change the value of (a global variable named) I such a change does not affect the value of the local version of I ; similarly, although C' may change the value of its local I it has no effect on the global version of I .

For example, the command

$$\text{new } x = 1 \text{ in } (y := x; x := x + 1; y := x)$$

has each of the traces

$$\begin{aligned} & ([y = 0], [y = 1])([y = 1], [y = 2]) \\ & ([x = 9, y = 0], [x = 9, y = 1]) \\ & ([x = 9, y = 1], [x = 9, y = 2]) \\ & ([x = 9, y = 0], [x = 9, y = 1]) \\ & ([x = 5, y = 1], [x = 5, y = 2]) \\ & ([y = 0], [y = 1])([y = 8], [y = 2]) \end{aligned}$$

but no traces of the forms

$$\begin{aligned} & \alpha([x = 9, y = 0], [x = 9, y = 9])\beta \\ & \alpha([x = 0, y = 0], [x = 1, y = 0])\beta \end{aligned}$$

for any $\alpha, \beta \in \Sigma^\infty$. It is also easy to see that $\mathcal{T}[\text{new } x = 0 \text{ in } x := x + 1] = \mathcal{T}[\text{skip}]$.

With this expansion the trace semantics is still fully abstract, and validates the following natural and intuitively valid equivalence:

$$\text{new } I = E \text{ in } C \equiv C \quad \text{if } I \notin \text{free}[C].$$

The proof relies on the following property: if I does not occur free in C then every trace α of C satisfies $\alpha \setminus I = \alpha$.

Similarly we can show that when I is not free in C_1 , and C_1 does not change any identifier occurring in E , we get

$$\begin{aligned} \text{new } I = E \text{ in } (C_1; C_2) &\equiv C_1; \text{new } I = E \text{ in } C_2 \\ \text{new } I = E \text{ in } (C_1 \parallel C_2) &\equiv C_1 \parallel \text{new } I = E \text{ in } C_2. \end{aligned}$$

Finally, if we define $[I'/I]C$ to be the command obtained by replacing all free occurrences of I in C by I' , renaming bound variables when necessary to avoid capture, we can show the validity of the law

$$\text{new } I = E \text{ in } C \equiv \text{new } I' = E \text{ in } [I'/I]C,$$

provided I' does not occur free in C .

As an example illustrating the use of local variables, assume that f denotes a boolean-valued total function on the integers and consider the following commands:

$$\begin{aligned} C_0 &= \text{while } found = 0 \text{ do (if } f(i) \\ &\quad \text{then } (found := 1; k := i) \text{ else } i := i + 2) \\ C_1 &= \text{while } found = 0 \text{ do (if } f(j) \\ &\quad \text{then } (found := 1; k := j) \text{ else } j := j + 2). \end{aligned}$$

Let SEARCH be the program

$$\begin{aligned} \text{new } found = 0 \text{ in} \\ (\text{new } i = 0 \text{ in } C_0) \parallel (\text{new } j = 1 \text{ in } C_1). \end{aligned}$$

Intuitively, SEARCH runs two loops in parallel, each looking for a "root" of f , and terminates if a root is found. Each loop uses a local counter variable, and the variables k and $found$ are shared. Assuming that f has a root and that execution is fair, this program should always terminate and the final value of k should be a root of f . This is shown by the trace set:⁹

$$\begin{aligned} \mathcal{T}[\text{SEARCH}] &= \{(s, [s \mid k = n]) \mid \\ &\quad n \in \mathbb{N} \ \& \ (s, \tau\tau) \in \mathcal{B}[f(n)]\}^\dagger. \end{aligned}$$

⁹ We tolerate a slight abuse of notation, using n to denote both an integer and the corresponding numeral.

12. SUMMARY AND CONCLUSIONS

We have introduced transition traces and used them as the basis for a variety of fully abstract semantics for a shared-variable parallel programming language. Our results apply in coarse- and fine-grained versions to yield full abstraction with respect to two forms of program behavior, corresponding to partial and strong correctness. In each case, extra language features may be added without invalidating full abstraction, provided certain general semantic properties are preserved; in particular, the trace semantics of the new features must be definable compositionally and monotonically. This shows the flexibility and generality of our ideas and results.

The idea of using sequences or traces of some kind to model the behavior of concurrent programs is widespread. For instance, several authors have used traces (sequences of communications) to build models of determinate or indeterminate dataflow networks, notably [11, 13, 21]. Indeed, others have also used sequences of pairs of states in imperative settings [2, 7, 10, 19].

In the earlier papers [2, 19] a pair of states was used to represent a *single* atomic action. Park's semantics [19] is obviously closely related to ours, since we adapt Park's definition of fairmerge; but Park's use of single-step traces causes his semantics to be too concrete, distinguishing between **skip** and **skip**; **skip** again. Abrahamson's semantics [2] is too concrete for the same reason, and ignores fairness. The key difference between this early work and ours is that we use a pair of states to represent an arbitrary *finite sequence* of atomic actions, and consequently develop a semantics based on *closed* sets of traces.

In the more recent papers [7, 10] a pair of states represents a single state-changing atomic action together with a finite sequence of idle steps. The semantics presented in [7, 10] models a program as a set of finite traces closed under a slightly weaker form of stuttering and mumbling than ours. In [7, 10] a set of traces T is said to be closed if and only if

$$\begin{aligned} \alpha\beta \in T &\Rightarrow \alpha(s, s)\beta \in T \\ \alpha(s, s)(s, s')\beta \in T &\Rightarrow \alpha(s, s')\beta \in T \\ \alpha(s, s')(s', s')\beta \in T &\Rightarrow \alpha(s, s')\beta \in T, \end{aligned}$$

so that mumbling is only allowed to absorb idle steps. The semantics given in [7] does not explicitly account for recursion (and therefore ignore issues of fairness). Their semantics achieves full abstraction with respect to a different notion of behavior, in which the observer of a program is assumed to be omniscient, able to observe *every* state-change occurring during a computation. For example the equivalence

$$(x := 1; x := x + 1) \text{ or } x := 2 \equiv x := 1; x := x + 1$$

fails in the semantics of [7], whereas it holds in our model. The treatment of local variables in [7] is essentially the same as ours.

Abadi and Plotkin [1] use a trace model (prefix-closed sets of finite sequences of pairs of states, also closed under stuttering and mumbling) for reasoning about safety properties of reactive systems and the study of composition rules.

Jonsson and Back [12] have independently constructed a trace model closely related to ours, concentrating on total correctness and dealing with unfair as well as fair execution; they consider a simple imperative language of guarded commands (with branches executed as atomic actions), extended with an interleaving parallel composition. Their model for partial correctness is essentially the same as our finite trace model, involving closure under stuttering and mumbling; for total correctness they introduce closure under "fair unstopability." In an earlier version of this paper [5], we mentioned that in a total correctness semantics trace sets would need to be closed under "chattering"; as Jonsson and Back show, this is not sufficient to achieve full abstraction. However, since the notion of termination considered in [12] is rather specialized, it is unclear if their ideas can be used to yield a total correctness semantics for our shared-variable language.

Program constructs or operational assumptions (such as fairness) that give rise to unbounded nondeterminism do not cause semantic problems in our framework. For instance, it is almost trivial to add a random assignment command $I := ?$ to the syntax, with the transition rule

$$\langle I := ?, s \rangle \rightarrow \langle \text{skip}, [s \mid I = n] \rangle \quad (n \in N)$$

and the following denotational semantics:

$$\mathcal{F}[I := ?] = \{(s, [s \mid I = n]) \mid s \in S \ \& \ n \in N\}^\dagger$$

This would again yield a fully abstract semantics.

In contrast unbounded non-determinism causes severe problems in traditional powerdomain semantics [9, 4]. Apt and Plotkin [4] proved that for a sequential while-loop language with random assignment there is no denotational continuous least fixed point semantics that is fully abstract with respect to a notion of behavior based on strong correctness, equivalently with respect to finite and infinite state traces. Our fair trace model provides a denotational continuous semantics for a parallel version of this language, is fully abstract with respect to the same notion of behavior, but is not a least fixed point semantics. The fair trace model can instead be characterized as the closure of a greatest fixed point semantics. We also showed that the corresponding least fixed point semantics (the finite traces model) is fully abstract with respect to finite state trace behavior, equivalently with respect to partial correctness behavior. For the

sequential language discussed by Apt and Plotkin there is no need to use traces to achieve full abstraction, since the relevant behavior functions can already be defined compositionally. When our semantic definitions are simplified and adapted to the sequential setting they yield two fully abstract semantics for the Apt-Plotkin language, with respect to partial correctness and strong correctness, respectively.

ACKNOWLEDGMENTS

Several people have made helpful comments and suggestions that led to improvements in the presentation of this paper, including Albert Meyer, Prakash Panangaden, Vaughan Pratt, John Reynolds, and the anonymous referees. Susan Older suggested using non-deterministic "guessing" to simplify the full abstraction proof for the fair trace model. This research was sponsored in part by the Avionics Laboratory, Wright Research and Development Center, Aeronautical Systems Division (AFSC), U.S. Air Force, Wright-Patterson AFB, OH 45433-6543 under Contract F33615-90-C-1465, Arpa Order 7597. Support also came from the National Science Foundation under Grant CCR-9006064. The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government.

Received January 1, 1994; final manuscript received February 1, 1996

REFERENCES

1. Abadi, M., and Plotkin, G. D. (1993), A logical view of composition, *Theoret. Comput. Sci.* **114**(1), 3-30.
2. Abrahamson, K. (1979), Modal logic of concurrent nondeterministic programs, in "Semantics of Concurrent Computation" (G. Kahn, Ed.), Lecture Notes in Computer Science, Vol. 70, pp. 21-33, Springer-Verlag, Berlin/New York.
3. Andrews, G. R., and Schneider, F. B. (1983), Concepts and notations for concurrent programming, *ACM Comput. Surveys* **15**(1), 3-43.
4. Apt, K. R., and Plotkin, G. D. (1986), Countable nondeterminism and random assignment, *J. Assoc. Comput. Mach.* **33** (4), 724-767.
5. Brookes, S. (1993), Full abstraction for a shared variable parallel language, in "Proc. 8th Annual IEEE Symposium on Logic in Computer Science," IEEE Comput. Soc. Press, Los Alamitos, CA.
6. de Bakker, J. W. (1971), Axiom systems for simple assignment statements, in "Symposium on Semantics of Algorithmic Languages" (E. Engeler, Ed.), Lecture Notes in Mathematics, Vol. 188, pp. 1-22, Springer-Verlag, Berlin/New York.
7. de Boer, F., Kok, J., Palamidessi, C., and Rutten, J. (1991), The failure of failures in a paradigm for asynchronous communication, in "Concur '91" (J. Baeten and J. Groote, Eds.), Lecture Notes in Computer Science, Vol. 527, pp. 111-126, Springer-Verlag, Berlin/New York.
8. Francez, N. (1986), "Fairness," Springer-Verlag, Berlin/New York.
9. Hennessy, M., and Plotkin, G. D. (1979), Full abstraction for a simple parallel programming language, in "Mathematical Foundations of Computer Science," Lecture Notes in Computer Science, Vol. 74, pp. 108-120, Springer-Verlag, Berlin/New York.
10. Horita, E., de Bakker, J., and Rutten, J. (1990), Fully Abstract Denotational Models for Nonuniform Concurrent Languages," Technical Report CS-R9027, Centre for Mathematics and Computer Science, Amsterdam.
11. Jonsson, B. (1989), A fully abstract trace semantics for dataflow networks, in "Sixteenth Annual ACM Symposium on Principles of Programming Languages," pp. 155-165, ACM Press, New York.

12. Jonsson, B., and Back, R. J. R. (1993), Fully abstract semantic orderings for shared-variable concurrent programs, draft paper.
13. Keller, R. M., and Panangaden, P. (1986), Semantics of digital networks containing indeterminate operators, *Distrib. Comput.* 1(4), 235–245.
14. Lamport, L. (1977), Proving the correctness of multiprocess programs, *IEEE Trans. Software Eng.* 3(2), 125–143.
15. Lamport, L. (1983), What good is temporal logic? in "Information Processing 83: Proceedings of the IFIP 9th World Congress" (R. E. A. Mason, Ed.), pp. 657–668, North-Holland, Amsterdam.
16. Lehmann, D., Pnueli, A., and Stavi, J. (1981), Impartiality, justice, and fairness: the ethics of concurrent termination, in "Proceedings of the Eighth International Conference on Automata, Languages and Programming," Lecture Notes in Computer Science, Vol. 115, pp. 264–277, Springer-Verlag, Berlin/New York.
17. Milner, R. (1977), Fully abstract models of typed lambda-calculi, *Theoret. Comput. Sci.* 4, 1–22.
18. Owicki, S. S., and Gries, D. (1976), An axiomatic proof technique for parallel programs, *Acta Inform.* 6, 319–340.
19. Park, D. (1979), On the semantics of fair parallelism, in "Abstract Software Specifications" (D. Bjørner, Ed.), Lecture Notes in Computer Science, Vol. 86, pp. 504–526, Springer-Verlag, Berlin/New York.
20. Plotkin, G. D. (1977), LCF considered as a programming language, *Theoret. Comput. Sci.* 5(3), 223–255.
21. Russell, J. R. (1989), Full abstraction for nondeterministic dataflow networks, in "Proceedings of the 30th Annual Symposium on Foundations of Computer Science," pp. 170–177, IEEE Press, New York.
22. Stoughton, A. (1988), "Fully Abstract Models of Programming Languages," Research Notes in Theoretical Computer Science, Pitman, London.

