
INDEPENDENCE AND CONCURRENT SEPARATION LOGIC*

JONATHAN HAYMAN AND GLYNN WINSKEL

Computer Laboratory, University of Cambridge, William Gates Building, 15 JJ Thomson Avenue,
Cambridge CB 0FD, United Kingdom

e-mail address: {jonathan.hayman, glynn.winskel}@cl.cam.ac.uk

ABSTRACT. A compositional Petri net-based semantics is given to a simple language allowing pointer manipulation and parallelism. The model is then applied to give a notion of validity to the judgements made by concurrent separation logic that emphasizes the process-environment duality inherent in such rely-guarantee reasoning. Soundness of the rules of concurrent separation logic with respect to this definition of validity is shown. The independence information retained by the Petri net model is then exploited to characterize the independence of parallel processes enforced by the logic. This is shown to permit a refinement operation capable of changing the granularity of atomic actions.

1. INTRODUCTION

The foundational work of Hoare on parallel programming [Hoa72] identified the fact that attributing an interleaved semantics to parallel languages is problematic. Three areas of difficulty were isolated, quoted directly:

- *That of defining a ‘unit of action’.*
- *That of implementing the interleaving on genuinely parallel hardware.*
- *That of designing programs to control the fantastic number of combinations involved in arbitrary interleaving.*

The significance of these problems increases with developments in hardware, such as multiple-core processors, that allow primitive machine actions to occur at the same time.

As Hoare went on to explain, a feature of concurrent systems in the physical world is that they are often *spatially separated*, operating on completely different resources and not interacting. When this is so, the systems are *independent* of each other, and therefore it is unnecessary to consider how they interact. This perspective can be extended by regarding computer processes as spatially separated if they operate on different memory locations. The problems above are resolved if the occurrence of non-independent parallel actions is prohibited except in rare cases where atomicity may be assumed, as might be enforced using the constructs proposed in [Dij68, Bri72].

1998 ACM Subject Classification: F.3.2, F.3.1, D.3.1, F.1.2.

Key words and phrases: separation logic, Petri nets, independence models, refinement, granularity.

* Extended version of [HW06].

Independence models for concurrency allow semantics to be given to parallel languages in a way that can tackle the problems associated with an interleaved semantics. The common core of independence models is that they record when actions are independent, and that independent actions can be run in either order or even concurrently with no consequence on their effect. This mitigates the increase in the state space since unnecessary interleavings of independent actions need not be considered (see *e.g.* [CGMP99] for applications to model checking). Independence models also permit easier notions of refinement which allow the assumed atomicity of actions to be changed.

It is surprising that, to our knowledge, there has been no *comprehensive* study of the semantics of programming languages inside an independence model. The first component of our work gives such a semantics in terms of a well-known independence model, namely Petri nets. Our model isolates the specification of the control flow of programs from their effect on the shared state. It indicates what appears to be a general method (an alternative to Plotkin’s structural operational semantics) for giving a *structural Petri net semantics* to a variety of languages — see the Conclusion, Section 7.

The language that we consider is motivated by the emergence of *concurrent separation logic* [O’H07], the rules of which form a partial correctness judgement about the execution of pointer-manipulating concurrent programs. Reasoning about such programs has traditionally proved difficult due to the problem of *variable aliasing*. For instance, Owicki and Gries’ system for proving properties of parallel programs that do not manipulate pointers [OG76] essentially requires that the programs operate on disjoint collections of variables, thereby allowing judgements to be composed. In the presence of pointers, the same syntactic condition cannot be imposed to yield a sound logic since distinct variables may point to the same memory location, thereby allowing arbitrary interaction between the processes. To give a specific example, Owicki and Gries’ system would allow a judgement of the form

$$\{x \mapsto 0 \wedge y \mapsto 0\} x := 1 \parallel y := 2 \{x \mapsto 1 \wedge y \mapsto 2\},$$

indicating that the result of assigning 1 to the program variable x concurrently with assigning 2 to y from a state where x and y both initially hold value 0 is a state where x holds value 1 and y holds value 2. The judgement is sound because the variables x and y are distinct. If pointers are introduced to the language, however, it is not sound to conclude that

$$\{[x] \mapsto 0 \wedge [y] \mapsto 0\} [x] := 1 \parallel [y] := 2 \{[x] \mapsto 1 \wedge [y] \mapsto 2\},$$

which would indicate that assigning 1 to the location pointed to by x and 2 to the location pointed to by y yields a state in which x points to a location holding 1 and y points to a location holding 2, since x and y may both point to the same location.

At the core of separation logic [Rey00, IO01], initially presented for non-concurrent programs, is the *separating conjunction*, $\varphi * \psi$, which asserts that the state in which processes execute may be split into two parts, one part satisfying φ and the other ψ . The separating conjunction was used by O’Hearn to adapt Owicki and Gries’ system to provide a rule for parallel composition suitable for pointer-manipulating programs [O’H07].

As we shall see, the rule for parallel composition is informally understood by splitting the initial state into two parts, one *owned* by the first process and the other by the second. Ownership can be seen as a dynamic constraint on the interference to be assumed: parallel processes always own disjoint sets of locations and only ever act on locations that they own. As processes evolve, ownership of locations may be transferred using a system of *invariants* (an example is presented in Section 4). A consequence of this notion of ownership is that

the rules discriminate between the parallel composition of processes and their interleaved expansion. For example, the logic does *not* allow the judgement

$$\{\ell \mapsto 0\} [\ell] := 1 \parallel [\ell] := 1 \{\ell \mapsto 1\},$$

which informally means that the effect of two processes acting in parallel which both assign the value 1 to the memory location ℓ from a state in which ℓ holds 0 is to yield a state in which ℓ holds 1. However, if we adopt the usual rule for the nondeterministic sum of processes, the corresponding judgement *is* derivable for their interleaved expansion,

$$([\ell] := 1; [\ell] := 1) + ([\ell] := 1; [\ell] := 1).$$

One would hope that the distinction that the logic makes between concurrent processes and their interleaved expansion is captured by the semantics; the Petri net model that we give does so directly.

The rules of concurrent separation logic contain a good deal of subtlety, and so lacked a completely formal account until the pioneering proof of their soundness due to Brookes [Bro07]. The proof that Brookes gives is based on a form of interleaved trace semantics. The presence of pointers within the model alongside the possibility that ownership of locations is transferred means, however, that the way in which processes are separated is absolutely non-trivial, which motivates strongly the study of the language within an independence model. We therefore give a proof of soundness using our net model and then characterize entirely semantically the independence of concurrent processes in Theorem 5.4.

It should be emphasized that the model that we present is different from Brookes' since it provides an *explicit* account of the intuitions behind ownership presented by O'Hearn. It involves taking the original semantics of the process and embellishing it to capture the semantics of the logic. The proof technique that we employ defines validity of assertions in a way that captures the rely-guarantee reasoning [Jon83] emanating from ownership in separation logic directly, and in a way that might be applied in other situations.

In [Rey04], Reynolds argues that the separation of parallel processes arising from the logic allows store actions that were assumed to be atomic, in fact, to be implemented as composite actions (seen as a change in their *granularity*) with no effect on the validity of the judgement. Independence models are suited to modeling situations where actions are not atomic, a perspective advocated by Lamport and Pratt [Pra86, Lam86]. We introduce a novel form of refinement, inspired by that of [vGG89], and show how this may be applied to address the issue of granularity using our characterization of the independence of processes arising from the logic.

2. TERMS AND STATES

Concurrent separation logic is a logic for programs that operate on a *heap*. A heap is a structure recording the values held by memory locations that allows the existence of pointers as well as providing primitives for the allocation and deallocation of memory locations. A heap can be seen as a finite partial function from a set of locations Loc to a set of values Val :

$$\text{Heap} \stackrel{\text{def}}{=} \text{Loc} \rightarrow_{\text{fin}} \text{Val}$$

We will use ℓ to range over elements of Loc and v to range over elements of Val . As stated, a heap location can point to another location, so we require that $\text{Loc} \subseteq \text{Val}$. We shall say that a location is *current* (or *allocated*) in a heap if the heap is defined at that location. The

procedure of making a non-current location current is *allocation*, and the reverse procedure is called *deallocation*. If h is a heap and $h(\ell) = \ell'$, there is no implicit assumption that $h(\ell')$ is defined. Consequently, heaps may contain *dangling* pointers.

In addition to operating on a heap, the programs that we shall consider shall make use of *critical regions* [Dij68] protected by *resources*. The mutual exclusion property that they provide is that no two parallel processes may be inside critical regions protected by the same resource. We will write Res for the set of resources and use r to range over its elements. Critical regions are straightforwardly implemented by recording, for each resource, whether the resource is available or unavailable. A process may enter a critical region protected by r only if r is available; otherwise it is blocked and may not resume execution until the resource becomes available. The process makes r unavailable upon entering the critical region and makes r available again when it leaves the critical region. The language also has a primitive, `resource w do t od`, which says that the variable w represents a resource local to t .

The syntax of the language that we will consider is presented in Figure 1. The symbol α is used to range over *heap actions*, which are actions on the heap that might change the values held at locations but do not affect the domain of definition of the heap. That is, they neither allocate nor deallocate locations. We reserve the symbol b for *boolean guards*, which are heap actions that may proceed without changing the heap if the boolean b holds.

Provision for allocation within our language is made via the `alloc(ℓ)` primitive for $\ell \in \text{Loc}$, which makes a location current and sets ℓ to point at this location. For symmetry, `dealloc(ℓ)` makes the location pointed to by ℓ non-current if ℓ points to a current location. Writing a heap as the set of values that it holds for each allocated location, the effect of the command `alloc(ℓ)` on the heap $\{\ell \mapsto 0\}$ might be to form a heap $\{\ell \mapsto \ell', \ell' \mapsto 1\}$ if the location ℓ' is chosen to be allocated and is assigned initial value 1. The effect of the command `dealloc(ℓ)` on the heap $\{\ell \mapsto \ell', \ell' \mapsto 1\}$ would be to form the heap $\{\ell \mapsto \ell'\}$.

The *guarded sum* $\alpha.t + \alpha'.t'$ is a process that executes as t if α takes place or as t' if α' takes place. We refer the reader to Section ? for a brief justification for disallowing non-guarded sums.

As mentioned earlier, critical regions are provided to control concurrency: the sub-process t inside `with r do t od` can only run when no other process is inside a critical region protected by r . The term `resource w do t od` has the resource variable w bound within t , asserting that a resource is to be chosen that is *local* to t and used for w . Consequently, in the process

$$(\text{resource } w \text{ do with } w \text{ do } t_1 \text{ od od}) \parallel (\text{resource } w \text{ do with } w \text{ do } t_2 \text{ od od})$$

the sub-processes t_1 and t_2 may run concurrently since they must be protected by different resources, one local to the process on the left and the other local to the process on the right. To model this, we shall say that the construct `resource w do t od` *binds* the variable w within t , and the variable w is *free* in `with w do t od`. We write $\text{fv}(t)$ for the free variables in t and say that a term *closed* if it contains no free resource variables; we shall restrict attention to such terms. We write $[r/w]t$ for the term obtained by substituting r for free occurrences of the variable w within t . As standard, we will identify terms ‘up to’ the standard alpha-equivalence \equiv induced by renaming bound occurrences of variables. The notation $\text{res}(t)$ is adopted to represent the resources occurring in t .

The semantics of the term `resource w do t od` will involve first picking a ‘fresh’ resource r and then running $[r/w]t$. It will therefore be necessary to record during the execution of

Terms:

$t ::=$	α	heap action
	$\text{alloc}(\ell)$	heap allocation
	$\text{dealloc}(\ell)$	heap disposal
	$t_1; t_2$	sequential composition
	$t_1 \parallel t_2$	parallel composition
	$\alpha_1.t_1 + \alpha_2.t_2$	guarded sum
	$\text{while } b \text{ do } t \text{ od}$	iteration
	$\text{resource } w \text{ do } t \text{ od}$	resource declaration
	$\text{with } r \text{ do } t \text{ od}$	critical region
	$\text{with } w \text{ do } t \text{ od}$	critical region (local).

Free variables and resources:

$\text{fv}(\alpha) = \emptyset$	$\text{res}(\alpha) = \emptyset$
$\text{fv}(\text{alloc}(\ell)) = \emptyset$	$\text{res}(\text{alloc}(\ell)) = \emptyset$
$\text{fv}(\text{dealloc}(\ell)) = \emptyset$	$\text{res}(\text{dealloc}(\ell)) = \emptyset$
$\text{fv}(t_1; t_2) = \text{fv}(t_1) \cup \text{fv}(t_2)$	$\text{res}(t_1; t_2) = \text{res}(t_1) \cup \text{res}(t_2)$
$\text{fv}(t_1 \parallel t_2) = \text{fv}(t_1) \cup \text{fv}(t_2)$	$\text{res}(t_1 \parallel t_2) = \text{res}(t_1) \cup \text{res}(t_2)$
$\text{fv}(\alpha_1.t_1 + \alpha_2.t_2) = \text{fv}(t_1) \cup \text{fv}(t_2)$	$\text{res}(\alpha_1.t_1; \alpha_2.t_2) = \text{res}(t_1) \cup \text{res}(t_2)$
$\text{fv}(\text{while } b \text{ do } t \text{ od}) = \text{fv}(t)$	$\text{res}(\text{while } b \text{ do } t \text{ od}) = \text{res}(t)$
$\text{fv}(\text{resource } w \text{ do } t \text{ od}) = \text{fv}(t) \setminus \{w\}$	$\text{res}(\text{resource } w \text{ do } t \text{ od}) = \text{res}(t)$
$\text{fv}(\text{with } r \text{ do } t \text{ od}) = \text{fv}(t)$	$\text{res}(\text{with } r \text{ do } t \text{ od}) = \text{res}(t) \cup \{r\}$
$\text{fv}(\text{with } w \text{ do } t \text{ od}) = \text{fv}(t) \cup \{w\}$	$\text{res}(\text{with } w \text{ do } t \text{ od}) = \text{res}(t)$

Substitution:

$[r/w] \alpha$	$= \alpha$	
$[r/w] \text{alloc}(\ell)$	$= \text{alloc}(\ell)$	
$[r/w] \text{dealloc}(\ell)$	$= \text{dealloc}(\ell)$	
$[r/w] t_1; t_2$	$= ([r/w] t_1); ([r/w] t_2)$	
$[r/w] t_1 \parallel t_2$	$= ([r/w] t_1) \parallel ([r/w] t_2)$	
$[r/w] \alpha_1.t_1 + \alpha_2.t_2$	$= \alpha_1.([r/w] t_1) + \alpha_2.([r/w] t_2)$	
$[r/w] \text{while } b \text{ do } t \text{ od}$	$= \text{while } b \text{ do } [r/w] t \text{ od}$	
$[r/w] \text{resource } w' \text{ do } t \text{ od}$	$= \text{resource } w' \text{ do } [r/w] t \text{ od}$	if $w \neq w'$
$[r/w] \text{resource } w \text{ do } t \text{ od}$	$= \text{resource } w \text{ do } t \text{ od}$	
$[r/w] \text{with } r' \text{ do } t \text{ od}$	$= \text{with } r' \text{ do } [r/w] t \text{ od}$	
$[r/w] \text{with } w' \text{ do } t \text{ od}$	$= \begin{cases} \text{with } r \text{ do } [r/w] t \text{ od} & \text{if } w = w' \\ \text{with } w' \text{ do } [r/w] t \text{ od} & \text{otherwise} \end{cases}$	

Figure 1: Syntax of terms

processes which resources are *current* (*i.e.* not fresh) as well as which current resources are *available* (*i.e.* not held by any process).

The way in which we shall formally model the state in which processes execute is motivated by the way in which we shall give the net semantics to closed terms. We begin

by defining the following sets:

$$\begin{aligned} \mathbf{D} &\stackrel{\text{def}}{=} \text{Loc} \times \text{Val} \\ \mathbf{L} &\stackrel{\text{def}}{=} \{\text{curr}(\ell) \mid \ell \in \text{Loc}\} \\ \mathbf{R} &\stackrel{\text{def}}{=} \text{Res} \\ \mathbf{N} &\stackrel{\text{def}}{=} \{\text{curr}(r) \mid r \in \text{Res}\}. \end{aligned}$$

A state σ is defined to be a tuple

$$(D, L, R, N)$$

where $D \subseteq \mathbf{D}$ represents the values held by locations in the heap; $L \subseteq \mathbf{L}$ represents the set of current, or allocated, locations of the heap; $R \subseteq \mathbf{R}$ represents the set of available resources; and $N \subseteq \mathbf{N}$ represents the set of current resources. The sets \mathbf{D} , \mathbf{L} , \mathbf{R} and \mathbf{N} are disjoint, so no ambiguity arises from writing, for example, $(\ell, v) \in \sigma$.

The interpretation of a state for the heap is that $(\ell, v) \in D$ if ℓ holds value v and that $\text{curr}(\ell) \in L$ if ℓ is current. For resources, $r \in R$ if the resource r is available and $\text{curr}(r) \in N$ if r is current. It is clear that only certain such tuples of subsets are sensible. In particular, the heap must be defined precisely on the set of current locations, and only current resources may be available.

Definition 2.1 (Consistent state). The state (D, L, R, N) is *consistent* if we have:

- the sets D , L , R and N are all finite,
- D is a partial function: for all ℓ, v and v' , if $(\ell, v) \in D$ and $(\ell, v') \in D$ then $v = v'$,
- L represents the domain of D : $L = \{\text{curr}(\ell) \mid \exists v : (\ell, v) \in D\}$, and
- all available resources are current: $R \subseteq \{r \mid \text{curr}(r) \in N\}$.

It is clear to see that the L component of any given consistent state may be inferred from the D component. It will, however, be useful to retain this information separately for when the net semantics is given. We shall call $D \subseteq \mathbf{D}$ a *heap* when it is a finite partial function from locations to values, and shall write $\ell \mapsto v$ for its elements rather than (ℓ, v) . We shall frequently make use of the following definition of the domain of a heap D :

$$\text{dom}(D) \stackrel{\text{def}}{=} \{\ell \mid \exists v. (\ell \mapsto v) \in D\}.$$

3. PROCESS MODELS

The definition of state that we have adopted permits a net semantics to be defined. Before doing so, we shall define how heap actions are to be interpreted and then give a transition semantics to closed terms.

3.1. Actions. The earlier definition of state allows a very general form of heap action to be defined that forms a basis for both the transition and net semantics. We assume that we are given the semantics of primitive actions α as $\mathcal{A}[\alpha]$ comprising a set of heap pairs:

$$\mathcal{A}[\alpha] \subseteq \text{Heap} \times \text{Heap}.$$

We require that whenever $(D_1, D_2) \in \mathcal{A}[\alpha]$, it is the case that D_1 and D_2 are (the graphs of) partial functions with the same domain.

The interpretation is that α can proceed in heap D if there are $(D_1, D_2) \in \mathcal{A}[\alpha]$ such that D has the same value as D_1 wherever D_1 is defined. The resulting heap is formed by

updating D to have the same value as D_2 wherever it is defined. It is significant that this definition allows us to infer precisely the set of locations upon which an action depends. The requirement on the domains of D_1 and D_2 ensures that actions preserve consistent markings (Lemma 3.25).

Example 3.1 (Assignment). For any two locations ℓ and ℓ' , let $[\ell] := [\ell']$ represent the action that copies the value held at location ℓ' to location ℓ . Its semantics is as follows:

$$\mathcal{A}[[\ell] := [\ell']] \stackrel{\text{def}}{=} \{(\{\ell \mapsto v, \ell' \mapsto v'\}, \{\ell \mapsto v', \ell' \mapsto v'\}) \mid v, v' \in \text{Val}\}$$

Following the informal account above of the semantics of actions, because in the semantics we have

$$(\{\ell_0 \mapsto 0, \ell_1 \mapsto 1\}, \{\ell_0 \mapsto 1, \ell_1 \mapsto 1\}) \in \mathcal{A}[[\ell_0] := [\ell_1]],$$

the state $\{\ell_0 \mapsto 0, \ell_1 \mapsto 1, \ell_2 \mapsto 2\}$ is updated by $[\ell_0] := [\ell_1]$ to $\{\ell_0 \mapsto 1, \ell_1 \mapsto 1, \ell_2 \mapsto 2\}$. \square

Example 3.2 (Booleans). Boolean guards b are actions that wait until the boolean expression holds and may then take place; they do not update the state. A selection of literals may be defined. For example:

$$\begin{aligned} \mathcal{A}[[\ell] = v] &\stackrel{\text{def}}{=} \{(\{\ell \mapsto v\}, \{\ell \mapsto v\})\} \\ \mathcal{A}[[\ell] = [\ell']] &\stackrel{\text{def}}{=} \{(\{\ell \mapsto v, \ell' \mapsto v\}, \{\ell \mapsto v, \ell' \mapsto v\}) \mid v \in \text{Val}\} \end{aligned}$$

The first gives the semantics of an action that proceeds only if ℓ holds value v and the second gives the semantics of an action that proceeds only if the locations ℓ and ℓ' hold the same value.

Since boolean actions shall not modify the heap, they shall possess the property that:

$$\text{if } (D_1, D_2) \in \mathcal{A}[[b]] \text{ then } D_1 = D_2.$$

This is preserved by the operations defined below. For heaps D and D' , we use $D \uparrow D'$ to mean that D and D' are compatible as partial functions and $D \not\uparrow D'$ otherwise, *i.e.* if they disagree on the values assigned to a common location.

$$\begin{aligned} \mathcal{A}[[\text{true}]] &\stackrel{\text{def}}{=} \{(\emptyset, \emptyset)\} \\ \mathcal{A}[[\text{false}]] &\stackrel{\text{def}}{=} \emptyset \\ \mathcal{A}[[b \wedge b']] &\stackrel{\text{def}}{=} \{(\{D \cup D'\}, \{D \cup D'\}) \mid D \uparrow D' \text{ and } (D, D) \in \mathcal{A}[[b]] \text{ and } (D', D') \in \mathcal{A}[[b']]\} \\ \mathcal{A}[[b \vee b']] &\stackrel{\text{def}}{=} \mathcal{A}[[b]] \cup \mathcal{A}[[b']] \\ \mathcal{A}[[\neg b]] &\stackrel{\text{def}}{=} \{(D, D) \mid D \text{ is a } \subseteq\text{-minimal heap s.t. } \forall D'. (D', D') \in \mathcal{A}[[b]] : D \not\uparrow D'\} \end{aligned}$$

By insisting on minimality in the clause for $\neg b$, we form an action that is defined at as few locations as possible to refute all grounds for b . \square

3.2. Transition semantics. As an aid to understanding the net model, and in particular to give a model with respect to which we can prove its correspondence, a transition semantics for *closed* terms (terms such that $\text{fv}(t) = \emptyset$) is given in Figure 2. A formal relationship between the two semantics is presented in Theorem 3.27. The transition semantics is given by means of labelled transition relations of the forms $\langle t, \sigma \rangle \xrightarrow{\lambda} \langle t', \sigma' \rangle$ and $\langle t, \sigma \rangle \xrightarrow{\lambda} \sigma'$. As usual, the first form of transition indicates that t performs an action labelled λ in state σ

to yield a resumption t' and a state σ' . The second indicates that t in state σ performs an action labelled λ to terminate and yields a state σ' . Labels follow the grammar

λ	$::=$	$\mathbf{act}(D_1, D_2)$	heap action
		$\mathbf{alloc}(\ell, v, \ell', v')$	heap allocation
		$\mathbf{dealloc}(\ell, \ell', v)$	heap disposal
		$\mathbf{decl}(r)$	resource declaration
		$\mathbf{end}(r)$	end of resource scope
		$\mathbf{acq}(r)$	resource acquisition (critical region entry)
		$\mathbf{rel}(r)$	resource release (critical region exit).

In the transition semantics, we write $\sigma \oplus \sigma'$ for the union of the components of two states where they are disjoint and impose the implicit side-condition that this is defined wherever it is used. For example, this implicit side-condition means, in the rule (ALLOC), that for $\mathbf{alloc}(\ell, v, \ell', v')$ to occur we must have $\mathbf{curr}(\ell') \notin \sigma$, and hence ℓ' was initially non-current. Similarly, the rule (RES) can only be applied to derive a transition labelled $\mathbf{decl}(r)$ if the resource r was not initially current.

The syntax of terms is extended temporarily to include $\mathbf{rel} r$ and $\mathbf{end} r$ which are special terms used in the rules (REL) and (END). These, respectively, are attached to the ends of terms protected by critical regions and the ends of terms in which a resource was declared.

For conciseness, we do not give an error semantics to situations in which non-current locations or resources are used; instead, the process will become stuck. We show in Section 4.3 that such situations are excluded by the logic.

3.3. Petri nets. Petri nets, introduced by Petri in his 1962 thesis [Pet62], are a well-known model for concurrent computation. It is beyond the scope of the current article to provide a full account of the many variants of Petri net and their associated theories; we instead refer the reader to [BRR87] for a good account. Roughly, a Petri net can be thought of as a transition system where, instead of a transition occurring from a single global state, an occurrence of an event is imagined to affect only the conditions in its neighbourhood. Petri nets allow a derived notion of *independence* of events; two events are independent if their neighbourhoods of conditions do not intersect.

We base our semantics on the following well-known variant of Petri net (*cf.* the ‘basic’ nets of [CW01] and [WN95]):

Definition 3.3 (Petri net). A Petri net is a five-tuple,

$$(B, E, \bullet(-), (-)^\bullet, M_0).$$

The set B comprises the *conditions* of the net, the set E consists of the *events* of the net, and M_0 is the subset of B of *marked conditions* (the *initial marking*). The maps

$$\bullet(-), (-)^\bullet : E \rightarrow \mathcal{P}ow(B)$$

are the *precondition* and *postcondition* maps, respectively.

Petri nets have an appealing graphical representation, with:

- circles to represent conditions,
- bold lines to represent events,
- arrows from conditions to events to represent the precondition map,

(ACT) :	$\frac{(D_1, D_2) \in \mathcal{A}[\alpha] \quad D_1 \subseteq D \quad D' = (D \setminus D_1) \cup D_2}{\langle \alpha, (D, L, R, N) \rangle \xrightarrow{\text{act}(D_1, D_2)} \langle D', L, R, N \rangle}$		
(ALLOC) :	$\langle \text{alloc}(\ell), \sigma \oplus \{\ell \mapsto v\} \rangle \xrightarrow{\text{alloc}(\ell, v, \ell', v')} \sigma \oplus \{\ell \mapsto \ell', \ell' \mapsto v', \text{curr}(\ell')\}$		
(DEALLOC) :	$\langle \text{dealloc}(\ell), \sigma \oplus \{\ell \mapsto \ell', \ell' \mapsto v', \text{curr}(\ell')\} \rangle \xrightarrow{\text{dealloc}(\ell, \ell', v')} \sigma \oplus \{\ell \mapsto \ell'\}$		
(SEQ) :	$\frac{\langle t_1, \sigma \rangle \xrightarrow{\lambda} \langle t'_1, \sigma' \rangle}{\langle t_1; t_2, \sigma \rangle \xrightarrow{\lambda} \langle t'_1; t_2, \sigma' \rangle}$	(SEQ') :	$\frac{\langle t_1, \sigma \rangle \xrightarrow{\lambda} \sigma'}{\langle t_1; t_2, \sigma \rangle \xrightarrow{\lambda} \langle t_2, \sigma' \rangle}$
(PAR-1) :	$\frac{\langle t_1, \sigma \rangle \xrightarrow{\lambda} \langle t'_1, \sigma' \rangle}{\langle t_1 \parallel t_2, \sigma \rangle \xrightarrow{\lambda} \langle t'_1 \parallel t_2, \sigma' \rangle}$	(PAR-2) :	$\frac{\langle t_2, \sigma \rangle \xrightarrow{\lambda} \langle t'_2, \sigma' \rangle}{\langle t_1 \parallel t_2, \sigma \rangle \xrightarrow{\lambda} \langle t_1 \parallel t'_2, \sigma' \rangle}$
(PAR'-1) :	$\frac{\langle t_1, \sigma \rangle \xrightarrow{\lambda} \sigma'}{\langle t_1 \parallel t_2, \sigma \rangle \xrightarrow{\lambda} \langle t_2, \sigma' \rangle}$	(PAR'-2) :	$\frac{\langle t_2, \sigma \rangle \xrightarrow{\lambda} \sigma'}{\langle t_1 \parallel t_2, \sigma \rangle \xrightarrow{\lambda} \langle t_1, \sigma' \rangle}$
(SUM-1) :	$\frac{\langle \alpha_1, \sigma \rangle \xrightarrow{\lambda} \sigma'}{\langle \alpha_1.t_1 + \alpha_2.t_2, \sigma \rangle \xrightarrow{\lambda} \langle t_1, \sigma' \rangle}$	(SUM-2) :	$\frac{\langle \alpha_2, \sigma \rangle \xrightarrow{\lambda} \sigma'}{\langle \alpha_1.t_1 + \alpha_2.t_2, \sigma \rangle \xrightarrow{\lambda} \langle t_2, \sigma' \rangle}$
(WHILE) :	$\frac{\langle b, \sigma \rangle \xrightarrow{\lambda} \sigma}{\langle \text{while } b \text{ do } t \text{ od}, \sigma \rangle \xrightarrow{\lambda} \langle p; \text{while } b \text{ do } t \text{ od}, \sigma \rangle}$	(WHILE') :	$\frac{\langle \neg b, \sigma \rangle \xrightarrow{\lambda} \sigma}{\langle \text{while } b \text{ do } t \text{ od}, \sigma \rangle \xrightarrow{\lambda} \sigma}$
(WITH) :	$\langle \text{with } r \text{ do } t \text{ od}, \sigma \oplus \{r\} \rangle \xrightarrow{\text{acq}(r)} \langle t; \text{rel } r, \sigma \rangle$		
(REL) :	$\langle \text{rel } r, \sigma \rangle \xrightarrow{\text{rel}(r)} \sigma \oplus \{r\}$		
(RES) :	$\langle \text{resource } w \text{ do } t \text{ od}, \sigma \rangle \xrightarrow{\text{decl}(r)} \langle [r/w]t; \text{end } r, \sigma \oplus \{r, \text{curr}(r)\} \rangle$		
(END) :	$\langle \text{end } r, \sigma \oplus \{r, \text{curr}(r)\} \rangle \xrightarrow{\text{end}(r)} \sigma$		

Figure 2: Transition semantics

- arrows from events to conditions to represent the postcondition map, and
- tokens (dots) inside conditions to represent the marking.

Action within nets is defined according to a *token game* which defines how the marking of the net changes according to *firing* of the events. An event e can fire if all its preconditions are marked and, following their un-marking, all the postconditions r, σ are not marked. That is, in marking M ,

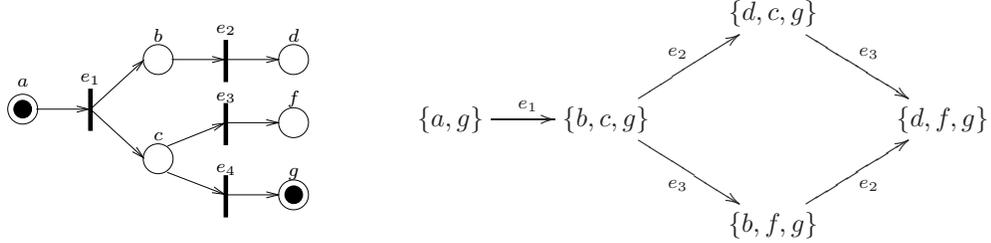
- (1) $\bullet e \subseteq M$
- (2) $(M \setminus \bullet e) \cap e^\bullet = \emptyset$.

Such an event is said to have *concession* or to be *enabled*. The marking following the occurrence of e is obtained by removing the tokens from the preconditions of e and placing a token in every postcondition of e . We write $M \xrightarrow{e} M'$ where

$$M' = (M \setminus \bullet e) \cup e \bullet.$$

If constraint (2) does not hold but constraint (1) does, so the preconditions are all marked (have a token inside) but following removal of the tokens from the preconditions there is a token in some postcondition, there is said to be *contact* in the marking and the event cannot fire.

Consider the following example Petri net, with its transition system between markings derived according to the token game.



The event e_1 is the only event with concession in the initial marking $\{a, g\}$. Its occurrence yields the marking obtained by un-marking its preconditions and marking its postconditions, namely $\{b, c, g\}$. In the marking $\{b, c, g\}$, contact prevents the occurrence of e_4 since its postcondition g is marked following removal of the token from its precondition c . However, in the marking $\{b, c, g\}$ both event e_2 and event e_3 can occur. Note that the occurrence of e_2 in marking $\{b, c, g\}$ does not affect the occurrence of e_3 and *vice versa* since the two events operate on completely disjoint sets of conditions.

For any event $e \in E$, define the notation

$$\bullet e \bullet \stackrel{\text{def}}{=} \bullet e \cup e \bullet.$$

The standard notion of independence within this form of Petri net is to say that two events e_1 and e_2 are independent, written $e_1 I e_2$, if their neighbourhoods are disjoint. That is,

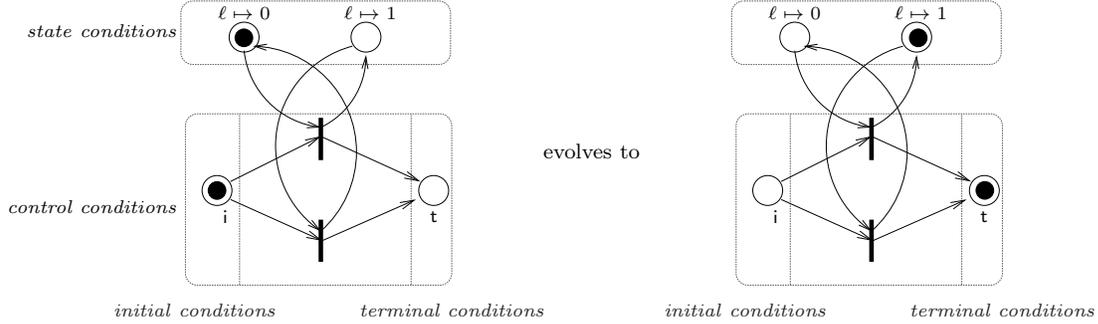
$$e_1 I e_2 \iff \bullet e_1 \bullet \cap \bullet e_2 \bullet = \emptyset.$$

It is easy to see in general that the occurrences of independent events in a marking do not affect each other.

Proposition 3.4. *Let e_1 and e_2 be events of the net N and suppose that $e_1 I e_2$.*

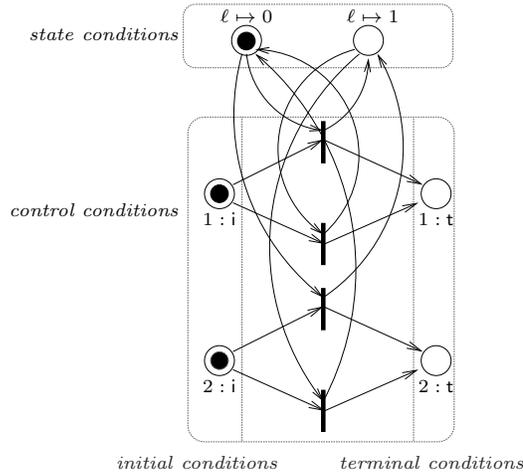
- *If there exist markings M , M' and M_1 of N such that $M \xrightarrow{e_1} M_1$ and $M_1 \xrightarrow{e_2} M'$ then there exists a marking M_2 such that $M \xrightarrow{e_2} M_2$ and $M_2 \xrightarrow{e_1} M'$.*
- *If there exist markings M , M_1 and M_2 of N such that $M \xrightarrow{e_1} M_1$ and $M \xrightarrow{e_2} M_2$ then there exists a marking M' such that $M_1 \xrightarrow{e_2} M'$ and $M_2 \xrightarrow{e_1} M'$. \square*

3.4. Overview of net semantics. Before giving the formal definition of the net semantics of closed terms, by means of an example we shall illustrate how our semantics shall be defined. First, we shall draw the semantics of an action $\text{toggle}(\ell, 0, 1)$ that toggles the value held at a location ℓ between 0 and 1.



Notice that in the above net there are conditions to represent the *shared state* in which processes execute, including for example the values held at locations (we have only drawn conditions that are actually used by the net). There are also conditions to represent the *control point* of the process. The net pictured on the left is in its *initial marking* of control conditions and the net on the right is in its *terminal marking* of control conditions, indicating successful completion of the process following the toggle of the value; the marking of the net initially had the state condition $\ell \mapsto 0$ marked and finished with the condition $\ell \mapsto 1$ marked. There is an event present in the net for each way that the action could take place: one event for toggling the value from 0 to 1 and another event for toggling the value from 1 to 0. Only the first event could occur in the initial marking of the net on the left, and no event can occur in the marking on the right since the control conditions are not appropriately marked.

The parallel composition $\text{toggle}(\ell, 0, 1) \parallel \text{toggle}(\ell, 0, 1)$ can be formed by taking two copies of the net $\text{toggle}(\ell, 0, 1)$ and forcing them to operate on disjoint sets of control conditions.



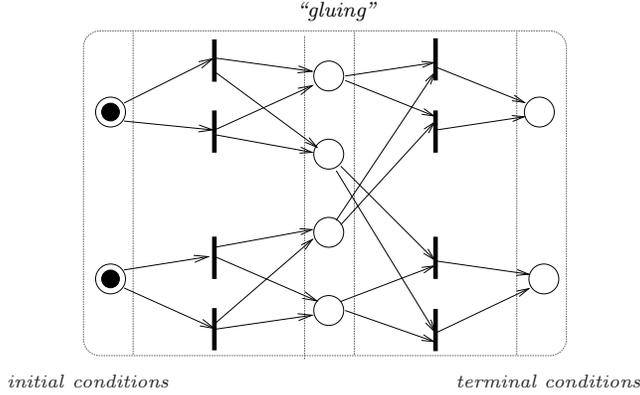
An example run of this net would involve first the top event changing the value of ℓ from 0 to 1 and then the bottom event changing ℓ back from 1 to 0. The resulting marking of

control conditions would be equal to the terminal conditions of the net, so no event would have concession in this marking.

The net representing the sequential composition

$$(\text{toggle}(\ell, 0, 1) \parallel \text{toggle}(\ell, 0, 1)); (\text{toggle}(\ell, 0, 1) \parallel \text{toggle}(\ell, 0, 1))$$

is formed by a ‘gluing’ operation that joins the terminal conditions of one copy of the net for $\text{toggle}(\ell, 0, 1)$ to the initial conditions of another copy of the net for $\text{toggle}(\ell, 0, 1)$. (In this example net, for clarity we shall not show the state conditions.)



3.5. Net structure. As outlined above, within the nets that we give for processes we distinguish two forms of condition, namely *control conditions* and *state conditions*. The markings of these sets of conditions determine the control point of the process and the state in which it is executing, respectively. When we give the net semantics, we will make use of the closure of the set of control conditions under various operations.

Definition 3.5 (Conditions). Define the set of control conditions \mathbf{C} , ranged over by c , to be the least set such that:

- \mathbf{C} contains distinguished elements i and t , standing for ‘initial’ and ‘terminal’, respectively.
- If $c \in \mathbf{C}$ then $r:c \in \mathbf{C}$ for all $r \in \text{Res}$ and $i:c \in \mathbf{C}$ for all $i \in \{1, 2\}$, to distinguish processes working on different resources or arising from different subterms.
- If $c, c' \in \mathbf{C}$ then $(c, c') \in \mathbf{C}$ to allow the ‘gluing’ operation above.

Define the set of *state conditions* \mathbf{S} to be $\mathbf{D} \cup \mathbf{L} \cup \mathbf{R} \cup \mathbf{N}$.

A state $\sigma = (D, L, R, N)$ corresponds to the marking $DULURUN$ of state conditions in the obvious way. Similarly, if C is a marking of control conditions and σ is a state, the pair (C, σ) corresponds to the marking $C \cup \sigma$. We therefore use the notations interchangeably.

The nets that we form shall be *extensional* in the sense that two events are equal if they have the same preconditions and the same postconditions. An event can therefore be regarded as a tuple

$$e = (C, \sigma, C', \sigma')$$

with preconditions $\bullet e \stackrel{\text{def}}{=} C \cup \sigma$ and postconditions $e \bullet \stackrel{\text{def}}{=} C' \cup \sigma'$. To obtain a concise notation for working with events, we write ${}^c e$ for the pre-control conditions of e :

$${}^c e \stackrel{\text{def}}{=} \bullet e \cap \mathbf{C}.$$

We likewise define notations e^C , ${}^D e$, ${}^L e$ etc., and call these the *components* of e by virtue of the fact that it is sufficient to define an event through the definition of its components. The pre-state conditions of e are ${}^S e = {}^D e \cup {}^L e \cup {}^R e \cup {}^N e$, and we define e^S similarly.

Two markings of control conditions are of particular importance: those marked when the process starts executing and those marked when the process has terminated. We call these the *initial* control conditions I and *terminal* control conditions T , respectively. We shall call a net with a partition of its conditions into control and state with the subsets of control conditions I and T an *embedded net*. For an embedded net N , we write $\text{Ic}(N)$ for I and $\text{Tc}(N)$ for T , and we write $\text{Ev}(N)$ for its set of events. Observe that no initial marking of state conditions is specified.

The semantics of a closed term t shall be an embedded net, written $\mathcal{N}[[t]]$. No confusion arises, so we shall write $\text{Ic}(t)$ for $\text{Ic}(\mathcal{N}[[t]])$, and $\text{Tc}(t)$ and $\text{Ev}(t)$ for $\text{Tc}(\mathcal{N}[[t]])$ and $\text{Ev}(\mathcal{N}[[t]])$, respectively. The nets formed shall always have the same sets of control and state conditions; the difference shall arise in the events present in the nets. It would be a trivial matter to restrict to the conditions that are actually used.

As we give the semantics of closed terms, we will make use of several constructions on nets. For example, we wish the events of parallel processes to operate on disjoint sets of control conditions. This is conducted using a *tagging* operation on events. We define $1:e$ to be the event e changed so that

$${}^C(1:e) \stackrel{\text{def}}{=} \{1:c \mid c \in {}^C e\} \quad (1:e)^C \stackrel{\text{def}}{=} \{1:c \mid c \in e^C\}$$

but otherwise unchanged in its action on state conditions. We define the notations $2:e$ and $r:e$ where $r \in \text{Res}$ similarly. The notations are extended pointwise to sets of events:

$$1:E \stackrel{\text{def}}{=} \{1:e \mid e \in E\}.$$

Another useful operation is what we call *gluing* two embedded nets together. For example, when forming the sequential composition of processes $t_1; t_2$, we want to enable the events of t_2 when t_1 has terminated. This is done by ‘gluing’ the two nets together at the terminal conditions of t_1 and the initial conditions of t_2 , having made them disjoint on control conditions using tagging. Wherever a terminal condition c of $\text{Tc}(t_1)$ occurs as a pre- or a postcondition of an event of t_1 , every element of the set $\{1:c\} \times \{2:\text{Ic}(t_2)\}$ would occur in its place. Similarly, the events of t_2 use the set of conditions $\{1:\text{Tc}(t_1)\} \times \{2:c'\}$ instead of an initial condition c' of $\text{Ic}(t_2)$. A variety of control properties that the nets we form possess (Lemma 3.11), such as that all events have at least one pre-control condition, allows us to infer that it is impossible for an event of t_2 to occur before t_1 has terminated, and thereon it is impossible for t_1 to resume. An example follows shortly.

Assume a set $P \subseteq \mathbf{C} \times \mathbf{C}$. Useful definitions to represent gluing are:

$$\begin{aligned} P \triangleleft C &\stackrel{\text{def}}{=} \{(c_1, c_2) \mid c_1 \in C \text{ and } (c_1, c_2) \in P\} \\ &\quad \cup \{c_1 \mid c_1 \in C \text{ and } \nexists c_2. (c_1, c_2) \in P\} \\ P \triangleright C &\stackrel{\text{def}}{=} \{(c_1, c_2) \mid c_2 \in C \text{ and } (c_1, c_2) \in P\} \\ &\quad \cup \{c_2 \mid c_2 \in C \text{ and } \nexists c_1. (c_1, c_2) \in P\} \end{aligned}$$

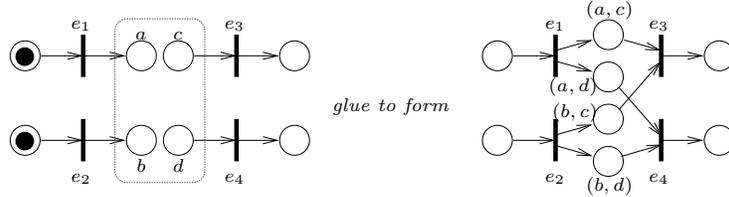
The first definition, $P \triangleleft C$, indicates that an occurrence of c_1 in C is to be replaced by occurrences of (c_1, c_2) for every c_2 such that (c_1, c_2) occurs in P . The second definition, $P \triangleright C$, indicates that an occurrence of c_2 in C is to be replaced by occurrences of (c_1, c_2) for every c_1 such that (c_1, c_2) occurs in P .

The notation is extended to events to give an event $P \triangleleft e$ in the following way, recalling that gluing will only affect the control conditions used by an event and in particular not its state conditions:

$$\begin{aligned} {}^c(P \triangleleft e) &\stackrel{\text{def}}{=} P \triangleleft ({}^c e) & (P \triangleleft e)^c &\stackrel{\text{def}}{=} P \triangleleft (e^c) \\ {}^s(P \triangleleft e) &\stackrel{\text{def}}{=} s_e & (P \triangleleft e)^s &\stackrel{\text{def}}{=} e^s \end{aligned}$$

The notation $P \triangleright e$ is defined similarly, and it is also extended to sets of events in the obvious pointwise manner. For any marking $M = (C, \sigma)$, we will write $P \triangleleft M$ for $(P \triangleleft C, \sigma)$ and similarly write $P \triangleright M$ for $(P \triangleright C, \sigma)$.

To give an example, consider the gluings $P \triangleleft C_1$ and $P \triangleright C_2$ where $C_1 = \{a, b\}$ and $C_2 = \{c, d\}$ are joined at $P = C_1 \times C_2$. Applying $P \triangleleft C_1$ to the left net and $P \triangleright C_2$ to the right net below, this indicates how gluing is used to sequentially compose embedded nets:



The operations of gluing and tagging affect only the control flow of events, not their effect on the marking of state conditions.

Lemma 3.6. *Let N be an embedded net with control conditions \mathbf{C} . Suppose that $P \subseteq \mathbf{C} \times \mathbf{C}$. For any marking M of N and tag $x \in \text{Res} \cup \{1, 2\}$:*

- $M \xrightarrow{e} M'$ iff $x:M \xrightarrow{x:e} x:M'$.
- $M \xrightarrow{e} M'$ iff $P \triangleleft M \xrightarrow{P \triangleleft e} P \triangleleft M'$, and
- $M \xrightarrow{e} M'$ iff $P \triangleright M \xrightarrow{P \triangleright e} P \triangleright M'$.

Furthermore:

- if $1:M \xrightarrow{1:e} M'_1$ then $M'_1 = 1:M'$ for some M' ,
- if $P \triangleleft M \xrightarrow{P \triangleleft e} M'_1$ then $M'_1 = P \triangleleft M'$ for some M' , and
- if $P \triangleright M \xrightarrow{P \triangleright e} M'_2$ then $M'_2 = P \triangleright M'$ for some M' .

Proof. The first and fourth items are straightforward to prove. The remaining properties may be shown using the following easily-demonstrated equations, along with their counterparts for \triangleright , for any subset of control conditions C :

- (1) $C = \emptyset$ iff $P \triangleleft C = \emptyset$,
- (2) $P \triangleleft (C \setminus C') = (P \triangleleft C) \setminus (P \triangleleft C')$,
- (3) $P \triangleleft (C \cup C') = (P \triangleleft C) \cup (P \triangleleft C')$, and
- (4) $P \triangleleft (C \cap C') = (P \triangleleft C) \cap (P \triangleleft C')$. □

3.6. Net semantics. The net semantics that we now give for closed terms is defined by induction on the *size* of terms, given in the obvious way. The reason why it is not given by induction on terms is that the semantics of **resource** w **do** t **od** is given according to the semantics of $[r/w]t$ for all resources r .

▷ **Heap action:** Let $\text{act}_{(C,C')}(D_1, D_2)$ denote an event e with

$${}^C e = C \quad e^C = C' \quad {}^D e = D_1 \quad e^D = D_2$$

and all other components empty, *i.e.* ${}^L e = e^L = {}^R e = e^R = {}^N e = e^N = \emptyset$. For an action α , we define:

$$\text{Ic}(\alpha) \stackrel{\text{def}}{=} \{i\}$$

$$\text{Tc}(\alpha) \stackrel{\text{def}}{=} \{t\}$$

$$\text{Ev}(\alpha) \stackrel{\text{def}}{=} \{\text{act}_{(\{i\},\{t\})}(D_1, D_2) \mid (D_1, D_2) \in \mathcal{A}[\alpha]\}.$$

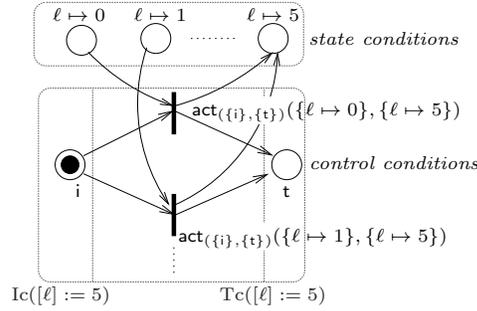
Example 3.7 ($\mathcal{N}[[\ell := 5]]$). Recall that

$$\mathcal{A}[[\ell := 5]] = \{(\{\ell \mapsto v\}, \{\ell \mapsto 5\}) \mid v \in \text{Val}\},$$

so

$$\text{Ev}([\ell := 5]) = \{\text{act}_{(\{i\},\{t\})}(\{\ell \mapsto v\}, \{\ell \mapsto 5\}) \mid v \in \text{Val}\}.$$

The definitions give the net $\mathcal{N}[[\ell := 5]]$:



▷ **Allocation and deallocation:** The command $\text{alloc}(\ell)$ activates, by making current and assigning an arbitrary value to, a non-current location and sets ℓ to point at it. For symmetry, $\text{dealloc}(\ell)$ deactivates the current location pointed to by ℓ .

We begin by defining two further event notations. First, $\text{alloc}_{(C,C')}(\ell, v, \ell', v')$ is the event e such that ${}^C e = C$ and $e^C = C'$ and

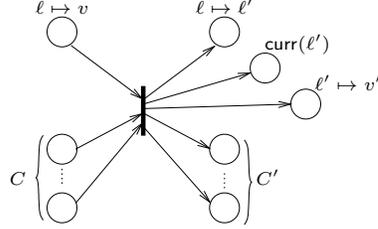
$${}^D e = \{\ell \mapsto v\} \quad e^D = \{\ell \mapsto \ell', \ell' \mapsto v'\} \quad {}^L e = \emptyset \quad e^L = \{\text{curr}(\ell')\},$$

and otherwise empty components, which changes ℓ' from being non-current to current, gives it value v' and changes the value held at ℓ from v to ℓ' . If the condition $\text{curr}(\ell')$ is marked before the event takes place, contact occurs, so the event has concession only if the location ℓ' is not initially current. Second, $\text{dealloc}_{(C,C')}(\ell, \ell', v')$ is the event e such that ${}^C e = C$ and $e^C = C'$ and

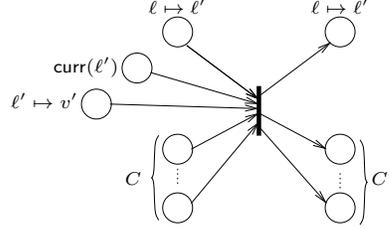
$${}^D e = \{\ell \mapsto \ell', \ell' \mapsto v'\} \quad e^D = \{\ell \mapsto \ell'\} \quad {}^L e = \{\text{curr}(\ell')\},$$

which does the converse of allocation. The location ℓ is left with a dangling pointer to ℓ' . The two events may be drawn as:

$\text{alloc}_{(C,C')}(\ell, v, \ell', v')$:



$\text{dealloc}_{(C,C')}(\ell, \ell', v')$:



The semantics of allocation is given by:

$$\begin{aligned} \text{Ic}(\text{alloc}(\ell)) &\stackrel{\text{def}}{=} \{i\} \\ \text{Tc}(\text{alloc}(\ell)) &\stackrel{\text{def}}{=} \{t\} \\ \text{Ev}(\text{alloc}(\ell)) &\stackrel{\text{def}}{=} \{\text{alloc}_{\{i\},\{t\}}(\ell, v, \ell', v') \mid \ell' \in \text{Loc and } v, v' \in \text{Val}\}. \end{aligned}$$

Note that there is an event present for every value that ℓ might initially hold and every value that ℓ' might be assumed to take initially.

The semantics of disposal is given by:

$$\begin{aligned} \text{Ic}(\text{dealloc}(\ell)) &\stackrel{\text{def}}{=} \{i\} \\ \text{Tc}(\text{dealloc}(\ell)) &\stackrel{\text{def}}{=} \{t\} \\ \text{Ev}(\text{dealloc}(\ell)) &\stackrel{\text{def}}{=} \{\text{dealloc}_{\{i\},\{t\}}(\ell, \ell', v') \mid \ell' \in \text{Loc and } v' \in \text{Val}\}. \end{aligned}$$

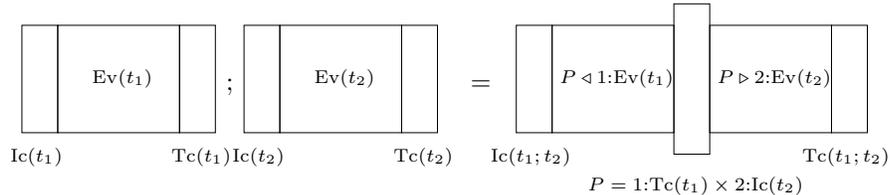
- ▷ **Sequential composition:** The sequential composition of terms involves gluing the terminal marking of the net for t_1 to the initial marking of the net for t_2 . The operation is therefore performed on the set

$$P = 1:\text{Tc}(t_1) \times 2:\text{Ic}(t_2).$$

Following the intuition above, we take

$$\begin{aligned} \text{Ic}(t_1; t_2) &\stackrel{\text{def}}{=} 1:\text{Ic}(t_1) \\ \text{Tc}(t_1; t_2) &\stackrel{\text{def}}{=} 2:\text{Tc}(t_2) \\ \text{Ev}(t_1; t_2) &\stackrel{\text{def}}{=} (P \triangleleft 1:\text{Ev}(t_1)) \cup (P \triangleright 2:\text{Ev}(t_2)). \end{aligned}$$

The formation of the sequential composition on control conditions may be drawn schematically as:



- ▷ **Parallel composition:** The control flow of the parallel composition of processes is autonomous; interaction occurs only through the state. We therefore force the events of the

two processes to work on disjoint sets of control conditions by giving them different tags:

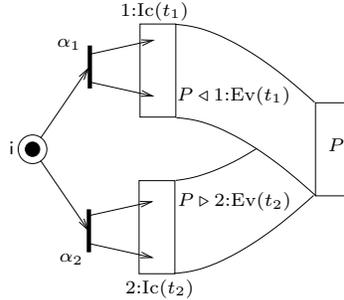
$$\begin{aligned} \text{Ic}(t_1 \parallel t_2) &\stackrel{\text{def}}{=} 1:\text{Ic}(t_1) \cup 2:\text{Ic}(t_2) \\ \text{Tc}(t_1 \parallel t_2) &\stackrel{\text{def}}{=} 1:\text{Tc}(t_1) \cup 2:\text{Tc}(t_2) \\ \text{Ev}(t_1 \parallel t_2) &\stackrel{\text{def}}{=} 1:\text{Ev}(t_1) \cup 2:\text{Ev}(t_2). \end{aligned}$$

Note that the definition of the semantics parallel composition is associative and commutative only if we regard nets up to isomorphism on the control conditions.

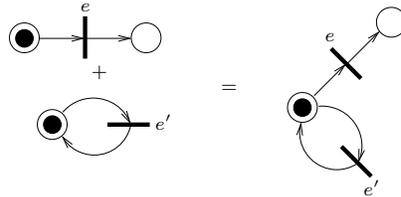
- ▷ **Guarded sum:** Let t be the term $\alpha_1.t_1 + \alpha_2.t_2$. The sum is formed by prefixing the actions onto the tagged nets representing the terms and then gluing the sets of terminal conditions. Let $P = (1:\text{Tc}(t_1)) \times (2:\text{Tc}(t_2))$. Define:

$$\begin{aligned} \text{Ic}(t) &\stackrel{\text{def}}{=} \{i\} \\ \text{Tc}(t) &\stackrel{\text{def}}{=} P \\ \text{Ev}(t) &\stackrel{\text{def}}{=} \{\text{act}_{(\{i\}, 1:\text{Ic}(t_1))}(D_1, D_2) \mid (D_1, D_2) \in \mathcal{A}[\alpha_1]\} \\ &\quad \cup \{\text{act}_{(\{i\}, 2:\text{Ic}(t_2))}(D_1, D_2) \mid (D_1, D_2) \in \mathcal{A}[\alpha_2]\} \\ &\quad \cup P \triangleleft (1:\text{Ev}(t_1)) \cup P \triangleright (2:\text{Ev}(t_2)). \end{aligned}$$

The net may be pictured schematically as follows, in which we have drawn only one representative event for each of α_1 and α_2 , and have elided the effect of these events on state conditions.



On a technical point, one may wonder why the syntax of the language requires that sums possess guards. This is seemingly curious since the category of safe Petri nets, which intuitively underlies a category of embedded nets, has a coproduct construction. However, as remarked in Section 5 of [Win87], there are cases where the coproduct of nets does not coincide with the usual interpretation of nondeterministic sum. In Section 3.3 of [Win86], this is explained as the occurrence net unfolding (the ‘behaviour’) of the coproduct of two nets not being equal to the coproduct of their respective unfoldings. To repeat an example given there, letting $+$ represent coproduct in the category of safe nets, we have:



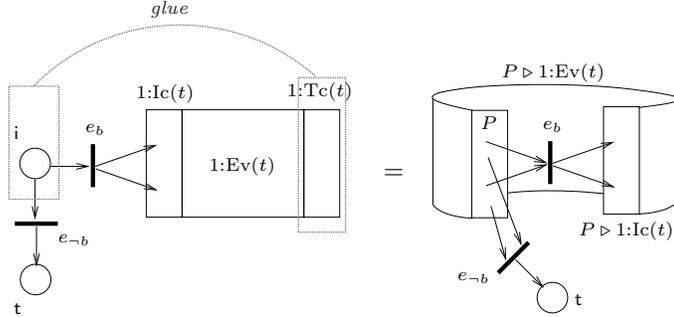
Consequently, using this coproduct as a definition of general sum, the runs of the net representing $\alpha + (\text{while true do } \alpha')$ would consist of some finite number of executions of α' followed, possibly, by one of α . Quite clearly, this does not correspond to the normal understanding of nondeterminism presented in the transition semantics.

The restriction of processes to only use guarded sums allows us to recover the standard interpretation of sums (hence allowing the standard structural operational rule for sums). As stated in [Win87, Win86], another alternative would be to ensure that no event has a postcondition inside the initial conditions of the net. This would necessitate a different semantics for **while** loops, possibly along the lines of [vGV87] which would unfold one iteration of any loop.

- ▷ **Iteration:** To form the net for **while** b **do** t **od** we glue the initial and the terminal conditions of $b.t$ together and then add events to exit the loop when $\neg b$ holds. Let $P = \{i\} \times 1:\text{Tc}(t)$. Define:

$$\begin{aligned} \text{Ic}(\text{while } b \text{ do } t \text{ od}) &\stackrel{\text{def}}{=} P \\ \text{Tc}(\text{while } b \text{ do } t \text{ od}) &\stackrel{\text{def}}{=} \{t\} \\ \text{Ev}(\text{while } b \text{ do } t \text{ od}) &\stackrel{\text{def}}{=} \{\text{act}_{(P, 1:\text{Ic}(t))}(D_t, D_t) \mid (D_t, D_t) \in \mathcal{A}[[b]]\} \\ &\quad \cup \{\text{act}_{(P, \{t\})}(D_f, D_f) \mid (D_f, D_f) \in \mathcal{A}[[\neg b]]\} \\ &\quad \cup P \triangleright (1:\text{Ev}(t)). \end{aligned}$$

The loop can be visualized in the following way (in which we only present one event, e_b , for the boolean b and one event, $e_{\neg b}$, for the boolean $\neg b$):



- ▷ **Critical regions and local resources:** We introduce the following notations for resource events.

$$\begin{aligned} \text{decl}_{(C, C')}(r): & e^{\mathbf{R}} = \{r\} \text{ and } e^{\mathbf{N}} = \{\text{curr}(r)\} \\ \text{end}_{(C, C')}(r): & e^{\mathbf{R}} = \{r\} \text{ and } e^{\mathbf{N}} = \{\text{curr}(r)\} \\ \text{acq}_{(C, C')}(r): & e^{\mathbf{R}} = \{r\} \\ \text{rel}_{(C, C')}(r): & e^{\mathbf{R}} = \{r\} \end{aligned}$$

These all have $e^{\mathbf{C}} = C$ and $e^{\mathbf{C}'} = C'$, and the components other than those listed are empty. Observe that the event $\text{decl}_{(C, C')}(r)$ will avoid contact, and thus be able to occur, only if the resource r is initially non-current.

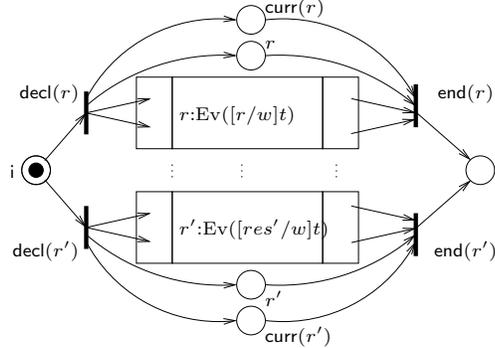
First consider **resource** w **do** t **od**. Its initial and terminal conditions are defined as:

$$\begin{aligned} \text{Ic}(\text{resource } w \text{ do } t \text{ od}) &\stackrel{\text{def}}{=} \{i\} \\ \text{Tc}(\text{resource } w \text{ do } t \text{ od}) &\stackrel{\text{def}}{=} \{t\}. \end{aligned}$$

Its events are defined as:

$$\begin{aligned} & \text{Ev}(\text{resource } w \text{ do } t \text{ od}) \\ \stackrel{\text{def}}{=} & \{ \text{decl}_{(\{i\}, r: \text{Ic}(t'))}(r), \text{end}_{(r: \text{Tc}(t'), \{t\})}(r) \mid r \in \text{Res and } t' = [r/w]t \} \\ & \cup \bigcup \{ r: \text{Ev}(t') \mid r \in \text{Res and } t' = [r/w]t \} \end{aligned}$$

The net formed can be depicted:



As such, the semantics of resource variable binding is a representation of the nondeterministic choice of resource to be selected to be used for the variable. Only one resource shall be chosen for the variable, and it will initially have been non-current thanks to contact described above. Note that the semantics is invariant under α -equivalence \equiv .

Now consider the term `with r do t od`. Its semantics is, informally, to acquire the resource r , then to execute t , and finally to release the resource r :

$$\begin{aligned} \text{Ic}(\text{with } r \text{ do } t \text{ od}) & \stackrel{\text{def}}{=} \{i\} \\ \text{Tc}(\text{with } r \text{ do } t \text{ od}) & \stackrel{\text{def}}{=} \{t\} \\ \text{Ev}(\text{with } r \text{ do } t \text{ od}) & \stackrel{\text{def}}{=} \{ \text{acq}_{(\{i\}, r: \text{Ic}(t))}(r) \} \cup r: \text{Ev}(t) \\ & \cup \{ \text{rel}_{(r: \text{Tc}(t), \{t\})}(r) \}. \end{aligned}$$

3.7. Runs of nets. A well-known property of independence models is that they support a form of run of the net in which independent actions are not interleaved: Given any sequence of events of the net between two markings, we can swap the consecutive occurrences of any two independent events to yield a run between the same two markings. As seen in for example [WN95], this allows us to form an equivalence class of runs between the same markings, generating a Mazurkiewicz trace. This yields a partially ordered multiset, or *pomset*, run [Pra86], in which the independence of event occurrences is captured through them being incomparable.

Definition 3.8. A pomset path of a net $N = (B, E, \bullet(-), (-)^\bullet, M_0)$ is a tuple $\pi = (X, \leq, \lambda)$ such that

- X is a finite set;
- \leq is a partial order on X ;
- $\lambda: X \rightarrow E$; and
- for all $x, x' \in X$, if $x \not\leq x'$ and $x' \not\leq x$ then $\lambda(x) \perp \lambda(x')$.

The elements of X can be thought of via λ as occurrences of events. Where two occurrences are unrelated through the order \leq , they can be thought of as occurring concurrently. Their independence ensures that the effect of this is defined simply as any sequential occurrence of the events.

Definition 3.9. A *sequence* is a path $\pi = (X, \leq, \lambda)$ in which \leq is a total order on X . Let x_1 be the event occurrence least in X according to \leq ; let x_2 be the least event occurrence strictly greater than x_1 ; and so on, all the way up to x_n which is the greatest event occurrence according to \leq for n equal to the size of X (assumed to be finite). The sequence π can be written as e_1, \dots, e_n , where $\lambda(x_i) = e_i$ for all $0 < i \leq n$. Say that a sequence $\pi = e_1, \dots, e_n$ is *from marking M to marking M' in N* if there exist M_0, \dots, M_n such that in N

$$M = M_0 \xrightarrow{e_1} M_1 \dots \xrightarrow{e_n} M_n = M'.$$

Note that the empty path is from marking M to marking M for any marking M . We shall say that a pomset path (X, \leq, λ) is from marking M to M' if there exists any extension of \leq to a total order \leq' such that (X, \leq', λ) is a sequence from M to M' . As discussed, it is a standard result that any other extension of \leq to a total order also yields a path from M to M' .

In fact, when we consider concurrent separation logic, we will only need to consider paths that are sequences, so in the rest of this paper we shall restrict attention to them; all our results generalize straightforwardly to pomsets. From now on, we shall therefore use the terms ‘sequence’, ‘path’ and ‘run’ interchangeably. We have chosen to highlight pomset runs (for conciseness, we have not presented other forms of ‘run’ of a net, such as causal nets) simply to show that Petri nets possess a notion of run that is non-interleaved.

Write $()$ for the path comprising no events and write e for the path with just a single event e . We introduce the notation $\pi : M \xrightarrow{*} M'$ to mean that π is a path from marking M to marking M' , and write $M \xrightarrow{*} M'$ if there exists a path from marking M to marking M' . We shall also write $\pi_1 \cdot \pi_2$ for the composition of sequential paths; clearly, $M \xrightarrow{\pi_1 \cdot \pi_2} M'$ iff there exists M'' such that $M \xrightarrow{\pi_1} M''$ and $M'' \xrightarrow{\pi_2} M'$.

Finally, the tagging and gluing operations are extended to paths pointwise:

$$\begin{aligned} x:(e_1, \dots, e_n) &\stackrel{\text{def}}{=} (x:e_1), \dots, (x:e_n) \\ P \triangleleft (e_1, \dots, e_n) &\stackrel{\text{def}}{=} (P \triangleleft e_1), \dots, (P \triangleleft e_n) \\ P \triangleright (e_1, \dots, e_n) &\stackrel{\text{def}}{=} (P \triangleright e_1), \dots, (P \triangleright e_n) \end{aligned}$$

3.8. Structural properties. Here we establish characterizations of the runs of the net $\mathcal{N}[[t]]$ according to the structure of t . The reader may wish to pass over these technical, but important, details and go directly to Section 3.9.

A complicating factor in characterizing the runs is that that we cannot describe *a priori* the markings reachable in the net for t from an initial state simply from the markings reachable from the nets representing the subterms of t (allowing for the substitution of resources for resource names) running from suitable initial states; this property, as one would expect, fails for parallel composition. However, we can establish properties about the *control flow* of programs. Since such properties are insensitive to the interaction through shared state of parallel processes, they may be established inductively on (the size of) terms. For an event e and markings of control conditions C and C' , we write $C \xrightarrow{e}_C C'$ if the event

e has concession in the marking C when considering only its control conditions, and its occurrence would result in the marking of control conditions C' :

$$C \xrightarrow{e}_C C' \iff {}^c e \subseteq C \text{ and } (C \setminus {}^c e) \cap e^C = \emptyset \text{ and } C' = (C \setminus {}^c e) \cup e^C.$$

We write $\sigma \xrightarrow{e}_S \sigma'$ if the event e has concession on state conditions in the marking σ and its occurrence yields the marking of state conditions σ'

Lemma 3.10. *For any event e and markings C, C' of control conditions and σ, σ' of state conditions, $(C, \sigma) \xrightarrow{e} (C', \sigma')$ iff $C \xrightarrow{e}_C C'$ and $\sigma \xrightarrow{e}_S \sigma'$. \square*

Following the above notation, we shall write $\pi: C \xrightarrow{*}_C C'$ if the path π is from the control marking C to C' , defined in the obvious way. We shall say that a marking C' is *control-reachable* from C , written $C \xrightarrow{*}_C C'$, if there exists a path π such that $\pi: C \xrightarrow{*}_C C'$. A particular consequence of the above lemma is that the marking (C', σ') is reachable from (C, σ) only if C' is control-reachable from C .

We begin with some fairly straightforward properties about the initial and terminal markings and the sets of pre- and postconditions of each event being nonempty. The first and second items of the lemma below could even be seen as part of the definition of embedded net since nonemptiness is necessary for the constructions above to result in nets with the expected behaviours. With the final property, they can be used to show that no event has concession in the terminal marking of the net. The third property eases the definitions constructing $\mathcal{N}[[t]]$.

Lemma 3.11. *For any closed term t and event $e \in \text{Ev}(t)$:*

- (1) $\text{Ic}(t) \neq \emptyset$ and $\text{Tc}(t) \neq \emptyset$,
- (2) ${}^c e \neq \emptyset$ and $e^C \neq \emptyset$,
- (3) $\text{Ic}(t) \cap \text{Tc}(t) = \emptyset$, and
- (4) ${}^c e \cap \text{Tc}(t) = \emptyset$

Proof. The proof follows a simple induction on the size of terms. \square

The following property, that any event occurring from the initial marking of a net has a precondition in the set of initial conditions (and the corresponding property that any event into the terminal marking of the net has a postcondition inside the terminal conditions), follows immediately from the previous lemma. It will be used frequently; for instance, to show that in the net $\mathcal{N}[[t_1; t_2]]$ if e_1 is an event from $\mathcal{N}[[t_1]]$ and e_2 is an event from $\mathcal{N}[[t_2]]$ and e_2 immediately follows e_1 in some sequential run, then there is a control condition that occurs in both the postconditions of e_1 and the preconditions of e_2 . This property is used in Theorem 5.4.

Lemma 3.12. *For any closed term t , event e and marking C of control conditions of $\mathcal{N}[[t]]$:*

- If $\text{Ic}(t) \xrightarrow{e}_C C$ then ${}^\bullet e \cap \text{Ic}(t) \neq \emptyset$.
- If $C \xrightarrow{e}_C \text{Tc}(t)$ then $e^\bullet \cap \text{Tc}(t) \neq \emptyset$. \square

Another important technical property that the embedded nets formed possess is that the marking of control conditions is equal to the set of initial conditions if either only initial conditions are marked or if all initial conditions are marked, for any reachable marking, and the similar statement for the terminal conditions of the net.

Definition 3.13. Say that an embedded net N is *clear* if, for any marking of control conditions C that is control-reachable from $\text{Ic}(N)$:

- (1) if either $C \subseteq \text{Ic}(t)$ or $\text{Ic}(t) \subseteq C$ then $C = \text{Ic}(t)$, and
- (2) if either $C \subseteq \text{Tc}(t)$ or $\text{Tc}(t) \subseteq C$ then $C = \text{Tc}(t)$.

This is used in the proofs characterizing the markings reachable in the net $\mathcal{N}[[t]]$ in terms of the markings reachable in the nets representing t 's subterms (for instance, to show that any run to completion of the net $\mathcal{N}[[t_1; t_2]]$ can be obtained as a run of the net $\mathcal{N}[[t_1]]$ followed by a run of the net $\mathcal{N}[[t_2]]$ since when t_1 in $\mathcal{N}[[t_1; t_2]]$ terminates, precisely the terminal control conditions of $\mathcal{N}[[t_1]]$ will be marked).

Some care is necessary since the proof that, for any closed term t , the net $\mathcal{N}[[t]]$ is clear itself requires understanding of the markings reachable in the net $\mathcal{N}[[t]]$. To resolve this apparent ‘circularity’, when proving the properties required of the net $\mathcal{N}[[t]]$ required to show that the net is clear we shall assume that the nets representing the subterms of t are clear. We shall *then* prove that any net $\mathcal{N}[[t]]$ is clear, allowing us to use elsewhere the properties relating runs of the net $\mathcal{N}[[t]]$ to the runs of the nets of subterms of t . In effect, we will be proving clearness and the structural properties simultaneously, by induction on the size of terms.

3.8.1. Sequential composition. The technique that we use to relate the runs of the net for a term t to the runs of the nets of its subterms is to establish a suitably strong invariant relating the markings arising before and after the occurrence of any event present in $\mathcal{N}[[t]]$, and then perform an induction on the length of sequence. For instance, for sequential composition, we prove:

Lemma 3.14. *Let $P = 1:\text{Tc}(t_1) \times 2:\text{Ic}(t_2)$. Assume that $\mathcal{N}[[t_1]]$ and $\mathcal{N}[[t_2]]$ are clear (Definition 3.13), and consider the net $\mathcal{N}[[t_1; t_2]]$. For any event $e \in \text{Ev}(t_1; t_2)$ and any markings of control conditions C_1 and C_2 :*

- $\text{Ic}(t_1; t_2) = P \triangleleft 1:\text{Ic}(t_1)$ and $\text{Tc}(t_1; t_2) = P \triangleright 2:\text{Tc}(t_2)$.
- $P = P \triangleleft 1:C_1$ iff $C_1 = \text{Tc}(t_1)$, and $P = P \triangleright 2:C_2$ iff $C_2 = \text{Ic}(t_2)$.
- Suppose that C_1 is control-reachable from $\text{Ic}(t_1)$ in $\mathcal{N}[[t_1]]$. If $P \triangleleft 1:C_1 \xrightarrow{e}_{\mathbf{C}} C'$ in $\mathcal{N}[[t_1; t_2]]$ then either $C_1 = \text{Tc}(t_1)$ or there exist C'_1 and e_1 such that $C_1 \xrightarrow{e_1}_{\mathbf{C}} C'_1$ in $\mathcal{N}[[t_1]]$ and $C' = P \triangleleft 1:C'_1$ and $e = P \triangleleft 1:e_1$.
- Suppose that C_2 is control-reachable from $\text{Ic}(t_2)$ in $\mathcal{N}[[t_2]]$. If $P \triangleright 2:C_2 \xrightarrow{e}_{\mathbf{C}} C'$ in $\mathcal{N}[[t_1; t_2]]$ then there exist C'_2 and e_2 such that $C_2 \xrightarrow{e_2}_{\mathbf{C}} C'_2$ in $\mathcal{N}[[t_2]]$ and $C' = P \triangleright 2:C'_2$ and $e = P \triangleright 2:e_2$.

Proof. The first item is simply a re-statement of part of the definition of $\mathcal{N}[[t_1; t_2]]$ and the second item is easy to show. The remaining parts follow an analysis of the events of the net. \square

Using this result, it can be shown that any state reached in $\mathcal{N}[[t_1; t_2]]$ is reached either as a run of $\mathcal{N}[[t_1]]$ or as a run of $\mathcal{N}[[t_1]]$ to a terminal marking followed by a run of $\mathcal{N}[[t_2]]$.

Lemma 3.15. *Suppose that the nets $\mathcal{N}[[t_1]]$ and $\mathcal{N}[[t_2]]$ are clear. If $\pi : \text{Ic}(t_1; t_2) \xrightarrow{*}_{\mathbf{C}} C$ in $\mathcal{N}[[t_1; t_2]]$ then either:*

- there exist C_1 and π_1 such that $C = P \triangleleft 1:C_1$ and $\pi = P \triangleleft 1:\pi_1$ and $\pi_1 : \text{Ic}(t_1) \xrightarrow{*}_{\mathbf{C}} C_1$ in $\mathcal{N}[[t_1]]$, or
- there exist C_2 , π_1 and π_2 such that $C = P \triangleright 2:C_2$ and $\pi = (P \triangleleft 1:\pi_1) \cdot (P \triangleright 2:\pi_2)$ and $\pi_1 : \text{Ic}(t_1) \xrightarrow{*}_{\mathbf{C}} \text{Tc}(t_1)$ in $\mathcal{N}[[t_1]]$ and $\pi_2 : \text{Ic}(t_2) \xrightarrow{*}_{\mathbf{C}} C_2$ in $\mathcal{N}[[t_2]]$,

where $P = 1:\text{Tc}(t_1) \times 2:\text{Ic}(t_2)$.

Proof. A straightforward induction on the length of π using Lemma 3.14. \square

The above lemma can be extended straightforwardly using Lemma 3.10 to obtain the following result involving states, using the fact that the operations of prefixing and tagging do not affect the action of events on state conditions:

Lemma 3.16. *Suppose that the nets $\mathcal{N}[[t_1]]$ and $\mathcal{N}[[t_2]]$ are clear. If $\pi : (\text{Ic}(t_1; t_2), \sigma_0) \longrightarrow^* (C, \sigma)$ in $\mathcal{N}[[t_1; t_2]]$ then either:*

- *there exist C_1 and π_1 such that $C = P \triangleleft 1:C_1$ and $\pi = P \triangleleft 1:\pi_1$ and $\pi_1 : (\text{Ic}(t_1), \sigma_0) \longrightarrow^* (C_1, \sigma)$ in $\mathcal{N}[[t_1]]$, or*
- *there exist C_2, σ', π_1 and π_2 such that $C = P \triangleright 2:C_2$ and $\pi = (P \triangleleft 1:\pi_1) \cdot (P \triangleright 2:\pi_2)$ and $\pi_1 : (\text{Ic}(t_1), \sigma_0) \longrightarrow^* (\text{Tc}(t_1), \sigma')$ in $\mathcal{N}[[t_1]]$ and $\pi_2 : (\text{Ic}(t_2), \sigma') \longrightarrow^* (C_2, \sigma)$ in $\mathcal{N}[[t_2]]$,*

where $P = 1:\text{Tc}(t_1) \times 2:\text{Ic}(t_2)$. \square

The converse result, that runs of the nets $\mathcal{N}[[t_1]]$ and $\mathcal{N}[[t_2]]$, with appropriate intermediate states, give rise to runs of the net $\mathcal{N}[[t_1; t_2]]$ can also be shown.

3.8.2. *Parallel composition.* Runs of control within the net $\mathcal{N}[[t_1 \parallel t_2]]$ are amenable to a similar (though in fact less complicated) analysis to that presented in Lemmas 3.14 and 3.15:

Lemma 3.17. *Consider the net $\mathcal{N}[[t_1 \parallel t_2]]$.*

- $\text{Ic}(t_1 \parallel t_2) = 1:\text{Ic}(t_1) \cup 2:\text{Ic}(t_2)$ and $\text{Tc}(t_1 \parallel t_2) = 1:\text{Tc}(t_1) \cup 2:\text{Tc}(t_2)$.
- *For any markings C_1, C_2 and C' of control conditions and any event $e \in \text{Ev}(t_1 \parallel t_2)$, if $1:C_1 \cup 2:C_2 \xrightarrow{e}_C C'$ in $\mathcal{N}[[t_1 \parallel t_2]]$ then either:*
 - *there exists $e_1 \in \text{Ev}(t_1)$ such that $e = 1:e_1$ and there exists C'_1 such that $C' = 1:C'_1 \cup 2:C_2$ and $C_1 \xrightarrow{e_1}_C C'_1$ in $\mathcal{N}[[t_1]]$, or*
 - *there exists $e_2 \in \text{Ev}(t_2)$ such that $e = 2:e_2$ and there exists C'_2 such that $C' = 1:C_1 \cup 2:C'_2$ and $C_2 \xrightarrow{e_2}_C C'_2$ in $\mathcal{N}[[t_2]]$.*

Proof. A straightforward examination of the events of $\mathcal{N}[[t_1 \parallel t_2]]$. \square

Using the preceding lemma, the paths of the net $\mathcal{N}[[t_1 \parallel t_2]]$ on control conditions can be characterized as:

Lemma 3.18. *If $\pi : \text{Ic}(t_1 \parallel t_2) \longrightarrow_C^* C$ in $\mathcal{N}[[t_1 \parallel t_2]]$ then any event e in π is either equal to $1:e_1$ for some event $e_1 \in \text{Ev}(t_1)$ or equal to $2:e_2$ for some event $e_2 \in \text{Ev}(t_2)$. Furthermore, there exist C_1 and C_2 such that $C = 1:C_1 \cup 2:C_2$ and*

$$\pi_1 : \text{Ic}(t_1) \longrightarrow_C^* C_1 \text{ and } \pi_2 : \text{Ic}(t_2) \longrightarrow_C^* C_2,$$

where π_1 is obtained by removing events equal to $2:e_2$ for some e_2 from π , and π_2 is obtained by removing events equal to $1:e_1$ for some e_1 from π .

Proof. Induction on the length of path π . \square

Notably there is no analogue to Lemma 3.16 involving the markings of state conditions for the parallel composition.

3.8.3. Iteration. The net $\mathcal{N}[\text{while } b \text{ do } t_0 \text{ od}]$ allows runs that start with an event that either shows that the boolean b holds or an event that shows that b fails. If b fails, the net enters its terminal marking and no further action occurs. If the boolean b passes, a run of the net $\mathcal{N}[t_0]$ occurs, followed by the net re-entering its initial control state. The following lemma captures this; it is proved by establishing an invariant in the same way as was done for the sequential composition, though for brevity we shall omit it.

Lemma 3.19. *Let $t \equiv \text{while } b \text{ do } t_0 \text{ od}$ and suppose that $\mathcal{N}[t_0]$ is clear. Let $P = \{i\} \times 1:\text{Tc}(t_1)$, and recall that $P = \text{Ic}(t)$. Assume that π is a path such that $\pi : \text{Ic}(t) \longrightarrow_C^* C$ in $\mathcal{N}[t]$ for some C . There exists a natural number $n \geq 0$ and a (possibly empty if $n = 0$) collection of paths π_1, \dots, π_n and heaps D_1, \dots, D_n such that, for each path π_i :*

$$\begin{array}{ll} \pi_i & : \text{Ic}(t_0) \longrightarrow_C^* \text{Tc}(t_0) & \text{in } \mathcal{N}[t_0] \\ \text{act}_{(\text{Ic}(b), \text{Tc}(b))}(D_i, D_i) & : \text{Ic}(b) \longrightarrow_C \text{Tc}(b) & \text{in } \mathcal{N}[b]. \end{array}$$

Write e_i for the event $\text{act}_{(P, 1:\text{Ic}(t_0))}(D_i, D_i)$. Either:

- $C = \text{Ic}(t)$ and $\pi = e_1 \cdot (P \triangleright 1:\pi_1) \cdot \dots \cdot e_n \cdot (P \triangleright 1:\pi_n)$;
- $C = P \triangleleft 1:C'$ for some marking of control conditions C' and there exists a path π' and heap D' such that

$$\pi = e_1 \cdot (P \triangleright 1:\pi_1) \cdot \dots \cdot e_n \cdot (P \triangleright 1:\pi_n) \cdot \text{act}_{(P, 1:\text{Ic}(t_0))}(D', D') \cdot (P \triangleright 1:\pi')$$

and

$$\begin{array}{ll} \pi' & : \text{Ic}(t_0) \longrightarrow_C^* C' & \text{in } \mathcal{N}[t_0] \\ \text{act}_{(\text{Ic}(b), \text{Tc}(b))}(D', D') & : \text{Ic}(b) \longrightarrow_C \text{Tc}(b) & \text{in } \mathcal{N}[b]; \text{ or} \end{array}$$

- $C = \text{Tc}(t)$ and there exists a heap D' such that

$$\pi = e_1 \cdot (P \triangleright 1:\pi_1) \cdot \dots \cdot e_n \cdot (P \triangleright 1:\pi_n) \cdot \text{act}_{(P, \text{Tc}(t))}(D', D')$$

and $\text{act}_{(\text{Ic}(\neg b), \text{Tc}(\neg b))}(D', D') : \text{Ic}(b) \longrightarrow_C \text{Tc}(b)$ in $\mathcal{N}[\neg b]$. \square

The three possible cases for the control marking C above correspond to net being in its initial control state (following some number of iterations), the net being in the body of the loop, and the net being in its terminal control state following exit of the loop.

3.8.4. Sums. The behaviour of the net $\mathcal{N}[\alpha_1.t_1 + \alpha_2.t_2]$ can be characterized as either the occurrence of an event of the action α_1 followed by a run of t_1 or the occurrence of an event of the action α_2 followed by a run of t_2 . Note that if $C = P \triangleleft 1:C_1$ then $C = \text{Tc}(\alpha_1.t_1 + \alpha_2.t_2)$ if, and only if, $C_1 = \text{Tc}(t_1)$, and the similar property for t_2 .

Lemma 3.20. *Let $t \equiv \alpha_1.t_1 + \alpha_2.t_2$ and $P = 1:\text{Tc}(t_1) \times 2:\text{Tc}(t_2)$ and suppose that the nets $\mathcal{N}[t_1]$ and $\mathcal{N}[t_2]$ are clear. If π is a path $\pi : \text{Ic}(t) \longrightarrow_C^* C$ in $\mathcal{N}[t]$ for some C then:*

- $C = \text{Ic}(t)$ and $\pi = ()$, or
- $C = P \triangleleft 1:C_1$ for some C_1 and $\pi = \text{act}_{(\text{Ic}(t), 1:\text{Ic}(t_1))}(D_1, D'_1) \cdot (P \triangleleft 1:\pi_1)$ for some π_1, D_1, D'_1 such that

$$\begin{array}{ll} \text{act}_{(\text{Ic}(\alpha_1), \text{Tc}(\alpha_1))}(D_1, D'_1) & : \text{Ic}(\alpha_1) \longrightarrow_C \text{Tc}(\alpha_1) & \text{in } \mathcal{N}[\alpha_1] \\ \pi_1 & : \text{Ic}(t_1) \longrightarrow_C^* C_1 & \text{in } \mathcal{N}[t_1]; \text{ or} \end{array}$$

- $C = P \triangleright 2 : C_2$ for some C_2 and $\pi = \mathbf{act}_{(\mathbf{Ic}(t), 2:\mathbf{Ic}(t_2))}(D_2, D'_2) \cdot (P \triangleright 2 : \pi_2)$ for some π_2, D_2, D'_2 such that

$$\begin{array}{lll} \mathbf{act}_{(\mathbf{Ic}(\alpha_2), \mathbf{Tc}(\alpha_2))}(D_2, D'_2) & : \mathbf{Ic}(\alpha_2) \longrightarrow_{\mathbf{C}} \mathbf{Tc}(\alpha_2) & \text{in } \mathcal{N}[[\alpha_2]] \\ \pi_2 & : \mathbf{Ic}(t_2) \longrightarrow_{\mathbf{C}}^* C_2 & \text{in } \mathcal{N}[[t_2]]. \end{array}$$

Proof. An induction following establishing an invariant in the style of Lemma 3.14. \square

3.8.5. *Resource declaration.* A consequence of the following result is that any complete run of the net $\mathbf{resource } w \text{ do } t_0 \text{ od}$ consists first of an event that chooses a resource r to be used for w , then a run of $[r/w]t_0$, and finally an event that records that r is no longer in use.

Lemma 3.21. *Suppose that the net $\mathcal{N}[[r/w]t_0]$ is clear for any resource r and let $t \equiv \mathbf{resource } w \text{ do } t_0 \text{ od}$. If in the net $\mathcal{N}[[t]$ we have $\pi : \mathbf{Ic}(t) \longrightarrow_{\mathbf{C}}^* C$ then either:*

- $C = \mathbf{Ic}(t)$ and $\pi = ()$, or
- there exist $r \in \mathbf{Res}$ and C' and π_0 such that $C = r : C'$ and

$$\pi = \mathbf{decl}_{(\{i\}, r : \mathbf{Ic}([r/w]t_0))}(r) \cdot (r : \pi_0)$$

and

$$\begin{array}{lll} \mathbf{decl}_{(\{i\}, r : \mathbf{Ic}([r/w]t_0))} & : \mathbf{Ic}(t) \longrightarrow_{\mathbf{C}} r : \mathbf{Ic}([r/w]t_0) & \text{in } \mathcal{N}[[t]] \text{ and} \\ \pi_0 & : \mathbf{Ic}([r/w]t_0) \longrightarrow_{\mathbf{C}}^* C' & \text{in } \mathcal{N}[[r/w]t_0], \text{ or} \end{array}$$

- $C = \mathbf{Tc}(t)$ and there exist $r \in \mathbf{Res}$ and π_0 such that

$$\pi = \mathbf{decl}_{(\{i\}, r : \mathbf{Ic}([r/w]t_0))}(r) \cdot (r : \pi_0) \cdot \mathbf{end}_{(r : \mathbf{Tc}([r/w]t_0), \{t\})}(r)$$

and

$$\begin{array}{lll} \mathbf{decl}_{(\{i\}, r : \mathbf{Ic}([r/w]t_0))} & : \mathbf{Ic}(t) \longrightarrow_{\mathbf{C}} r : \mathbf{Ic}([r/w]t_0) & \text{in } \mathcal{N}[[t]], \\ \pi_0 & : \mathbf{Ic}([r/w]t_0) \longrightarrow_{\mathbf{C}}^* \mathbf{Tc}([r/w]t_0) & \text{in } \mathcal{N}[[r/w]t_0], \text{ and} \\ \mathbf{end}_{(r : \mathbf{Tc}([r/w]t_0), \{t\})} & : r : \mathbf{Tc}([r/w]t_0) \longrightarrow_{\mathbf{C}} \mathbf{Tc}(t) & \text{in } \mathcal{N}[[t]]. \end{array}$$

Proof. By establishing an invariant on markings between the occurrences of single events, as in Lemma 3.14. \square

3.8.6. *Critical regions.* The net $\mathcal{N}[[\mathbf{with } r \text{ do } t_0 \text{ od}]]$ starts by acquiring the resource r . If this action cannot proceed because the resource is unavailable, no event will occur. If the resource is available, the process behaves as t_0 , and then releases the resource r if t_0 terminates.

Lemma 3.22. *Let $t \equiv \mathbf{with } r \text{ do } t_0 \text{ od}$ and suppose that the net $\mathcal{N}[[t_0]]$ is clear. If in the net $\mathcal{N}[[t]$ we have $\pi : \mathbf{Ic}(t) \longrightarrow_{\mathbf{C}}^* C$ then either:*

- $C = \mathbf{Ic}(t)$ and $\pi = ()$,
- $C = r : C_0$ for some marking of control conditions C_0 and $\pi = \mathbf{acq}_{(\mathbf{Ic}(t), r : \mathbf{Ic}(t_0))}(r) \cdot (r : \pi_0)$ for some path π_0 such that $\pi_0 : \mathbf{Ic}(t_0) \longrightarrow_{\mathbf{C}}^* C_0$ in $\mathcal{N}[[t_0]]$, or
- $C = \mathbf{Tc}(t)$ and $\pi = \mathbf{acq}_{(\mathbf{Ic}(t), r : \mathbf{Ic}(t_0))}(r) \cdot (r : \pi_0) \cdot \mathbf{rel}_{(r : \mathbf{Tc}(t_0), \mathbf{Tc}(t))}(r)$ for some path π_0 such that $\pi_0 : \mathbf{Ic}(t_0) \longrightarrow_{\mathbf{C}}^* \mathbf{Tc}(t_0)$ in $\mathcal{N}[[t_0]]$. \square

3.8.7. *Clearness.* Now that we have established these control properties of the runs of processes, we can show that the clearness property of Definition 3.13 does indeed hold in the net $\mathcal{N}[[t]]$ for any term t .

Lemma 3.23. *For any closed term t , the net $\mathcal{N}[[t]]$ is clear.*

Proof. Following the observation that

$$\begin{aligned} 1:C \subseteq 1:C' & \text{ iff } C \subseteq C' \\ P \triangleleft C \subseteq P \triangleleft C' & \text{ iff } C \subseteq C' \\ P \triangleright C \subseteq P \triangleright C' & \text{ iff } C \subseteq C', \end{aligned}$$

the property can be proved by induction on the size of terms using the above control properties. \square

3.8.8. *Preservation of consistency.* The final attribute that we aim towards is that any marking of state conditions σ reachable in $\mathcal{N}[[t]]$ from a consistent initial marking of state conditions σ_0 is itself consistent. The only challenge here will be showing that if $r \in \sigma$ then $\text{curr}(r) \in \sigma$, which shall require some understanding of the nature of the critical regions present in our semantics; the other requirements for consistency are straightforwardly shown to be preserved through the occurrence of the events present in $\mathcal{N}[[t]]$.

We shall first show that any release of a resource is dependent on the prior acquisition of that resource: for any sequence π and any resource there exists an injection f that associates any occurrence of a release event to a prior occurrence of an acquisition event of that resource, and between the two occurrences there are no other actions on that resource.

Lemma 3.24. *Let π be a sequence of events, $\pi = (e_1, \dots, e_n)$. For any closed term t , resource r and marking of control conditions C such that $\pi : \text{Ic}(t) \xrightarrow{*}_C C$ in $\mathcal{N}[[t]]$, there exists a partial function $f : \mathbb{N} \rightarrow \mathbb{N}$ satisfying, for all $i, j \in \mathbb{N}$:*

- f is injective,
- if there exist sets of control conditions C_1, C_2 such that $e_i = \text{rel}_{(C_1, C_2)}(r)$ then $f(i)$ defined, and
- if $f(i)$ defined then $f(i) < i$ and there exist sets of control conditions C_1, C_2 such that $e_{f(i)} = \text{acq}_{(C_1, C_2)}(r)$.

Moreover, if there exist markings of state conditions $\sigma_0, \dots, \sigma_n$ and markings of control conditions C_0, \dots, C_n such that $(C_{i-1}, \sigma_{i-1}) \xrightarrow{e_i} (C_i, \sigma_i)$ for all i with $0 < i \leq n$ and $C_0 = \text{Ic}(t)$, then there exists an f satisfying the above constraints and such that, for all k with $i < k < f(i)$, there exist no C' and C'' such that either $e_k = \text{acq}_{(C', C'')}(r)$ or $e_k = \text{rel}_{(C', C'')}(r)$.

Proof. The first property is shown, using the control properties of sequences established above, by induction on the size of terms. The second property arises since if $e_i = \text{acq}_{(C_i, C'_i)}(r)$ and $e_j = \text{acq}_{(C_j, C'_j)}(r)$ for $i < j$ then there must exist k such that $i < k < j$ and $e_k = \text{rel}_{(C_k, C'_k)}(r)$, and the symmetric property for release events. \square

We are now able to show that the nets formed preserve the consistency of the markings of state conditions.

Lemma 3.25 (Preservation of consistent markings). *For any closed term t , if $(\text{Ic}(t), \sigma_0) \longrightarrow^* (C, \sigma)$ in the net $\mathcal{N}[[t]]$ and the marking σ_0 of state conditions is consistent then σ is consistent.*

Proof. It is straightforward to prove by induction on the size of the term t that the events present in that net $\mathcal{N}[[t]]$ are all of one of the following forms:

$$\begin{array}{cccc} \text{act}_{(C, C')}(D, D') & \text{alloc}_{(C, C')}(\ell, v, \ell', v') & \text{decl}_{(C, C')}(r) & \text{acq}_{(C, C')}(r) \\ & \text{dealloc}_{(C, C')}(\ell, \ell', v') & \text{end}_{(C, C')}(r) & \text{rel}_{(C, C')}(r) \end{array}$$

It is readily shown that each form of event preserves the consistency of the marking of state conditions, apart from showing that if $r \in \sigma$ then $\text{curr}(r) \in \sigma$.

Suppose, for contradiction, that π' is a path such that $\pi' : (\text{Ic}(t), \sigma_0) \longrightarrow^* (C, \sigma)$ in $\mathcal{N}[[t]]$ and that $r \in \sigma$ but $\text{curr}(r) \notin \sigma$. Assume, furthermore, and without loss of generality, that any other marking of state conditions σ' along π has the property that if $r \in \sigma$ then $\text{curr}(r) \in \sigma$. It must be the case that $\pi' = \pi \cdot \text{rel}_{(D_1, D'_1)}(r)$ for some D_1, D'_1 and π . By Lemma 3.24, there exist D_2, D'_2, π_1 and π_2 such that $\pi = \pi_1 \cdot \text{acq}_{(D_2, D'_2)}(r) \cdot \pi_2$ and no event in π_2 is an $\text{acq}(r)$ or $\text{rel}(r)$ event. Let $\pi_1 : (\text{Ic}(t), \sigma_0) \longrightarrow^* (C_1, \sigma_1)$. We must have $r \in \sigma_1$, and by assumption $\text{curr}(r) \in \sigma_1$. It can be seen that we must have $\text{curr}(r) \in \sigma'$ and $r \notin \sigma'$ for all states σ' reached along $\text{acq}_{(D_2, D'_2)}(r) \cdot \pi_2$ from (C_1, σ_1) since no $\text{end}(r)$ event can have concession in such markings. Consequently, we must have $\text{curr}(r) \in \sigma_2$ for σ_2 obtained by following the path $\pi : (\text{Ic}(t), \sigma_0) \longrightarrow^* (C_2, \sigma_2)$, and therefore $\text{curr}(r) \in \sigma$. \square

The structure of processes ensures that any resource initially current remains current through the execution of the net. The same property working backwards from the terminal marking of the net also holds.

Lemma 3.26. *Let σ, σ' be a consistent markings of state conditions. For any markings of control conditions C, C' :*

- (1) *If $(\text{Ic}(t), \sigma) \longrightarrow^* (C', \sigma')$ in $\mathcal{N}[[t]]$ and $\text{curr}(r) \in \sigma$ then $\text{curr}(r) \in \sigma'$.*
- (2) *If $(C, \sigma) \longrightarrow^* (\text{Tc}(t), \sigma')$ in $\mathcal{N}[[t]]$ and $\text{curr}(r) \in \sigma'$ then $\text{curr}(r) \in \sigma$.*

Proof. We shall only show (1) since (2) is similar. An induction on the size of terms using the control properties above gives the following:

- If there exists a sequence π such that $\pi \cdot \text{end}_{(C_1, C_2)}(r) : \text{Ic}(t) \longrightarrow_C^* C$ for some C_1, C_2 then there exists an event $\text{decl}_{(C'_1, C'_2)}(r)$ in π for some C'_1, C'_2 .

Let π' be a sequence $\pi' : (\text{Ic}(t), \sigma) \longrightarrow^* (C', \sigma')$ and assume that $\text{curr}(r) \in \sigma$. Without loss of generality, suppose that (C', σ') is the earliest marking along π' from $(\text{Ic}(t), \sigma)$ such that $\text{curr}(r) \notin \sigma'$; otherwise, we can take the initial segment of π' with this property. Examination of the events given by our semantics reveals that the last event in π' is an $\text{end}_{(C_1, C_2)}(r)$ event, since otherwise $\text{curr}(r)$ is not in the state prior to σ' . Now, applying the result above informs that there is an event $\text{decl}_{(C'_1, C'_2)}(r)$ in π' and this must occur before $\text{end}_{(C_1, C_2)}(r)$. Now, the event $\text{decl}_{(C'_1, C'_2)}(r)$ can only occur in a marking σ_0 of state conditions such that $\text{curr}(r) \notin \sigma_0$, but this contradicts our assumption that σ' was the first marking of state conditions reachable along π' from $(\text{Ic}(t), \sigma)$ with $\text{curr}(r) \notin \sigma'$. \square

3.9. Correspondence of semantics. As we have progressed, the event notations introduced have corresponded to labels of the transition semantics. Write $|e|$ for the label corresponding to event e . Before progressing to consider separation logic, we shall give a theorem¹ that shows how the net and transition semantics correspond. It assumes a definition of open map bisimulation [JNW93, NW96] based on paths as pomsets, $(N, M) \sim (N', M')$, relating paths of net N from marking M to paths of N' from M' . The bisimulations that we form respect terminal markings and markings of state conditions.

Theorem 3.27 (Correspondence). *Let t be a closed term and σ be a consistent state.*

- If $\langle t, \sigma \rangle \xrightarrow{\lambda} \sigma'$ then there exists e such that $|e| = \lambda$ and $(\text{Ic}(t), \sigma) \xrightarrow{e} (\text{Tc}(t), \sigma')$ in $\mathcal{N}[[t]]$.
- If $\langle t, \sigma \rangle \xrightarrow{\lambda} \langle t', \sigma' \rangle$ then there exists e such that $|e| = \lambda$ and $(\text{Ic}(t), \sigma) \xrightarrow{e} (C', \sigma')$ in $\mathcal{N}[[t]]$ and $(\mathcal{N}[[t]], C', \sigma') \sim (\mathcal{N}[[t']], \text{Ic}(t'), \sigma')$.
- If $(\text{Ic}(t), \sigma) \xrightarrow{e} (C', \sigma')$ in $\mathcal{N}[[t]]$ then either there exists t' such that $\langle t, \sigma \rangle \xrightarrow{|e|} \langle t', \sigma' \rangle$ and $(\mathcal{N}[[t]], C', \sigma') \sim (\mathcal{N}[[t']], \text{Ic}(t'), \sigma')$, or $\langle t, \sigma \rangle \xrightarrow{|e|} \sigma'$ and $C' = \text{Tc}(t)$. \square

Write $(t, \sigma) \sim (t', \sigma')$ iff there exist a label-preserving bisimulation (in the standard sense) between the transitions systems for t from initial state σ and t' from σ' . From the preceding result, we obtain adequacy of our semantics:

Corollary 3.28 (Adequacy). *Let t, t' be closed terms and σ, σ' be consistent states. If $(\mathcal{N}[[t]], \text{Ic}(t), \sigma) \sim (\mathcal{N}[[t']], \text{Ic}(t'), \sigma')$ then $(t, \sigma) \sim (t', \sigma')$. \square*

The converse property with respect to \sim fails. For instance, for any σ we have

$$(\alpha_1 \parallel \alpha_2, \sigma) \sim (\alpha_1.\alpha_2 + \alpha_2.\alpha_1, \sigma).$$

However, the definition of open bisimulation on the nets with pomsets as paths yields

$$(\mathcal{N}[[\alpha_1 \parallel \alpha_2]], \text{Ic}(\alpha_1 \parallel \alpha_2), \sigma) \not\sim (\mathcal{N}[[\alpha_1.\alpha_2 + \alpha_2.\alpha_1]], \text{Ic}(\alpha_1.\alpha_2 + \alpha_2.\alpha_1), \sigma).$$

The reason why the property fails is that the transition system does not capture the independence of actions.

4. SEPARATION LOGIC

As discussed in the introduction, concurrent separation logic establishes partial correctness assertions about concurrent heap-manipulating programs; that whenever a given program running from a heap satisfying a *heap formula* φ terminates, the resulting heap satisfies a heap formula ψ . The semantics of the heap logic arises as an instance of the logic of Bunched Implications [OP99]. At its core are the associated notions of heap composition and the separating conjunction. Two heaps may be composed if they are defined over disjoint sets of locations:

$$D_1 \cdot D_2 \stackrel{\text{def}}{=} D_1 \cup D_2 \text{ if } \text{dom}(D_1) \cap \text{dom}(D_2) = \emptyset.$$

¹The proof of this theorem is rather technical and requires a presentation of open maps on the category of embedded Petri nets, so we shall not present the proof here. It shall appear, with the other omitted results, in the first author's PhD thesis.

A heap satisfies the separating conjunction $\varphi_1 * \varphi_2$ if it can be split into two parts, one satisfying φ_1 and the other φ_2 :

$$D \models \varphi_1 * \varphi_2 \quad \text{iff} \quad \text{there exist } D_1, D_2 \text{ such that } D_1 \cdot D_2 \text{ defined and} \\ D = D_1 \cdot D_2 \text{ and } D_1 \models \varphi_1 \text{ and } D_2 \models \varphi_2.$$

The semantics of the other parts of the heap logic is of little significance when considering the semantics of the program logic. For completeness, however, it is defined by induction on the size of formulæ in Figure 3 where the full syntax also appears. Unlike the heap logic presented in [Bro07], we do not allow arithmetic on memory locations; this is just to simplify the presentation, and such arithmetic could easily be added. Since we distinguish the types of locations and values, we use x_{loc} as the logical variable for locations and x_{val} for the logical variable for values. We adopt the usual binding precedences, and $*$ binds more tightly the standard logical connectives. We define the shorthand notation $\ell \mapsto -$ for $\exists x_{\text{val}}(\ell \mapsto x_{\text{val}})$. We shall write $\models \varphi$ if $D \models \varphi$ for all heaps D , and write $\varphi \implies \psi$ if $\models \varphi \rightarrow \psi$.

We now present the intuition for the key judgement of concurrent separation logic, $\Gamma \vdash \{\varphi\} t \{\psi\}$, where φ and ψ are formulæ of the heap logic, and Γ is a *environment of resource invariants*, of the form $r_1 : \chi_1, \dots, r_n : \chi_n$, associating invariants χ_i with resources r_i . (We refer the reader to [O’H07] for a fuller introduction.) Informally, the judgement means:

In any run from a heap satisfying φ and the invariants Γ , the process t never accesses locations that it does not own, and if the process t terminates then it does so in a heap satisfying ψ and the invariants Γ .

Central to this understanding is the notion of *ownership*, which we capture formally in Section 4.1. Initially the process t is considered to *own* that part of the heap which satisfies φ , and accordingly to *own* the locations in that subheap. As t runs the locations it owns may change as it acquires and releases resources, and correspondingly the locations used in justifying their invariants.

Ownership plays a key role in making the judgements of concurrent separation logic compositional: a judgement $\Gamma \vdash \{\varphi\} t \{\psi\}$ should hold even if other (unknown) processes are to execute in the same heap. It is therefore necessary to make certain assumptions about the ways in which these other processes might interact with the process t . This is achieved through ownership, by assuming that each process owns, throughout its execution, a separate, though possibly changing, part of the heap; the part of the heap that each process owns must not be accessed by any other process; moreover a process must not access locations it does not own.

The rules of concurrent separation logic are presented in Figure 4 in the style of [Bro07]. The only significant difference between the two systems is that we omit the rules for auxiliary variables and for existential quantification. Both are omitted for simplicity since they are peripheral to the focus of our work.

As a first example, the rule for heap actions (L-ACT) would allow the judgement

$$\Gamma \vdash \{\ell \mapsto 0\} [\ell] := 1 \{\ell \mapsto 1\}$$

since the process is initially assumed to own the location ℓ because the part of the heap that the process initially owns satisfies $\ell \mapsto 0$. The resulting part of the heap owned by the

<i>Variables:</i>	$x ::= x_{\text{loc}}$	Location variable
	$ x_{\text{val}}$	Value variable
<i>Location expressions:</i>	$e_{\text{loc}} ::= x_{\text{loc}}$	Location variable
	$ \ell$	Location, $\ell \in \text{Loc}$
<i>Expressions:</i>	$e ::= e_{\text{loc}}$	Location expression
	$ x_{\text{val}}$	Value variable
	$ v$	Value, $v \in \text{Val}$
<i>Formulae:</i>	$\varphi ::= e_{\text{loc}} \mapsto e_{\text{val}}$	heap location
	$ \varphi * \varphi$	separating conjunction
	$ \mathbf{empty}$	empty heap
	$ \varphi \wedge \varphi$	conjunction
	$ \varphi \vee \varphi$	disjunction
	$ \varphi \rightarrow \varphi$	implication
	$ \neg \varphi$	negation
	$ \exists x. \varphi$	existential quantification
	$ \forall x. \varphi$	universal quantification
	$ e = e$	equality
	$ \top$	true
	$ \perp$	false
<i>Semantics of closed formulae:</i>		
$D \models \ell \mapsto v$	iff	$D = \{\ell \mapsto v\}$
$D \models \varphi_1 * \varphi_2$	iff	there exist D_1, D_2 such that $D_1 \cdot D_2$ defined and $D = D_1 \cdot D_2$ and $D_1 \models \varphi_1$ and $D_2 \models \varphi_2$
$D \models \mathbf{empty}$	iff	$D = \emptyset$
$D \models \varphi_1 \wedge \varphi_2$	iff	$D \models \varphi_1$ and $D \models \varphi_2$
$D \models \varphi_1 \vee \varphi_2$	iff	$D \models \varphi_1$ or $D \models \varphi_2$
$D \models \varphi_1 \rightarrow \varphi_2$	iff	$D \models \varphi_1$ implies $D \models \varphi_2$
$D \models \neg \varphi$	iff	not $D \models \varphi$
$D \models \exists x_{\text{loc}}. \varphi$	iff	there exists $\ell \in \text{Loc}$ such that $D \models [\ell/x_{\text{loc}}]\varphi$
$D \models \exists x_{\text{val}}. \varphi$	iff	there exists $v \in \text{Val}$ such that $D \models [v/x_{\text{val}}]\varphi$
$D \models \forall x_{\text{loc}}. \varphi$	iff	for all $\ell \in \text{Loc}$: $D \models [\ell/x_{\text{loc}}]\varphi$
$D \models \forall x_{\text{val}}. \varphi$	iff	for all $v \in \text{Val}$: $D \models [v/x_{\text{val}}]\varphi$
$D \models v = v'$	iff	$v = v'$
$D \models \top$	always	
$D \models \perp$	never	

Figure 3: Syntax and semantics of the heap logic

	for all $D \models \varphi$ and $(D_1, D_2) \in \mathcal{A}[\alpha]$: $\text{dom}(D_1) \subseteq \text{dom}(D)$ and $D_1 \subseteq D$ implies $(D \setminus D_1) \cup D_2 \models \psi$
(L-ACT) :	$\frac{\Gamma \vdash \{\varphi\} \alpha \{\psi\}}{\Gamma \vdash \{\varphi\} \alpha \{\psi\}}$
(L-ALLOC) :	$\Gamma \vdash \{\ell \mapsto -\} \text{alloc}(\ell) \{\exists x_{\text{loc}}(\ell \mapsto x_{\text{loc}} * x_{\text{loc}} \mapsto -)\}$
(L-DEALLOC) :	$\Gamma \vdash \{\exists x_{\text{loc}}(\ell \mapsto x_{\text{loc}} * x_{\text{loc}} \mapsto -)\} \text{dealloc}(\ell) \{\exists x_{\text{loc}}(\ell \mapsto x_{\text{loc}})\}$
(L-SEQ) :	$\frac{\Gamma \vdash \{\varphi\} t_1 \{\varphi'\} \quad \Gamma \vdash \{\varphi'\} t_2 \{\psi\}}{\Gamma \vdash \{\varphi\} t_1; t_2 \{\psi\}}$
(L-SUM) :	$\frac{\Gamma \vdash \{\varphi\} \alpha_1 \{\varphi_1\} \quad \Gamma \vdash \{\varphi\} \alpha_2 \{\varphi_2\} \quad \Gamma \vdash \{\varphi_1\} t_1 \{\psi\} \quad \Gamma \vdash \{\varphi_2\} t_2 \{\psi\}}{\Gamma \vdash \{\varphi\} \alpha_1.t_1 + \alpha_2.t_2 \{\psi\}}$
(L-WHILE) :	$\frac{\Gamma \vdash \{\varphi\} b \{\varphi'\} \quad \Gamma \vdash \{\varphi\} \neg b \{\psi\} \quad \Gamma \vdash \{\varphi'\} t \{\varphi\}}{\Gamma \vdash \{\varphi\} \text{while } b \text{ do } t \text{ od } \{\psi\}}$
(L-RES) :	$\frac{\Gamma, r : \chi \vdash \{\varphi\} [r/w]t \{\psi\}}{\Gamma \vdash \{\varphi * \chi\} \text{resource } w \text{ do } t \text{ od } \{\psi * \chi\}} \left(\begin{array}{l} \chi \text{ precise} \\ r \notin \text{dom}(\Gamma) \end{array} \right)$
(L-CR) :	$\frac{\Gamma, r : \chi \vdash \{\varphi * \chi\} t \{\psi * \chi\}}{\Gamma, r : \chi \vdash \{\varphi\} \text{with } r \text{ do } t \text{ od } \{\psi\}}$
(L-PAR) :	$\frac{\Gamma \vdash \{\varphi_1\} t_1 \{\psi_1\} \quad \Gamma \vdash \{\varphi_2\} t_2 \{\psi_2\}}{\Gamma \vdash \{\varphi_1 * \varphi_2\} t_1 \parallel t_2 \{\psi_1 * \psi_2\}}$
(L-FRAME) :	$\frac{\Gamma \vdash \{\varphi\} t \{\psi\}}{\Gamma \vdash \{\varphi * \varphi'\} t \{\psi * \varphi'\}}$
(L-CONSEQUENCE) :	$\frac{\varphi \implies \varphi' \quad \Gamma \vdash \{\varphi'\} t \{\psi'\} \quad \psi' \implies \psi}{\Gamma \vdash \{\varphi\} t \{\psi\}}$
(L-CONJUNCTION) :	$\frac{\Gamma \vdash \{\varphi_1\} t \{\psi_1\} \quad \Gamma \vdash \{\varphi_2\} t \{\psi_2\}}{\Gamma \vdash \{\varphi_1 \wedge \varphi_2\} t \{\psi_1 \wedge \psi_2\}}$
(L-DISJUNCTION) :	$\frac{\Gamma \vdash \{\varphi_1\} t \{\psi_1\} \quad \Gamma \vdash \{\varphi_2\} t \{\psi_2\}}{\Gamma \vdash \{\varphi_1 \vee \varphi_2\} t \{\psi_1 \vee \psi_2\}}$
(L-EXPANSION) :	$\frac{\Gamma \vdash \{\varphi\} t \{\psi\}}{\Gamma, \Gamma' \vdash \{\varphi\} t \{\psi\}}$
(L-CONTRACTION) :	$\frac{\Gamma, \Gamma' \vdash \{\varphi\} t \{\psi\}}{\Gamma \vdash \{\varphi\} t \{\psi\}} (\text{res}(t) \subseteq \text{dom}(\Gamma))$

Figure 4: Rules of concurrent separation logic

process satisfies $\ell \mapsto 1$. The judgement

$$\Gamma \vdash \{\text{empty}\} [\ell] := 2 \{\top\}$$

is *not* derivable however: the part of the heap initially owned by the process satisfies **empty**, and therefore the process initially does not own the location ℓ . Assignment to ℓ violates the principle that the process may only act on locations that it owns — the so-called *frame property*.

An instance of the separating conjunction is seen in the rule for parallel composition, (L-PAR):

$$\frac{\Gamma \vdash \{\varphi_1\} t_1 \{\psi_1\} \quad \Gamma \vdash \{\varphi_2\} t_2 \{\psi_2\}}{\Gamma \vdash \{\varphi_1 * \varphi_2\} t_1 \parallel t_2 \{\psi_1 * \psi_2\}}$$

Informally, the rule is sound because the part of the initial heap that is owned by the process $t_1 \parallel t_2$ can be split into two parts, one part satisfying φ_1 owned by t_1 and the other satisfying φ_2 owned by t_2 ; as the processes execute the subheaps that we see each as owning remain disjoint from each other and end up separately satisfying ψ_1 and ψ_2 .

It is vital that the logic enforces the requirement that processes only act on locations that they own. If this requirement were not imposed, so that the judgement

$$\Gamma \vdash \{\text{empty}\} [\ell] := 2 \{\top\}$$

were derivable, then the rule for parallel composition could be applied with the other judgement above to conclude that

$$\Gamma \vdash \{\ell \mapsto 0 * \text{empty}\} [\ell] := 1 \parallel [\ell] := 2 \{\ell \mapsto 1 * \top\}.$$

This flawed assertion would imply that whenever the process $[\ell] := 1 \parallel [\ell] := 2$ runs from a state satisfying $\ell \mapsto 0$, the resulting state has $\ell \mapsto 1$, which is obviously wrong.

The notion of ownership is subtle since the collection of locations that a process owns may change as the process evolves. As seen in the rule (L-ALLOC), the intuitive reading is that after an allocation event has taken place the process owns the newly current location. Similarly, deallocation of a location leads to loss of ownership. For example, it is possible to make the judgement

$$\Gamma \vdash \{\ell \mapsto -\} \text{alloc}(\ell) \{\exists x_{\text{loc}}. \ell \mapsto x_{\text{loc}} * x_{\text{loc}} \mapsto -\}.$$

If the new location were ℓ' which initially held value v , this would mean that in the the (fragment of the) resulting heap $\{\ell \mapsto \ell' * \ell' \mapsto v\}$, the locations ℓ and ℓ' would be owned by the process. Consequently, an action $[[\ell]] := 0$ which assigns 0 to the location pointed to by ℓ resulting in the heap $\{\ell \mapsto \ell', \ell' \mapsto 0\}$ allows the judgement

$$\Gamma \vdash \{\exists x_{\text{loc}}. \ell \mapsto x_{\text{loc}} * x_{\text{loc}} \mapsto -\} [[\ell]] := 0 \{\exists x_{\text{loc}}. \ell \mapsto x_{\text{loc}} * x_{\text{loc}} \mapsto 0\}$$

by (L-ACT) since both locations would be owned by the process. The rule (L-SEQ) can now be applied to obtain

$$\Gamma \vdash \{\ell \mapsto -\} \text{alloc}(\ell); [[\ell]] := 0 \{\exists x_{\text{loc}}. \ell \mapsto x_{\text{loc}} * x_{\text{loc}} \mapsto 0\},$$

indicating that the process has ownership of the location ℓ' , seen in the ability to write to ℓ' , once it has been allocated.

To allow the logic to make judgements beyond those applicable to the almost ‘disjointly concurrent’ programs outlined so far, further interaction is allowed through a system of *invariants*. The judgement environment Γ records a formula called an invariant for each resource in its domain, which contains all the resources occurring in the term. The intuition

is that, whenever a resource r with an invariant χ is available, there is part of the heap unowned by any other process and protected by the resource that satisfies χ . In such a situation, we shall say that the locations used to satisfy χ are ‘owned’ by the invariant for r . Processes may gain ownership of these locations, and thereby the right to access them, by entering a critical region protected by the resource. When the process leaves the critical region, the invariant must be restored and the ownership of the locations used to satisfy the invariant is relinquished. This is reflected in the rule (L-CR). As an example, we have the following derivation:

$$\frac{\text{(L-ACT)} \frac{}{r : \ell \mapsto 0 \vdash \{\ell' \mapsto - * \ell \mapsto 0\}} [\ell'] := [\ell] \{\ell' \mapsto 0 * \ell \mapsto 0\}}{\text{(L-CR)} \frac{}{r : \ell \mapsto 0 \vdash \{\ell' \mapsto -\} \text{ with } r \text{ do } [\ell'] := [\ell] \text{ od } \{\ell' \mapsto 0\}}}$$

The process initially owns the location ℓ' , and the location ℓ is protected by the resource r . We reason about the process inside the critical region running from a state with ownership of the locations governed by the invariant in addition to those that it owned before entering the critical region since no other process can be operating on them; that is, we reason about $[\ell'] := [\ell]$ with locations ℓ and ℓ' owned by the process. However, when the process leaves the critical region, ownership of the locations used to satisfy the invariant is lost, indicated by the conclusion $\ell' \mapsto 0$ in the judgement rather than $\ell' \mapsto 0 * \ell \mapsto 0$.

An invariant is required to be a *precise* heap logic formula.

Definition 4.1 (Precision). A heap logic formula χ is *precise* if for any heap D there is at most one subheap $D_0 \subseteq D$ such that $D_0 \models \chi$.

We leave discussion of the rôle of precision to the conclusion, though it might be seen to be of use since it identifies uniquely the part of the heap that is owned by the invariant if the resource is available. Formally, Γ ranges over finite partial functions from resources to precise heap formulæ. We write $\text{dom}(\Gamma)$ for the set of resources on which Γ is defined, and write Γ, Γ' for the union of the two partial functions, defined only if $\text{dom}(\Gamma) \cap \text{dom}(\Gamma') = \emptyset$. We write $r : \chi$ for the singleton environment taking resource r to χ , and we allow ourselves to write $r : \chi \in \Gamma$ if $\Gamma(r) = \chi$.

The rules allow ownership of locations to be transferred through invariants. Consider the invariant χ defined as $\ell' \mapsto 0 \vee (\ell' \mapsto 1 * \ell \mapsto 0)$. If the resource is available, the invariant is satisfied: it either protects the location ℓ' , which has value 0, or it protects location ℓ' , which has value 1, as well as location ℓ . A process can acquire ownership of ℓ across a critical region by changing the value of ℓ' from 1 to 0 and may leave ownership of ℓ inside the invariant by changing the value of ℓ' from 0 to 1.

Assume, for example, that the process owns location ℓ . The only way in which the invariant χ can be satisfied disjointly from the locations that the process owns is for ℓ' to hold value 0. That is, we have

$$\ell \mapsto 0 * (\ell' \mapsto 0 \vee (\ell' \mapsto 1 * \ell \mapsto 0)) \implies \ell \mapsto 0 * \ell' \mapsto 0$$

which is implicitly used in the instance of the rule (L-CONSEQUENCE) below. Consequently, as the process enters a critical region protected by r , it gains ownership of location ℓ' . If the process sets the value of ℓ' to 1, when the process leaves the critical region it must restore the invariant to the resource, and so relinquish ownership of both ℓ' and ℓ . This is seen in the derivation of the following judgement, in which we take $\Gamma = r : \chi$.

$$\begin{array}{c}
\text{(L-ACT)} \\
\text{(L-CONSEQUENCE)} \frac{\Gamma \vdash \{\ell \mapsto 0 * \ell' \mapsto 0\} [\ell'] := 1 \{\ell \mapsto 0 * \ell' \mapsto 1\}}{\Gamma \vdash \{\ell \mapsto 0 * \chi\} [\ell'] := 1 \{\text{empty} * \chi\}} \\
\text{(L-CR)} \frac{\Gamma \vdash \{\ell \mapsto 0\} \text{ with } r \text{ do } [\ell'] := 1 \text{ od } \{\text{empty}\}}{\Gamma \vdash \{\ell \mapsto 0\} \text{ with } r \text{ do } [\ell'] := 1 \text{ od } \{\text{empty}\}}
\end{array}$$

With this derivation, we can derive

$$\Gamma \vdash \{\ell \mapsto 2\} [\ell] := 0; \text{ with } r \text{ do } [\ell'] := 1 \{\text{empty}\}.$$

It is also possible to acquire ownership of locations through an invariant. Let the action **diverge** have the same semantics as that of the boolean guard **false**, which is an action that can never occur *i.e.* the process is stuck. We have the following derivation:

$$\begin{array}{c}
\Gamma \vdash \{\chi\} [\ell'] = 0 \{\ell' \mapsto 0\} \\
\Gamma \vdash \{\ell' \mapsto 0\} \text{ diverge } \{\ell' \mapsto 0 * \ell \mapsto 0\} \\
\Gamma \vdash \{\chi\} [\ell'] = 1 \{\ell' \mapsto 1 * \ell \mapsto 0\} \\
\Gamma \vdash \{\ell' \mapsto 1 * \ell \mapsto 0\} [\ell'] := 0 \{\ell' \mapsto 0 * \ell \mapsto 0\} \\
\text{(L-SUM)} \frac{\Gamma \vdash \{\chi\} ([\ell'] = 0.\text{diverge}) + ([\ell'] = 1.[\ell'] := 0) \{\ell \mapsto 0 * \ell' \mapsto 0\}}{\Gamma \vdash \{\text{empty} * \chi\} ([\ell'] = 0.\text{diverge}) + ([\ell'] = 1.[\ell'] := 0) \{\ell \mapsto 0 * \chi\}} \\
\text{(L-CONSEQUENCE)} \frac{\Gamma \vdash \{\text{empty} * \chi\} ([\ell'] = 0.\text{diverge}) + ([\ell'] = 1.[\ell'] := 0) \{\ell \mapsto 0 * \chi\}}{\Gamma \vdash \{\text{empty}\} \text{ with } r \text{ do } ([\ell'] = 0.\text{diverge}) + ([\ell'] = 1.[\ell'] := 0) \text{ od } \{\ell \mapsto 0\}} \\
\text{(L-CR)}
\end{array}$$

The undischarged hypotheses at the top of the derivation are all proved by the rule (L-ACT). Let t_0 denote the process $([\ell'] = 0.\text{diverge}) + ([\ell'] = 1.[\ell'] := 0)$. Observe that the process **with** r **do** t_0 **od** is considered to own no part of the initial heap. As the process enters the critical region, it is considered to take ownership of the part of the heap satisfying the invariant for r , *viz* χ . There are two ways in which χ might be satisfied:

- (1) It may be that the process gains ownership of the location ℓ' which holds value 0. In this case, only the guard $[\ell'] = 0$ of t_0 can pass, so the process must evolve to **diverge** and therefore never terminates. It is therefore trivially true that the remainder of the derivation, that if the process t_0 terminates then the part of the heap that it owns satisfies $\ell \mapsto 0 * \chi$ and therefore after leaving the critical region and losing ownership of the locations satisfying χ that the process owns location ℓ , is sound.
- (2) The process might have taken control of the locations ℓ , holding value 0, and ℓ' , holding value 1. Inside the critical region, the process t_0 can be seen to change the value of ℓ' from 1 to 0. The only way that the invariant χ can then be satisfied is by the location ℓ' holding 0, so ownership of ℓ' is lost as the process leaves the critical region. Importantly, the process retains ownership of location ℓ .

Using the derivations given above, we can give an example of ownership of ℓ , as exhibited by the right to write to ℓ , being transferred (we have annotated internal assertions arising

from the proofs above inside the program):

$$\Gamma \vdash \begin{array}{l} \{\ell \mapsto 2\} \\ [\ell] := 0; \\ \{\ell \mapsto 0\} \\ \text{with } r \text{ do} \\ \quad [\ell'] := 1 \\ \text{od} \\ \{\text{empty}\} \end{array} \quad \begin{array}{l} \{\ell \mapsto 2\} \\ \parallel \\ \{\text{empty}\} \\ \text{with } r \text{ do} \\ \quad [\ell'] = 0. \text{diverge} \\ \quad + [\ell'] = 1. [\ell'] := 0 \\ \text{od;} \\ \{\ell \mapsto 0\} \\ [\ell] := 1 \\ \{\ell \mapsto 1\} \\ \parallel \\ \{\ell \mapsto 1\} \end{array}$$

We also see that, in any terminating run of this process, it must be the case that the process on the left terminates strictly before the process on the right begins.

The final remark to be made on the rules of the logic is that (L-RES) allows invariants to be established for newly declared resources. We reason about the closed term $[r/w]t$, for an arbitrary ‘fresh’ resource r ; it is sufficient to consider only one such resource, as shall be seen in Lemma 4.25. The resource r is known not to occur in the domain of Γ and hence does not occur in the term t thanks to the following lemma, proved straightforwardly by induction on the judgement.

Lemma 4.2. *If $\Gamma \vdash \{\varphi\} t \{\psi\}$ then $\text{res}(t) \subseteq \text{dom}(\Gamma)$.* □

4.1. Ownership model. We now progress to give a formal interpretation of the rules presented in the previous section. The key idea is that the judgement $\Gamma \vdash \{\varphi\} t \{\psi\}$ is robust against the operation of other ‘external’ processes (which have themselves been subject to a judgement in the logic) on the state, so that the rule for parallel composition is valid. From the account presented earlier, external processes may act on the heap providing they do not access the locations ‘owned’ by the process t , and they may act to acquire and release resources providing they respect the invariants in Γ . External processes may also make non-current resources current through the instantiation of a resource variable and might make such resources non-current. The semantics of judgements must therefore keep a record of how each current location in the heap and each current resource is owned: whether the process might access the location, whether it forms part of an invariant protected by a resource, or whether external processes might act on that location, along with a similar record for resources. The semantics will include *interference events* to represent such forms of action by external processes.

Capturing these requirements, we construct an *interference net* with respect to the environment Γ to represent the execution of suitable external processes proved against Γ . This involves creating *ownership conditions* $\omega_{\text{proc}}(\ell)$, $\omega_{\text{inv}}(\ell)$ and $\omega_{\text{oth}}(\ell)$ for each location ℓ . The intuition is that $\omega_{\text{proc}}(\ell)$ is marked if ℓ is owned by the process, $\omega_{\text{inv}}(\ell)$ if ℓ is used to satisfy the invariant for an available open resource, and $\omega_{\text{oth}}(\ell)$ is marked if ℓ is current but owned by another process.

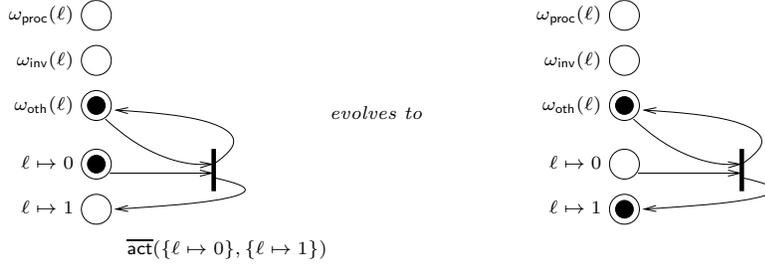
To give an example, suppose that we have the judgement

$$\Gamma \vdash \{k \mapsto 1\} [k] := 0 \{k \mapsto 0\}.$$

The proof can be composed with the judgement $\Gamma \vdash \{\ell \mapsto 0\} [\ell] := 0 \{\ell \mapsto 1\}$ to obtain

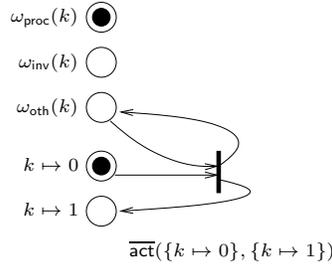
$$\Gamma \vdash \{k \mapsto 1 * \ell \mapsto 0\} [k] := 0 \parallel [\ell] := 1 \{k \mapsto 0 * \ell \mapsto 1\}.$$

The first proof, that the assignment $[k] := 0$ changes the value at k from 1 to 0, must take into account the possibility that the values held at other locations may change. In particular, it must take into account the possibility that the value at ℓ (not to equal k) changes from 0 to 1. We therefore reason about the net $\mathcal{N}[[k] := 0]$ in the presence of the following interference event, which changes the value held at ℓ from 0 to 1:



Notably, the above event requires that the location ℓ is owned by an external process, *i.e.* the condition $\omega_{\text{oth}}(\ell)$ is marked.

Since we do not know with which other judgements $\Gamma \vdash \{k \mapsto 1\} [k] := 0 \{k \mapsto 0\}$ may be composed, there are interference events present in the net for all the forms of interference permissible according to the notion of ownership. For instance, the interference event which changes the value of k from 0 to 1



is present in the net. However, the judgement asserts that k is owned by the process, so this interference event (and indeed any other interference event that affects k) will not be able to occur because the condition $\omega_{\text{proc}}(k)$ will be marked, not $\omega_{\text{oth}}(k)$.

As mentioned above, we introduce interference events to mimic the action of external processes on resources. The notion of ownership is therefore extended in this setting to resources, for example so that an external process cannot be allowed to release a resource held by the current process. It is important to make a distinction between resources in the domain of the environment Γ (called *open* resources) and those that are not (called *closed* resources): Open resources have invariants associated with them, so the ownership of the heap is affected by events that acquire or release them, as presented earlier in this section; this is not the case for closed resources. Closed resources are those resources made current to instantiate a local resource variable. They may either be used by the process being considered if it declared the resource, or be used by some external process if some external process declared the resource. We shall introduce conditions $\omega_{\text{proc}}(r)$, $\omega_{\text{inv}}(r)$ and $\omega_{\text{oth}}(r)$ for each resource r . The condition $\omega_{\text{proc}}(r)$ will be marked if either the resource is closed and was made current by the process or if the resource is open and is held by the process.

The condition $\omega_{\text{inv}}(r)$ will be marked if r is open and available. The condition $\omega_{\text{oth}}(r)$ will be marked if either the resource is closed and was made current by an external process or if the resource is both open and the external process holds it.

The set of ownership conditions is denoted \mathbf{W} :

$$\mathbf{W} \stackrel{\text{def}}{=} \{\omega_{\text{proc}}(\ell), \omega_{\text{inv}}(\ell), \omega_{\text{oth}}(\ell) \mid \ell \in \text{Loc}\} \\ \cup \{\omega_{\text{proc}}(r), \omega_{\text{inv}}(r), \omega_{\text{oth}}(r) \mid r \in \text{Res}\}.$$

We use W to range over markings of ownership conditions and introduce the notations ${}^{\mathbf{w}}e$ and $e^{\mathbf{w}}$, as before, for the sets of pre-ownership conditions of e and post-ownership conditions of e , respectively. For a set of locations L , we define the notation

$$\omega_{\text{proc}}(L) \stackrel{\text{def}}{=} \{\omega_{\text{proc}}(\ell) \mid \ell \in L\},$$

and define $\omega_{\text{inv}}(L)$ and $\omega_{\text{oth}}(L)$ similarly. Only certain markings of ownership conditions are consistent with a state σ :

Definition 4.3 (Consistent marking). The marking of state and ownership conditions (σ, W) of $\mathcal{W}[[t]]_{\Gamma}$ is *consistent* if:

- (1) σ is a consistent state in $\mathcal{N}[[t]]$,
- (2) for each $z \in \text{Loc} \cup \text{Res}$, at most one of $\{\omega_{\text{proc}}(z), \omega_{\text{inv}}(z), \omega_{\text{oth}}(z)\}$ is marked,
- (3) for each $z \in \text{Loc} \cup \text{Res}$, the ownership condition $\text{curr}(z)$ is in σ iff precisely one of $\{\omega_{\text{proc}}(z), \omega_{\text{inv}}(z), \omega_{\text{oth}}(z)\}$ is in W ,
- (4) if $r \in \text{dom}(\Gamma)$ and $r \in R$ then $\omega_{\text{inv}}(r) \in W$,
- (5) if $r \in \text{dom}(\Gamma)$ and $r \notin R$ then either $\omega_{\text{proc}}(r) \in W$ or $\omega_{\text{oth}}(r) \in W$, and
- (6) if $\text{curr}(r) \in \sigma$ and $r \notin \text{dom}(\Gamma)$ then either $\omega_{\text{proc}}(r) \in W$ or $\omega_{\text{oth}}(r) \in W$.

Requirements (2) and (3) assert that W is essentially a function from the set of current locations and resources to describe their ownership. Requirement (4) states that any available open resource is owned as an invariant: it can be accessed either by the process being considered or by an external process, and there is an invariant associated with r . Requirement (5) states that any unavailable open resource is either held by the process or by an external process. Requirement (6) asserts that any closed resource is owned either by the current process or by an external process.

Table 1 defines a number of notations for events corresponding to the permitted interference described. To summarize, there will be interference events to represent the following kinds of action by external processes:

- $\overline{\text{act}}(D_1, D_2)$: Arbitrary action on the heap (excluding allocation or deallocation) owned by external processes.
- $\overline{\text{alloc}}(\ell, v, \ell', v')$: Allocation of a new location ℓ' by an external process, storing the result in the location ℓ . The location ℓ must initially have been owned by an external process. Ownership of the new location ℓ' is taken by the external process.
- $\overline{\text{dealloc}}(\ell, v, \ell', v')$: Disposal of the location ℓ' pointed to by ℓ . Both locations are initially owned by external processes, so $\omega_{\text{oth}}(\ell)$ and $\omega_{\text{oth}}(\ell')$ are preconditions to the event.
- $\overline{\text{decl}}(r)$: Declaration of a resource r . The condition $\text{curr}(r)$ is marked by the event, so the resource was not initially current. Ownership of r is taken by the external process, so $\omega_{\text{oth}}(r)$ is in the postconditions of the event.
- $\overline{\text{end}}(r)$: End of scope of a resource r , only permissible if the resource was initially declared by an external process and therefore $\omega_{\text{oth}}(r)$ is marked.

Abbreviation u	Preconditions		Postconditions	
	$^s u$	$^w u$	u^s	u^w
$\overline{\text{act}}(D_1, D_2)$	D_1	$\omega_{\text{oth}}(\text{dom}(D_1))$	D_2	$\omega_{\text{oth}}(\text{dom}(D_2))$
$\overline{\text{alloc}}(\ell, v, \ell', v')$	$\{\ell \mapsto v\}$	$\{\omega_{\text{oth}}(\ell)\}$	$\{\text{curr}(\ell')\} \cup$ $\{\ell \mapsto \ell', \ell' \mapsto v'\}$	$\{\omega_{\text{oth}}(\ell), \omega_{\text{oth}}(\ell')\}$
$\overline{\text{dealloc}}(\ell, \ell', v')$	$\{\text{curr}(\ell')\} \cup$ $\{\ell \mapsto \ell', \ell' \mapsto v'\}$	$\{\omega_{\text{oth}}(\ell), \omega_{\text{oth}}(\ell')\}$	$\{\ell \mapsto \ell'\}$	$\{\omega_{\text{oth}}(\ell)\}$
$\overline{\text{decl}}(r)$	$\{\}$	$\{\}$	$\{\text{curr}(r), r\}$	$\{\omega_{\text{oth}}(r)\}$
$\overline{\text{end}}(r)$	$\{\text{curr}(r), r\}$	$\{\omega_{\text{oth}}(r)\}$	$\{\}$	$\{\}$
$\overline{\text{acq}}(r)$	$\{r\}$	$\{\omega_{\text{oth}}(r)\}$	$\{\}$	$\{\omega_{\text{oth}}(r)\}$
$\overline{\text{rel}}(r)$	$\{\}$	$\{\omega_{\text{oth}}(r)\}$	$\{r\}$	$\{\omega_{\text{oth}}(r)\}$
$\overline{\text{acq}}(r, D_0)$	$D_0 \cup \{r\}$	$\omega_{\text{inv}}(\text{dom}(D_0)) \cup$ $\{\omega_{\text{inv}}(r)\}$	D_0	$\omega_{\text{oth}}(\text{dom}(D_0)) \cup$ $\{\omega_{\text{oth}}(r)\}$
$\overline{\text{rel}}(r, D_0)$	D_0	$\{\omega_{\text{oth}}(r)\} \cup$ $\omega_{\text{oth}}(\text{dom}(D_0))$	$D_0 \cup \{r\}$	$\{\omega_{\text{inv}}(r)\} \cup$ $\omega_{\text{inv}}(\text{dom}(D_0))$

Table 1: Interference events

- $\overline{\text{acq}}(r)$: For a closed resource r , the external process may acquire the resource if it is not local to the process being considered and therefore $\omega_{\text{oth}}(r)$ is marked.
- $\overline{\text{rel}}(r)$: For a closed resource r , the external process may release the resource if it is not local to the process being considered and therefore $\omega_{\text{oth}}(r)$ is marked.
- $\overline{\text{acq}}(r, D_0)$: For an open resource r with an invariant χ in Γ , if $D_0 \models \chi$ and D_0 is part of the current heap then ownership of the locations in the domain of D_0 is changed from being protected by the resource to being owned by the external process, *i.e.* un-marking $\omega_{\text{inv}}(\ell)$ and marking $\omega_{\text{oth}}(\ell)$ for each location $\ell \in \text{dom}(D_0)$. The ownership of r also changes, from $\omega_{\text{inv}}(r)$ being marked to $\omega_{\text{oth}}(r)$ being marked.
- $\overline{\text{rel}}(r, D_0)$: The corresponding release action.

Definition 4.4 (Interference net). The interference net for Γ has conditions \mathbf{S} , the state conditions, and \mathbf{W} , the ownership conditions. It has the following events:

- $\overline{\text{act}}(D_1, D_2)$ for all D_1 and D_2 forming partial functions with the same domain
- $\overline{\text{alloc}}(\ell, v, \ell', v')$ and $\overline{\text{dealloc}}(\ell, \ell', v')$ for all locations ℓ and ℓ' and values v and v'
- $\overline{\text{decl}}(r)$ and $\overline{\text{end}}(r)$ for all resources r
- $\overline{\text{acq}}(r)$ and $\overline{\text{rel}}(r)$ for all closed resources r
- $\overline{\text{acq}}(r, D_0)$ and $\overline{\text{rel}}(r, D_0)$ for all $r \in \text{dom}(\Gamma)$ and D_0 such that $D_0 \models \chi$, for χ the unique formula such that $r : \chi \in \Gamma$

We use the symbol u to range over interference events.

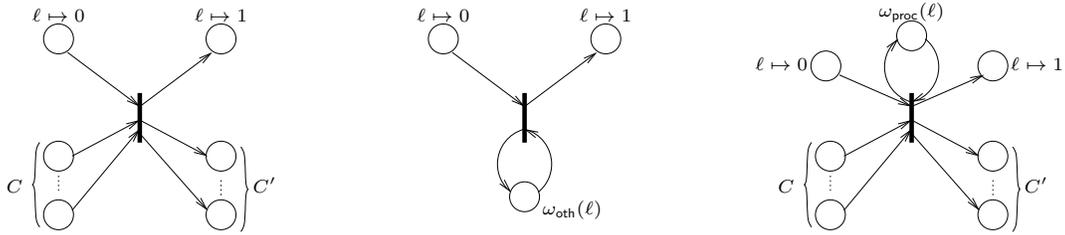
The interference events illustrate how the ownership of locations is *dynamic* and how this constrains the possible forms of interference. The rule for parallel composition requires that the behaviour of the process being reasoned about itself conforms to these constraints, allowing its action to be seen as interference when reasoning about the other process. This requirement may be captured by *synchronizing* the events of the process with those from the interference net in the following way:

- The process event $\text{act}_{(C,C')}(D, D')$ synchronizes with $\overline{\text{act}}(D, D')$
- The process event $\text{alloc}_{(C,C')}(\ell, v, \ell', v')$ synchronizes with $\overline{\text{alloc}}(\ell, v, \ell', v')$
- The process event $\text{dealloc}_{(C,C')}(\ell, \ell', v')$ synchronizes with $\overline{\text{dealloc}}(\ell, \ell', v')$
- The process event $\text{decl}_{(C,C')}(r)$ synchronizes with $\overline{\text{decl}}(r)$
- The process event $\text{end}_{(C,C')}(r)$ synchronizes with $\overline{\text{end}}(r)$
- The process event $\text{acq}_{(C,C')}(r)$ synchronizes with $\overline{\text{acq}}(r)$ for any closed resource r , *i.e.* for any $r \notin \text{dom}(\Gamma)$
- The process event $\text{rel}_{(C,C')}(r)$ synchronizes with $\overline{\text{rel}}(r)$ for any closed resource r
- If r is an open resource with $r : \chi \in \Gamma$, the process event $\text{acq}_{(C,C')}(r)$ synchronizes with every $\overline{\text{acq}}(r, D_0)$ such that $D_0 \models \chi$. Similarly, $\text{rel}_{(C,C')}(r)$ synchronizes with every $\overline{\text{rel}}(r, D_0)$ such that $D_0 \models \chi$.

Suppose that two events synchronize, e from the process and u from the interference net. The event u is the event that would fire in the net for the other parallel process to simulate the event e ; it is its dual. Let $e \cdot u$ be the event formed by taking the union of the pre- and postconditions of e and u , other than using $\omega_{\text{proc}}(\ell)$ in place of $\omega_{\text{oth}}(\ell)$, and similarly $\omega_{\text{proc}}(r)$ in place of $\omega_{\text{oth}}(r)$.

$$\begin{aligned} \bullet(e \cdot u) &\stackrel{\text{def}}{=} \{b \mid b \in \bullet e \cup \bullet u \text{ and } \nexists z.b = \omega_{\text{oth}}(z)\} \cup \{\omega_{\text{proc}}(z) \mid \omega_{\text{oth}}(z) \in \bullet u\} \\ (e \cdot u)^\bullet &\stackrel{\text{def}}{=} \{b \mid b \in e^\bullet \cup u^\bullet \text{ and } \nexists z.b = \omega_{\text{oth}}(z)\} \cup \{\omega_{\text{proc}}(z) \mid \omega_{\text{oth}}(z) \in u^\bullet\} \end{aligned}$$

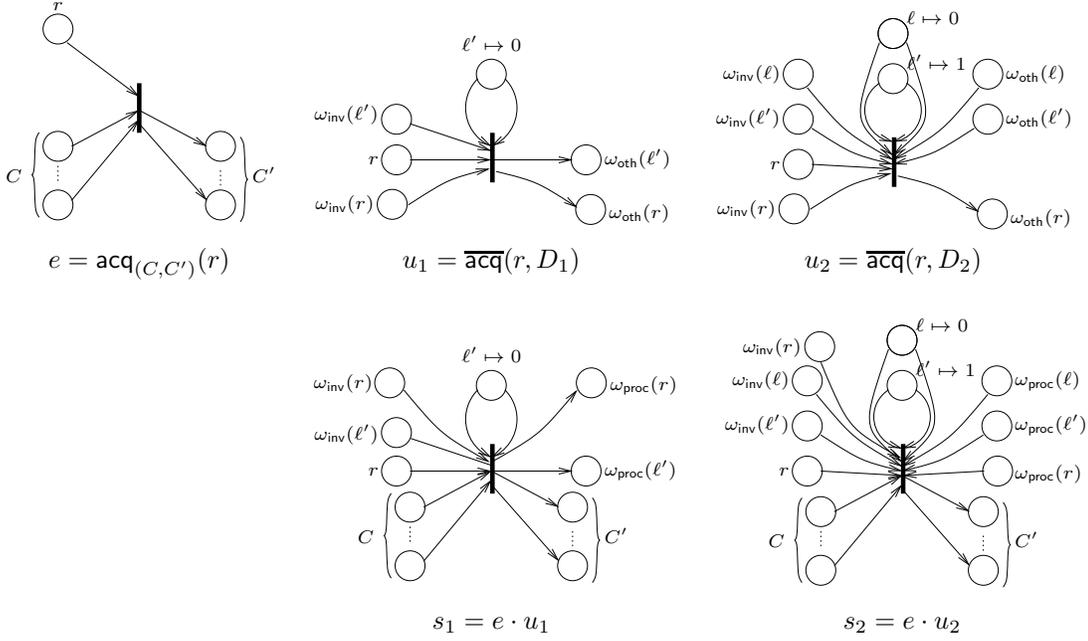
Example 4.5 (Synchronization of heap actions). Define the following events:



$$e = \text{act}_{(C,C')}(\{l \mapsto 0\}, \{l \mapsto 1\}) \quad u = \overline{\text{act}}(\{l \mapsto 0\}, \{l \mapsto 1\}) \quad e \cdot u$$

The event e is an event inside the process net, with pre-control conditions C and post-control conditions C' , that changes the value of l from 0 to 1. It synchronizes with only one event, u , which performs the corresponding interference action. For the event u to occur, the condition $\omega_{\text{oth}}(\ell)$ must be marked *i.e.* the location ℓ must be seen as owned by an ‘external’ process. The event formed by synchronizing e and u is $e \cdot u$, which requires the location ℓ to be owned by the current process for it to occur. \square

Example 4.6 (Synchronization of critical regions). Define the following events, where the event e is an event inside the process net, with pre-control conditions C and post-control conditions C' , that acquires the open resource r .



Recall the invariant $\ell' \mapsto 0 \vee (\ell' \mapsto 1 * \ell \mapsto 0)$ used above. There are two heaps, $D_1 = \{\ell' \mapsto 0\}$ and $D_2 = \{\ell' \mapsto 1, \ell \mapsto 0\}$ that satisfy this formula. There are correspondingly two interference events u_1 and u_2 that synchronize with e : the event u_1 acquires the resource r and transfers the ownership of ℓ' and r to the external process from the invariant, whereas the event u_2 acquires the resource r and transfers ownership of ℓ , ℓ' and r to the external process from the invariant. The event u_1 requires that the heap initially has value 0 at ℓ' ; the event u_2 requires that the heap initially has value 1 at ℓ' and 0 at ℓ . The synchronized events $e \cdot u_1$ and $e \cdot u_2$ are similar, transferring ownership from the invariant to the process being considered. \square

The semantics of judgements made using the rules of concurrent separation logic will consider a net $\mathcal{W}[[t]]_\Gamma$ with both interference events to represent external processes running and synchronized events to represent the process t .

Definition 4.7 (Ownership net). The ownership net for t in Γ , denoted $\mathcal{W}[[t]]_\Gamma$, is the net formed with the previous definitions of control conditions \mathbf{C} , state conditions \mathbf{S} and ownership conditions \mathbf{W} , and events:

- Every event u from the interference net for Γ , and
- Every event $e \cdot u$ where e is an event of $\mathcal{N}[[t]]$ and u from the interference net such that e and u synchronize.

We shall continue to use the symbol e to refer to any kind of event in ownership nets, but shall reserve the symbol s for those events known in particular to be synchronized events.

A consequence of the precision of invariants is that at most one of the synchronized events corresponding to an event in $\mathcal{N}[[t]]$ may be enabled in any marking of the ownership net $\mathcal{W}[[t]]_\Gamma$.

Lemma 4.8. *For any marking σ of state conditions, let (C, σ, W) and (C', σ, W') be consistent markings of the net $\mathcal{W}[[t]]_\Gamma$. For any event e in $\mathcal{N}[[t]]$ and any interference events u*

and u' in $\mathcal{W}[[t]]_\Gamma$, if $e \cdot u$ has concession in (C, σ, W) and $e \cdot u'$ has concession in (C', σ, W') then $u = u'$.

Proof. Straightforwardly seen to follow from precision by an analysis of the possible forms of the event e . \square

The occurrence of a synchronized event $e \cdot u$ in a marking (C, σ, W) of the net $\mathcal{W}[[t]]_\Gamma$ clearly gives rise to the occurrence of the event e in $\mathcal{N}[[t]]$. The earlier results describing the behaviour of $\mathcal{N}[[t]]$ in terms of the behaviour of the nets representing its subterms can therefore be applied to the net $\mathcal{W}[[t]]_\Gamma$.

Lemma 4.9. *If $M = (C, \sigma, W)$ and $M' = (C', \sigma', W')$ are markings of $\mathcal{W}[[t]]_\Gamma$ and $M \xrightarrow{e} M'$ then either e is an interference event and $C = C'$ or $e = e_1 \cdot u$ for an event e_1 of $\mathcal{N}[[t]]$ and an interference event u and $(C, \sigma) \xrightarrow{e_1} (C', \sigma')$ in $\mathcal{N}[[t]]$.*

Proof. The events of $\mathcal{W}[[t]]_\Gamma$ are, by definition, only interference events or synchronized events. If e is an interference event, $C = C'$ because ${}^c e = \emptyset$ and $e^c = \emptyset$. For a synchronized event $e_1 \cdot u$, observe that ${}^c(e_1 \cdot u) = {}^c e_1$ and that $(e_1 \cdot u)^c = e_1^c$, and similarly for ${}^l e_1$, ${}^r e_1$, ${}^n e_1$, e_1^l , e_1^r and e_1^n . The only cases where either ${}^d(e_1 \cdot u) \neq {}^d e_1$ or $(e_1 \cdot u)^d \neq e_1^d$ are acquisition or release of an open resource, but in these cases ${}^d e_1 = \emptyset = e_1^d$ and ${}^d(e_1 \cdot u) = (e_1 \cdot u)^d$. The result follows as a straightforward calculation. \square

The proof that consistent markings are preserved in the net $\mathcal{W}[[t]]_\Gamma$ is similar to that of Lemma 3.25; the additional requirements on the marking of ownership conditions are readily seen to be preserved by both interference and synchronized events.

Lemma 4.10 (Preservation of consistent markings). *For any closed term t , if in the net $\mathcal{W}[[t]]_\Gamma$ it is the case that $(\text{Ic}(t), \sigma_0, W_0) \xrightarrow{*} (C, \sigma, W)$ and (σ_0, W_0) is consistent then (σ, W) is consistent.* \square

The formulation of the ownership net permits a fundamental understanding of when a process acts in a way that cannot be seen as any form of interference; that is, when the process has violated its guarantees.

Definition 4.11 (Violating marking). Let (C, σ, W) be a consistent marking of $\mathcal{W}[[t]]_\Gamma$. We say that M is *violating* if there exists an event e of $\mathcal{N}[[t]]$ that has concession in marking (C, σ) but there is no event u from the interference net that synchronizes with e such that $e \cdot u$ has concession in (C, σ, W) .

We shall give two examples of violating markings. The first shall be an example of action on an unowned location, and the second shows how release of an open resource will cause a violation if the invariant is not restored.

Example 4.12. Let $(\{i\}, \sigma, W)$ be a consistent marking of $\mathcal{W}[[\ell] := 1]_\Gamma$ with $\ell \mapsto 0 \in \sigma$ and $\omega_{\text{oth}}(\ell) \in W$. The event $e = \text{act}_{(\{i\}, \{t\})}(\{\ell \mapsto 0\}, \{\ell \mapsto 1\})$ has concession in (C, σ) , but the only interference event that can synchronize with e is $u = \overline{\text{act}}(\{\ell \mapsto 0\}, \{\ell \mapsto 1\})$. We have $\omega_{\text{oth}}(\ell) \in {}^w u$ and therefore $\omega_{\text{proc}}(\ell) \in {}^w(e \cdot u)$, so the event $e \cdot u$ does not have concession in the marking (C, σ, W) which is therefore violating: the process acted on the unowned location ℓ' . \square

Example 4.13. Let r be an open resource with the invariant $\chi = \ell' \mapsto 0 \vee (\ell' \mapsto 1 * \ell \mapsto 0)$, and let (C, σ, W) be a consistent marking of $\mathcal{W}[[t]]_{\Gamma, r: \chi}$ with $\{\ell \mapsto 1, \ell' \mapsto 1\} \subseteq \sigma$ and $\omega_{\text{proc}}(\ell), \omega_{\text{proc}}(\ell') \in W$. Suppose further that the event $e = \text{rel}_{(C_1, C_2)}(r)$ has concession in (C, σ) in the net $\mathcal{N}[[t]]$. The only two interference events in $\mathcal{W}[[t]]_{\Gamma, r: \chi}$ that synchronize with e are

$$\begin{aligned} u_1 &= \overline{\text{rel}}(r, \{\ell' \mapsto 0\}) \\ u_2 &= \overline{\text{rel}}(r, \{\ell' \mapsto 1, \ell \mapsto 0\}), \end{aligned}$$

corresponding to the two ways in which χ can be satisfied. The invariant is not satisfied in the heap component of σ , so the preconditions of the two events

$$\begin{aligned} \bullet(e \cdot u_1) &= C \cup \{\ell' \mapsto 0, \omega_{\text{proc}}(\ell')\} \\ \bullet(e \cdot u_2) &= C \cup \{\ell' \mapsto 1, \ell \mapsto 0, \omega_{\text{proc}}(\ell'), \omega_{\text{proc}}(\ell)\} \end{aligned}$$

are not contained in the marking (C, σ, W) , which is therefore a violating marking because there was no part of the owned heap that satisfied the invariant yet the resource was released. \square

If no violating marking is ever encountered, the behaviour of $\mathcal{W}[[t]]_{\Gamma}$ encapsulates all that of $\mathcal{N}[[t]]$.

Lemma 4.14. *For any consistent marking (C, σ, W) of the net $\mathcal{W}[[t]]_{\Gamma}$ and any event $e \in \text{Ev}(t)$, if $(C, \sigma) \xrightarrow{e} (C', \sigma')$ in $\mathcal{N}[[t]]$ then either (C, σ, W) is violating or there exists a marking of ownership conditions W' and an interference event u that synchronizes with e such that $(C, \sigma, W) \xrightarrow{e \cdot u} (C', \sigma', W')$ in $\mathcal{W}[[t]]_{\Gamma}$.*

Proof. Immediate from the definition of violating marking and the fact that, for any e and u that synchronize and any state σ

$${}^c(e \cdot u) = {}^c e \quad (e \cdot u)^c = e^c \quad \sigma \setminus {}^s(e \cdot u) \cup (e \cdot u)^s = \sigma \setminus {}^s e \cup e^s$$

which is easily proved by inspection of the forms that $e \cdot u$ may take. \square

4.2. Soundness and validity. The rule for parallel composition permits the view that the ownership of the heap is initially split between the two processes, so that what one process owns is seen as owned by an external process by the other.

Definition 4.15 (Ownership split). Let W be a marking of ownership conditions. Markings of ownership conditions W_1 and W_2 form an *ownership split* of W if for all $z \in \text{Loc} \cup \text{Res}$:

- $\omega_{\text{oth}}(z) \in W$ iff $\omega_{\text{oth}}(z) \in W_1$ and $\omega_{\text{oth}}(z) \in W_2$,
- $\omega_{\text{inv}}(z) \in W$ iff $\omega_{\text{inv}}(z) \in W_1$ and $\omega_{\text{inv}}(z) \in W_2$, and
- $\omega_{\text{proc}}(z) \in W$ iff either $\omega_{\text{proc}}(z) \in W_1$ and $\omega_{\text{oth}}(z) \in W_2$,
or $\omega_{\text{proc}}(z) \in W_2$ and $\omega_{\text{oth}}(z) \in W_1$.

If W_1 and W_2 form an ownership split of W , then fewer locations and resources are owned by the process in W_1 than in W , and similarly for W_2 . As one would expect, a process can act in the same way without causing a violation if it owns more, and more interference can occur if the process owns less. This is the essence of the frame property referred to earlier.

Lemma 4.16. *Consider markings of the net $\mathcal{W}[[t]]_\Gamma$. Let W_1 and W_2 form an ownership split of W .*

- *For any synchronized event $s = e \cdot u$, if $(C, \sigma, W_1) \xrightarrow{s} (C', \sigma', W'_1)$ then there exist W' and W'_2 such that $(C, \sigma, W) \xrightarrow{s} (C', \sigma', W')$ and $(C, \sigma, W_2) \xrightarrow{u} (C, \sigma', W'_2)$, and furthermore W'_1 and W'_2 form an ownership split of W' .*
- *For any interference event u , if $(C, \sigma, W) \xrightarrow{u} (C, \sigma', W')$ then there exist W'_1 and W'_2 such that $(C, \sigma, W_1) \xrightarrow{u} (C, \sigma, W'_1)$ and $(C, \sigma, W_2) \xrightarrow{u} (C, \sigma', W'_2)$, and furthermore W'_1 and W'_2 form an ownership split of W' .*

Proof. A straightforward (but long) analysis of the possible forms of s and u . □

Following Brookes' lead, we are now able to prove the key lemma upon which the proof of soundness lies. The effect of this lemma is that the terminal states of parallel processes may be determined simply by observing the terminal markings of the net of each parallel process running in isolation if we split the ownership of the initial state correctly. For convenience, the lemma is stated without intimating the particular event that takes place on the net transition relation.

Lemma 4.17 (Parallel decomposition). *Let $M = (1:C_1 \cup 2:C_2, \sigma, W)$ be a consistent marking of the net $\mathcal{W}[[t_1 \parallel t_2]]_\Gamma$, and let W_1 and W_2 form an ownership split of W . The markings $M_1 = (C_1, \sigma, W_1)$ and $M_2 = (C_2, \sigma, W_2)$ are consistent, and furthermore:*

- *If the marking M is violating in $\mathcal{W}[[t_1 \parallel t_2]]_\Gamma$ then either M_1 is violating in $\mathcal{W}[[t_1]]_\Gamma$ or M_2 is violating in $\mathcal{W}[[t_2]]_\Gamma$.*
- *If neither M_1 nor M_2 is violating and $(1:C_1 \cup 2:C_2, \sigma, W) \longrightarrow (1:C'_1 \cup 2:C'_2, \sigma', W')$ in $\mathcal{W}[[t_1 \parallel t_2]]_\Gamma$ then there exist W'_1 and W'_2 forming an ownership split of W' such that $(C_1, \sigma, W_1) \longrightarrow (C'_1, \sigma', W'_1)$ in $\mathcal{W}[[t_1]]_\Gamma$ and $(C_2, \sigma, W_2) \longrightarrow (C'_2, \sigma', W'_2)$ in $\mathcal{W}[[t_2]]_\Gamma$.*

Proof. It is straightforward from Definition 4.3 to see that M_i is a consistent marking for both $i \in \{1, 2\}$.

- (1) Suppose that the marking M is violating in $\mathcal{W}[[t_1 \parallel t_2]]_\Gamma$. Without loss of generality, assume that this is because there exists an event $1:e_1$ of $\mathcal{N}[[t_1 \parallel t_2]]$ that has concession in marking $(1:C_1 \cup 2:C_2, \sigma)$ but there is no event interference event u such that $1:e_1$ synchronizes with u and $(1:e_1) \cdot u$ has concession in M . Assume, for contradiction, that the marking M_1 is non-violating in $\mathcal{W}[[t_1]]_\Gamma$. The event e_1 has concession in marking (C_1, σ) of $\mathcal{N}[[t_1]]$ by the first part of Lemma 3.6, so there must exist u_1 an interference event of $\mathcal{W}[[t_1]]_\Gamma$ such that $e_1 \cdot u_1$ has concession in M_1 . The interference events of $\mathcal{W}[[t_1]]_\Gamma$ are precisely the interference events of $\mathcal{W}[[t_1 \parallel t_2]]_\Gamma$ and the tagging of control conditions has no effect on whether events may synchronize, so the event $(1:e_1) \cdot u_1$ is in $\mathcal{W}[[t_1 \parallel t_2]]_\Gamma$. From Lemmas 4.16 and 3.6, the event $1:e_1 \cdot u_1$ has concession in marking M , which is therefore not violating — a contradiction.
- (2) It is a straightforward consequence of Lemma 4.16 that the second property holds if the transition $(1:C_1 \cup 2:C_2, \sigma, W) \longrightarrow (1:C'_1 \cup 2:C'_2, \sigma', W')$ is induced by the occurrence of an interference event. Suppose instead that it is induced by a synchronized event.

Without loss of generality, suppose that in $\mathcal{W}[[t_1 \parallel t_2]]_\Gamma$ we have $M \xrightarrow{(1:e_1) \cdot u} M'$ for $M' = (1:C'_1 \cup 2:C'_2, \sigma', W')$, for some event e_1 in $\mathcal{N}[[t_1]]$. We shall show that $M_1 \xrightarrow{e_1 \cdot u} (C_1, \sigma', W'_1)$ in $\mathcal{W}[[t_1]]_\Gamma$ and $M_2 \xrightarrow{u} (C'_2, \sigma', W'_2)$ in $\mathcal{W}[[t_2]]_\Gamma$ for some W'_1, W'_2 such that W'_1 and W'_2

form an ownership split of W' . Since we have $M \xrightarrow{(1:e_1)\cdot u} M'$ in $\mathcal{W}[[t_1 \parallel t_2]]_\Gamma$, it is easy to see that we have $(1:C_1 \cup 2:C_2, \sigma) \xrightarrow{1:e_1} (1:C'_1 \cup 2:C'_2, \sigma')$ in $\mathcal{N}[[t_1 \parallel t_2]]$ and $C_2 = C'_2$. Hence in $\mathcal{N}[[t_1]]$ we have $(C_1, \sigma) \xrightarrow{e_1} (C'_1, \sigma')$. By assumption, the marking (C_1, σ, W_1) is not a violating marking of $\mathcal{W}[[t_1]]_\Gamma$, so there exists an interference event u_1 that synchronizes with e_1 such that $(C_1, \sigma, W_1) \xrightarrow{e_1 \cdot u_1} (C'_1, \sigma', W''_1)$ for some W''_1 in $\mathcal{W}[[t_1]]_\Gamma$, so in $\mathcal{W}[[t_1 \parallel t_2]]_\Gamma$ we therefore have $(1:C_1 \cup 2:C_2, \sigma, W_1) \xrightarrow{(1:e_1)\cdot u_1} (1:C'_1 \cup 2:C_2, \sigma', W''_1)$. By Lemma 4.8, we have $u_1 = u$ and therefore $W''_1 = W_1$ because the occurrence of an event in a marking yields a unique marking. Now, by Lemma 4.16 there exist W'' and W'_2 such that W'_1 and W'_2 form an ownership split of W'' and $(1:C_1 \cup 2:C_2, \sigma, W) \xrightarrow{(1:e_1)\cdot u} (1:C'_1 \cup 2:C_2, \sigma', W'')$ and $(1:C_1 \cup 2:C_2, \sigma, W_2) \xrightarrow{u} (1:C_1 \cup 2:C_2, \sigma', W'_2)$. The occurrence of an event in a marking leads to a unique marking, so $W'' = W'$. It is easy to see that $(C_1, \sigma, W_1) \xrightarrow{e_1 \cdot u} (C'_1, \sigma', W'_1)$ in $\mathcal{W}[[t_1]]_\Gamma$ and that $(C_2, \sigma, W_2) \xrightarrow{u} (C_2, \sigma, W'_2)$ in $\mathcal{W}[[t_2]]_\Gamma$, so the proof is complete. \square

The ownership semantics described above has been carefully defined to explicitly take into account the intuitions behind the rule for parallel composition, resulting in the short proof of the parallel decomposition lemma above. The remaining complexity in the proof of soundness lies in the rule for establishing an invariant associated with a resource:

$$\text{(L-RES)} : \frac{\Gamma, r : \chi \vdash \{\varphi\} [r/w]t \{\psi\}}{\Gamma \vdash \{\varphi * \chi\} \text{ resource } w \text{ do } t \text{ od } \{\psi * \chi\}} \left(\begin{array}{l} \chi \text{ precise} \\ r \notin \text{dom}(\Gamma) \end{array} \right)$$

It is quite easy to see why this rule follows the intuitive semantics for judgements presented above: Any run of the net $\mathcal{W}[[\text{resource } w \text{ do } t \text{ od}]]_\Gamma$ to a terminal marking from a state with the heap owned by the process initially satisfying $\varphi * \chi$ can be seen, in conjunction with Lemma 3.21, as consisting first of an event that declares a fresh resource r current, then a run of $\mathcal{W}[[[r/w]t]]_\Gamma$, followed by an event that makes r non-current. The run of $\mathcal{W}[[[r/w]t]]_\Gamma$ from a state where the part of the heap that the process owns satisfies $\varphi * \chi$ is simulated by a run of $\mathcal{W}[[[r/w]t]]_{\Gamma, r : \chi}$ along which the locations satisfying χ are owned by the invariant χ in an environment where r is an open resource. In particular, the run obtained has no interference on the resource r or the locations that it protects and r is available in the terminal state of the run. Assuming the validity of the judgement $\Gamma, r : \chi \vdash \{\varphi\} [r/w]t \{\psi\}$, the resulting state owned by the process is therefore seen to satisfy the formula $\varphi * \chi$. Similarly, if there were a reachable marking in $\mathcal{W}[[\text{resource } w \text{ do } t \text{ od}]]_\Gamma$ where the process accesses a location or resource that it does not own would result in there being a reachable marking in $\mathcal{W}[[[r/w]t]]_{\Gamma, r : \chi}$ where the process accesses an unowned location or resource. The more formal presentation of this intuition follows.

We shall begin by explicitly characterizing the runs of the net $\mathcal{W}[[\text{resource } w \text{ do } t_0 \text{ od}]]_\Gamma$. The result is again a little technical, as is the following lemma, Lemma 4.21; they are used in the proof of soundness of the rule (L-RES). The reader may wish to pass through this result and Lemma 4.21 and only take note of the following definitions of $\text{inv}(\Gamma, R)$ and $D \downarrow_w \text{proc}$, $D \downarrow_w \text{inv}$ and $D \downarrow_w \text{oth}$.

Lemma 4.18. *Suppose that σ_0 and W_0 form a consistent marking of state and ownership conditions and let $t \equiv \text{resource } w \text{ do } t_0 \text{ od}$. For a resource r , define the synchronized*

events

$$\begin{aligned} s_r &= \text{decl}_{(\{i\}, r: \text{Ic}([r/w]t_0))}(r) \cdot \overline{\text{decl}}(r) \\ s'_r &= \text{end}_{(r: \text{Tc}([r/w]t_0), \{t\})}(r) \cdot \overline{\text{end}}(r) \end{aligned}$$

If in the net $\mathcal{W}[[t]]_\Gamma$ we have $\pi : (\text{Ic}(t), \sigma_0, W_0) \longrightarrow^* (C, \sigma, W)$ then either:

- $C = \text{Ic}(t)$ and π consists only of interference events, or
- there exist $r, C', \sigma', W', \pi_0$ and π_1 such that π_0 comprises only interference events, $C = r:C'$ and

$$\pi = \pi_0 \cdot s_r \cdot (r:\pi_1)$$

and

$$\begin{aligned} \pi_0 \cdot s_r : (\text{Ic}(t), \sigma_0, W_0) &\longrightarrow^* (r:\text{Ic}([r/w]t_0), \sigma', W') && \text{in } \mathcal{W}[[t]]_\Gamma \text{ and} \\ \pi_1 : (\text{Ic}([r/w]t_0), \sigma', W') &\longrightarrow^* (C', \sigma, W) && \text{in } \mathcal{W}[[r/w]t_0]_\Gamma, \text{ or} \end{aligned}$$

- $C = \text{Tc}(t)$ and there exist $r, \sigma', \sigma'', W', W'', \pi_0, \pi_1, \pi_2$ such that π_0 and π_2 comprise only interference events,

$$\pi = \pi_0 \cdot s_r \cdot (r:\pi_1) \cdot s'_r \cdot \pi_2,$$

and

$$\begin{aligned} \pi_0 \cdot s_r : (\text{Ic}(t), \sigma_0, W_0) &\longrightarrow^* (r:\text{Ic}([r/w]t_0), \sigma', W') && \text{in } \mathcal{W}[[t]]_\Gamma, \\ \pi_1 : (\text{Ic}([r/w]t_0), \sigma', W') &\longrightarrow^* (\text{Tc}([r/w]t_0), \sigma'', W'') && \text{in } \mathcal{W}[[r/w]t_0]_\Gamma, \text{ and} \\ s'_r \cdot \pi_2 : (r:\text{Tc}([r/w]t_0), \sigma'', W'') &\longrightarrow^* (\text{Tc}(t), \sigma, W) && \text{in } \mathcal{W}[[t]]_\Gamma. \end{aligned}$$

Proof. Readily seen to be a consequence of Lemmas 3.23, 3.21, 3.10 and 4.9. \square

It can be shown, as a consequence of the preceding lemma, that during the run of the net following the declaration event, the resource r chosen for w is owned by the process until it is made non-current at the end of the variable w 's scope.

Lemma 4.19. *Let $t \equiv \text{resource } w \text{ do } t_0 \text{ od}$. If $(r:C_0, \sigma, W)$ is reachable from $(\text{Ic}(t), \sigma_0, W_0)$, which is a consistent marking of $\mathcal{W}[[t]]_\Gamma$, then $\omega_{\text{proc}}(r) \in W$. \square*

We write $\text{inv}(\Gamma, R)$ for the formula $\chi_1 * \dots * \chi_n$ formed as the separating conjunction of the invariants of all the *available*, according to R , open resources. It is defined by induction on the size of the domain of Γ :

$$\begin{aligned} \text{inv}(\emptyset, R) &\stackrel{\text{def}}{=} \text{empty} \\ \text{inv}((\Gamma, r:\chi), R) &\stackrel{\text{def}}{=} \begin{cases} \text{inv}(\Gamma, R), & \text{if } r \notin R \\ \chi * \text{inv}(\Gamma, R), & \text{if } r \in R. \end{cases} \end{aligned}$$

Define the notations

$$\begin{aligned} D \upharpoonright_w \text{proc} &\stackrel{\text{def}}{=} \{\ell \mapsto v \in D \mid \omega_{\text{proc}}(\ell) \in W\} \\ D \upharpoonright_w \text{inv} &\stackrel{\text{def}}{=} \{\ell \mapsto v \in D \mid \omega_{\text{inv}}(\ell) \in W\} \\ D \upharpoonright_w \text{oth} &\stackrel{\text{def}}{=} \{\ell \mapsto v \in D \mid \omega_{\text{oth}}(\ell) \in W\} \end{aligned}$$

to represent the heap at locations owned by the process, invariants and other processes, respectively. In any state that we consider, we would expect $D \upharpoonright_w \text{inv} \models \text{inv}(\Gamma, R)$. A marking of the net $\mathcal{W}[[t]]_\Gamma$ can be converted to a marking of $\mathcal{W}[[t]]_{\Gamma, r:\chi}$ by, if r is available, regarding ownership of the locations satisfying the invariant χ as being owned by the invariant rather than by the process.

Definition 4.20. Suppose that χ is a precise heap formula. Let $M = (C, (D, L, R, N), W)$ be a consistent marking of $\mathcal{W}[[t]]_\Gamma$ such that if $r \in R$ then there exists (necessarily unique) $D_0 \subseteq D \upharpoonright_W \text{proc}$ such that $D_0 \models \chi$. Define the projection of M into the net $\mathcal{W}[[t]]_{\Gamma, r: \chi}$ to be

$$\pi_r^\chi(M) \stackrel{\text{def}}{=} (C, (D, L, R, N), W'),$$

where:

- if $r \notin R$: $W' = W$
- if $r \in R$: Let $D_0 \subseteq D$ be such that $D_0 \models \chi$.

$$\begin{aligned} W' &= \{\omega_{\text{oth}}(\ell) \mid \omega_{\text{oth}}(\ell) \in W\} \\ &\cup \{\omega_{\text{oth}}(r') \mid \omega_{\text{oth}}(r') \in W\} \\ &\cup \{\omega_{\text{inv}}(\ell) \mid \omega_{\text{inv}}(\ell) \in W \text{ or } \ell \in \text{dom}(D_0)\} \\ &\cup \{\omega_{\text{inv}}(r') \mid \omega_{\text{inv}}(r') \in W \text{ or } r' = r\} \\ &\cup \{\omega_{\text{proc}}(\ell) \mid \omega_{\text{proc}}(\ell) \in W \text{ and } \ell \notin \text{dom}(D_0)\} \\ &\cup \{\omega_{\text{proc}}(r') \mid \omega_{\text{proc}}(r') \in W \text{ and } r' \neq r\} \end{aligned}$$

It is clear that if M is a consistent marking of $\mathcal{W}[[t]]_\Gamma$ then $\pi_r^\chi(M)$ is a consistent marking of $\mathcal{W}[[t]]_{\Gamma, r: \chi}$. The key lemma representing the account above, that behaviour in the net where a resource is closed is simulated by the net where the resource is open, is now stated, though we shall not show its proof here.

Lemma 4.21. *Let r be a resource such that $r \notin \text{dom}(\Gamma)$ and let χ be a precise heap logic formula. Let $M = (C, (D, L, R, N), W)$ be a consistent marking of $\mathcal{W}[[t]]_\Gamma$ such that:*

- $\omega_{\text{proc}}(r) \in W$,
- $D \upharpoonright_W \text{inv} \models \text{inv}(\Gamma, R)$, and
- if $r \in R$ then there exists $D_0 \subseteq D \upharpoonright_W \text{proc}$ such that $D_0 \models \chi$.

Then

- (1) *If M is a violating marking in $\mathcal{W}[[t]]_\Gamma$ then $\pi_r^\chi(M)$ is a violating marking in $\mathcal{W}[[t]]_{\Gamma, r: \chi}$.*
- (2) *For any event u of $\mathcal{W}[[t]]_\Gamma$ that is an interference event, if M is not a violating marking and $M \xrightarrow{u} M'$ where $M' = (C', (D', L', R', N'), W')$ and $\omega_{\text{proc}}(r) \in W'$ then:*
 - $\pi_r^\chi(M) \xrightarrow{u} \pi_r^\chi(M')$ in $\mathcal{W}[[t]]_{\Gamma, r: \chi}$ and:
 - $D' \upharpoonright_{W'} \text{inv} \models \text{inv}(\Gamma, R')$
 - if $r \in R'$ then there exists $D_0 \subseteq D' \upharpoonright_{W'} \text{proc}$ such that $D_0 \models \chi$.
- (3) *For any synchronized event $s = e_1 \cdot u$ of $\mathcal{W}[[t]]_\Gamma$, if M is not a violating marking and $M \xrightarrow{s} M'$ where $M' = (C', (D', L', R', N'), W')$ and $\omega_{\text{proc}}(r) \in W'$ then either:*
 - $\pi_r^\chi(M)$ is violating in $\mathcal{W}[[t]]_{\Gamma, r: \chi}$, or
 - there exists u' such that $\pi_r^\chi(M) \xrightarrow{e_1 \cdot u'} \pi_r^\chi(M')$ in $\mathcal{W}[[t]]_{\Gamma, r: \chi}$ and:
 - $D' \upharpoonright_{W'} \text{inv} \models \text{inv}(\Gamma, R')$
 - if $r \in R'$ then there exists $D_0 \subseteq D' \upharpoonright_{W'} \text{proc}$ such that $D_0 \models \chi$. □

We shall say that a state σ with an ownership marking W satisfies the formula φ and the invariants in Γ if the heap restricted to the owned locations satisfies φ and the invariants are met for all the available resources. The rest of the heap, seen as owned by external processes, is unconstrained.

Definition 4.22. A marking (C, σ, W) of $\mathcal{W}[[t]]_\Gamma$ satisfies φ in Γ if:

- the marking (C, σ, W) is consistent,

- $D \upharpoonright_w \text{proc} \models \varphi$, and
 - $D \upharpoonright_w \text{inv} \models \text{inv}(\Gamma, R)$,
- where $\sigma = (D, L, R, N)$.

We now attach a notion of validity to judgements $\Gamma \vdash \{\varphi\} t \{\psi\}$. It shall assert that no violating marking is ever reached and that whenever the process t runs to completion from a state where the part of the heap that it owns satisfies φ then the part of the resulting heap that it owns satisfies ψ .

Definition 4.23 (Validity). Let t be a closed term. Define $\Gamma \models \{\varphi\} t \{\psi\}$ if, for any σ and W such that the marking $(\text{Ic}(t), \sigma, W)$ satisfies φ in Γ :

- any marking reachable in $\mathcal{W}[[t]]_\Gamma$ from $(\text{Ic}(t), \sigma, W)$ is non-violating, and
- for any σ' and W' , if the marking $(\text{Tc}(t), \sigma', W')$ is reachable in $\mathcal{W}[[t]]_\Gamma$ from $(\text{Ic}(t), \sigma, W)$ then $(\text{Tc}(t), \sigma', W')$ satisfies ψ in Γ .

It is useful to note that the occurrence of an interference event does not affect whether a marking satisfies φ in Γ or whether it is violating. Consequently, when considering validity it is unnecessary to account for runs of the net $\mathcal{W}[[t]]_\Gamma$ that start or end with an interference event.

Lemma 4.24. *Let M be a consistent marking of $\mathcal{W}[[t]]_\Gamma$ that satisfies φ in Γ and is non-violating. If u is an interference event and $M \xrightarrow{u} M'$ then M' satisfies φ in Γ and is non-violating.*

Proof. Straightforward from the definition of satisfaction of φ in Γ by considering the possible forms of u . \square

In the rule (L-RES) which allows invariants to be established for resources, only one resource is considered for substitution for the variable. The following lemma shows that this is sufficient; the semantics of judgements is unaffected by the choice of resource.

Lemma 4.25. *For any resources r, r' such that $r, r' \notin \text{dom}(\Gamma)$ and any term t with $\text{fv}(t) \subseteq \{w\}$ and $\text{res}(t) \subseteq \text{dom}(\Gamma)$,*

$$\Gamma, r : \chi \models \{\varphi\} [r/w]t \{\psi\} \text{ iff } \Gamma, r' : \chi \models \{\varphi\} [r'/w]t \{\psi\}.$$

Proof. The net $\mathcal{W}[[r/w]t]_{\Gamma, r : \chi}$ is clearly isomorphic to $\mathcal{W}[[r'/w]t]_{\Gamma, r' : \chi}$ through interchanging the conditions

$$\begin{aligned} r \leftrightarrow r' & & \text{curr}(r) \leftrightarrow \text{curr}(r') \\ \omega_{\text{proc}}(r) \leftrightarrow \omega_{\text{proc}}(r') & & \omega_{\text{inv}}(r) \leftrightarrow \omega_{\text{inv}}(r') & & \omega_{\text{oth}}(r) \leftrightarrow \omega_{\text{oth}}(r'). \end{aligned}$$

The result follows from the definition of validity being insensitive to such permutations. \square

We are now in a position where we the rules of concurrent separation logic can be proved sound. Only two important cases of the proof shall be presented here; full details will be available in the first author's PhD thesis.

Theorem 4.26 (Soundness). *For any closed term t , if $\Gamma \vdash \{\varphi\} t \{\psi\}$ then $\Gamma \models \{\varphi\} t \{\psi\}$.*

Proof. By rule induction on the judgement $\Gamma \vdash \{\varphi\} t \{\psi\}$. Note that, due to Lemma 4.24, we shall only consider runs of $\mathcal{W}[[t]]_\Gamma$ that do not start or end with an interference event.

(L-PAR): Suppose that we have $\Gamma \vdash \{\varphi_1 * \varphi_2\} t_1 \parallel t_2 \{\psi_1 * \psi_2\}$ because $\Gamma \vdash \{\varphi_1\} t_1 \{\psi_1\}$ and $\Gamma \vdash \{\varphi_2\} t_2 \{\psi_2\}$. Assume that marking $M = (\text{Ic}(t_1 \parallel t_2), \sigma, W)$ satisfies $\varphi_1 * \varphi_2$ in Γ .

It can be seen from the definitions that there exist W_1 and W_2 forming an ownership split of W such that $(\text{Ic}(t_1), \sigma, W_1)$ is a marking of $\mathcal{W}[[t_1]]_\Gamma$ that satisfies φ_1 in Γ and $(\text{Ic}(t_2), \sigma, W_2)$ satisfies φ_2 in Γ . Let marking $M' = (C', \sigma', W')$ be reachable from M ; a simple induction on the length of path to M using Lemma 4.17 and Lemma 3.17 shows that there exist C'_1, C'_2, W'_1 and W'_2 such that $C' = 1:C'_1 \cup 2:C'_2$ and W'_1 and W'_2 form an ownership split of W' . Furthermore, the marking (C'_1, σ', W'_1) is reachable from $(\text{Ic}(t_1), \sigma, W_1)$ in $\mathcal{W}[[t_1]]_\Gamma$ and (C'_2, σ', W'_2) is reachable from $(\text{Ic}(t_2), \sigma, W_2)$ in $\mathcal{W}[[t_2]]_\Gamma$.

Suppose that the marking M' is violating. Using Lemma 4.17, it follows that either (C'_1, σ', W'_1) or (C'_2, σ', W'_2) is a violating marking. This contradicts either the induction hypothesis for $\Gamma \vdash \{\varphi_1\} t_1 \{\psi_1\}$ or the induction hypothesis for $\Gamma \vdash \{\varphi_2\} t_2 \{\psi_2\}$, so M' cannot be violating.

Now suppose that the marking M' is terminal: we have $C'_1 = \text{Tc}(t_1)$ and $C'_2 = \text{Tc}(t_2)$. From the induction hypotheses, we obtain that (C'_1, σ', W'_1) satisfies ψ_1 in Γ and that (C'_2, σ', W'_2) satisfies ψ_2 in Γ . It is easy to see from the definition of ownership split that therefore (C', σ', W') satisfies $\psi_1 * \psi_2$ in Γ .

(L-RES): Let $t \equiv \text{resource } w \text{ do } t_0 \text{ od}$. Suppose that $\Gamma \vdash \{\varphi * \chi\} t \{\psi * \chi\}$ because $\Gamma, r_0 : \chi \vdash \{\varphi\} [r_0/w]t_0 \{\psi\}$ for some $r_0 \notin \text{dom}(\Gamma)$. Assume that the marking $M = (\text{Ic}(t), \sigma, W)$ satisfies $\varphi * \chi$ in Γ , and let $M' = (C', \sigma', W')$ be reachable from M in $\mathcal{W}[[t]]_\Gamma$. According to Lemma 4.18, there are three cases to consider for the marking M' .

- The first case has $M = M'$ (we need not consider runs starting with an interference event according to Lemma 4.24). Since $\text{Ic}(t) \neq \text{Tc}(t)$, all that we must show is that M is non-violating. Using Lemma 3.21, we can infer that the only events with concession in the marking $(\text{Ic}(t), \sigma)$ of $\mathcal{N}[[t]]$ are equal to $\text{decl}_{(\text{Ic}(t), r: \text{Ic}([r/w]t_0))}(r)$ for some $r \in \text{Res}$ such that $\text{curr}(r) \notin \sigma$. The marking M is assumed to be consistent, so for each such r we have $\omega_{\text{proc}}(r) \notin W$ and hence the synchronized event $\text{decl}_{(\text{Ic}(t), r: \text{Ic}([r/w]t_0))}(r) \cdot \overline{\text{decl}}(r)$ has concession in M . The marking M' cannot therefore be violating.
- Secondly, there exists a resource r , markings σ_0, W_0, C_1 and a path π_1 such that $C' = r:C_1$ and

$$\begin{array}{lcl} s_r : (C, \sigma, W) & \longrightarrow & (r:\text{Ic}([r/w]t_0), \sigma_0, W_0) \text{ in } \mathcal{W}[[t]]_\Gamma \\ \pi_1 : (\text{Ic}([r/w]t_0), \sigma_0, W_0) & \longrightarrow^* & (C_1, \sigma', W') \text{ in } \mathcal{W}[[r/w]t_0]_\Gamma, \end{array}$$

where $s_r = \text{decl}_{(\text{Ic}(t), r: \text{Ic}([r/w]t_0))}(r) \cdot \overline{\text{decl}}(r)$. The marking (C', σ', W') cannot be a terminal marking of the net $\mathcal{W}[[t]]_\Gamma$, so all that we must show is that it is non-violating. We have $r, \text{curr}(r) \in \sigma_0$ and $\omega_{\text{proc}}(r) \in W_0$ since they are in the postconditions of s_r . A simple induction on the length of π using Lemmas 4.19 and 4.21 informs that $\pi_r^\chi(C_1, \sigma', W')$ is reachable from $\pi_r^\chi(\text{Ic}([r/w]t_0), \sigma_0, W_0)$ in $\mathcal{W}[[r/w]t_0]_{\Gamma, r: \chi}$. We have $\text{curr}(r) \notin \sigma$ because the event s_r has concession in M , so $r \notin \text{dom}(\Gamma)$ because the marking M is consistent. Since $\text{res}(t) \subseteq \text{dom}(\Gamma)$ by Lemma 4.2, we may use Lemma 4.25 in conjunction with the induction hypothesis to obtain $\Gamma, r: \chi \models \{\varphi\} [r/w]t_0 \{\psi\}$. It is an easy calculation to show that $\pi_r^\chi(\text{Ic}([r/w]t_0), \sigma_0, W_0)$ satisfies φ in $\Gamma, r: \chi$, so the marking $\pi_r^\chi(C_1, \sigma', W')$ is non-violating. By Lemma 4.21, the marking (C_1, σ', W') of $\mathcal{W}[[r/w]t_0]_\Gamma$ is therefore non-violating. According to Lemma 3.21, there are two possible ways in which the marking (C', σ', W') of $\mathcal{W}[[t]]_\Gamma$ might be violating. Firstly, there might exist an event e of $\mathcal{N}[[r/w]t_0]$ that has concession in the marking (C_1, σ') but there is no interference event u that synchronizes with e such that $e \cdot u$ has concession in the marking (C_1, σ', W') . We have shown, however, that this is not the case since the marking (C_1, σ', W') is non-violating. Alternatively, the event $e'_r = \text{end}_{(r: \text{Tc}([r/w]t_0), \text{Tc}(t))}(r)$ might have concession in

the marking (C', σ') of $\mathcal{N}[[t]]$ but the event $s'_r = e'_r \cdot \overline{\text{end}}(r)$ might not have concession in (C', σ', W') ; that is, $\omega_{\text{proc}}(r) \notin W'$. However, we have $\omega_{\text{proc}}(r) \in \sigma_0$ so by applying Lemma 4.19 along path π_1 we obtain $\omega_{\text{proc}}(r) \in W'$. So the event s'_r has concession in the marking, which is therefore not violating.

- The final case is where $C' = \text{Tc}(t)$ and there exist $\sigma_0, \sigma_1, W_0, W_1$ and a path π_1 such that

$$\begin{aligned} s_r &: (\text{Ic}(t), \sigma, W) && \longrightarrow && (r:\text{Ic}([r/w]t), \sigma_0, W_0) && \text{ in } \mathcal{W}[[t]]_\Gamma \\ \pi_1 &: (\text{Ic}([r/w]t_0), \sigma_0, W_0) && \longrightarrow^* && (\text{Tc}([r/w]t_0), \sigma_1, W_1) && \text{ in } \mathcal{W}[[r/w]t_0]_\Gamma \\ s'_r &: (r:\text{Tc}([r/w]t_0), \sigma_1, W_1) && \longrightarrow && (\text{Tc}(t), \sigma', W') = M' && \text{ in } \mathcal{W}[[t]]_\Gamma, \end{aligned}$$

where $s_r = \text{decl}_{(\text{Ic}(t), r:\text{Ic}([r/w]t_0))}(r) \cdot \overline{\text{decl}}(r)$. The marking M' is readily seen to be non-violating since no event of $\mathcal{N}[[t]]$ has concession if the marking of control conditions is $\text{Tc}(t)$. All that remains is to show that M' satisfies ψ in Γ . As in the previous case, we have $r, \text{curr}(r) \in \sigma_0$ and $\omega_{\text{proc}}(r) \in W_0$ and $r \notin \text{dom}(\Gamma)$. It is easily seen that the marking $\pi_r^\chi(\text{Ic}([r/w]t_0), \sigma_0, W_0)$ of $\mathcal{W}[[r/w]t_0]_{\Gamma, r:\chi}$ satisfies φ in $\Gamma, r:\chi$. A simple induction on the length of the path π using Lemmas 4.19 and 4.21 shows that the marking $\pi_r^\chi(\text{Tc}([r/w]t_0), \sigma_1, W_1)$ is reachable in $\mathcal{W}[[r/w]t_0]_{\Gamma, r:\chi}$ from $\pi_r^\chi(\text{Ic}([r/w]t_0), \sigma_0, W_0)$. Using Lemmas 4.25 and 4.2, from the induction hypothesis $\Gamma, r_0:\chi \models \{\varphi\} [r_0/w]t_0 \{\psi\}$, the marking $\pi_r^\chi(\text{Tc}([r/w]t_0), \sigma_1, W_1)$ satisfies ψ in $\Gamma, r:\chi$. We have $r \in \sigma_1$ since the event s'_r has concession in the marking $(r:\text{Tc}([r/w]t_0), \sigma_1, W_1)$, so the marking $(\text{Tc}([r/w]t_0), \sigma_1, W_1)$ satisfies $\psi * \chi$ in Γ , from which it is easily seen that $(\text{Tc}(t), \sigma', W')$ also satisfies $\psi * \chi$ in Γ . \square

The following result connects the definition of validity to the execution of processes without interference or ownership.

Corollary 4.27 (Connection). *Let t be a closed term with $\text{res}(t) = \emptyset$ and let $\sigma = (D, L, \emptyset, \emptyset)$ be a consistent marking of state conditions for which $D \models \varphi$. If $\emptyset \models \{\varphi\} t \{\psi\}$ then whenever a terminal marking $(\text{Tc}(t), \sigma')$ is reachable from $(\text{Ic}(t), \sigma)$ in $\mathcal{N}[[t]]$, the resulting heap D' satisfies ψ , where $\sigma' = (D', L', R', N')$.*

Proof. A consequence of soundness and Lemma 4.14. \square

4.3. Fault. It can be seen that the rules of concurrent separation logic ensure that processes, running from suitable initial states, only access current locations. The syntax of the language ensures that processes only access current resources and that they are never blocked when releasing a resource through it already being available. We shall now demonstrate that processes avoid such ‘faults’, in which we shall say that an event e is *control-enabled* in a marking C of control conditions if there exists a marking C' such that $C \xrightarrow{e}_C C'$.

Definition 4.28 (Fault). There is a *fault* in a marking $M = (C, \sigma)$ of the net $\mathcal{N}[[t]]$ if there exists a control-enabled event e in $\mathcal{N}[[t]]$ with ${}^c e = C_1$ and $e^c = C_2$ for some C_1, C_2 such that either:

- (1) there exist D, D' such that $e = \text{act}_{(C_1, C_2)}(D, D')$ and there exists $\ell \in \text{dom}(D)$ with $\text{curr}(\ell) \notin \sigma$,
- (2) there exist ℓ, v, ℓ', v' such that $e = \text{alloc}_{(C_1, C_2)}(\ell, v, \ell', v')$ and $\text{curr}(\ell) \notin \sigma$,
- (3) there exist ℓ, v, ℓ', v' such that $e = \text{dealloc}_{(C_1, C_2)}(\ell, v, \ell')$ and either $\text{curr}(\ell) \notin \sigma$ or $\text{curr}(\ell') \notin \sigma$,
- (4) there exists r such that either $e = \text{acq}_{(C_1, C_2)}(r)$ or $e = \text{rel}_{(C_1, C_2)}(r)$ and $\text{curr}(r) \notin \sigma$, or

(5) there exists r such that $e = \text{rel}_{(C_1, C_2)}(r)$ and $r \in \sigma$.

This definition also applies to markings (C, σ, W) of $\mathcal{W}[[t]]_\Gamma$ in the by ignoring the marking of ownership conditions W and considering synchronized events $e \cdot u$.

Theorem 4.29 (Fault avoidance). *Suppose that $\Gamma \vdash \{\varphi\} t \{\psi\}$ and that the marking $(\text{Ic}(t), \sigma_0, W_0)$ satisfies φ in Γ . If (C, σ, W) is reachable from $(\text{Ic}(t), \sigma_0, W_0)$ then there is not a fault in (C, σ, W) .*

Proof. By rule induction on the judgement $\Gamma \vdash \{\varphi\} t \{\psi\}$. □

A corollary of this result and Lemma 4.14 is that if $\emptyset \vdash \{\varphi\} t \{\psi\}$ then no fault is reachable from an initial marking of $\mathcal{N}[[t]]$ if the heap initially satisfies φ .

5. SEPARATION

As mentioned in the introduction, the logic discriminates between the parallel composition of processes and their interleaved expansion. In Brookes' trace semantics [Bro07], this was accounted for by making the notion of a *race* primitive within the semantics: when forming the parallel composition of processes, if two processes concurrently write to the same location, a special 'race' action occurs and the trace proceeds no further. Our approach when defining the semantics has been different; we do not regard a race as 'catastrophic' and have not embellished our semantics with special race states. Instead, we shall prove, using the semantics directly, that races do not occur for proved processes running from suitable initial states.

Generally, a race can be said to occur when two interacting heap actions occur concurrently. Recall that a heap action is represented in the net semantics by a set of events, with common pre- and post-control conditions, representing each way in which the action can affect the heap. According to the net model, two actions may be allowed to run concurrently if their events do not overlap on their pre- or post-control conditions. In such a situation, where ${}^c e_1 \cap {}^c e_2 = \emptyset$, we shall say that e_1 and e_2 are *control-independent*.

One way of capturing the race freedom of a process running from an initial state is to show that there is no reachable marking in the net where two control-independent events are control-enabled but access a common heap location, except interaction through allocation. We, however, shall prove a result based on the *behaviour* of processes: that whenever two events are control-independent and can occur, then either they are independent or they lie within a form of prescribed class of action.

Definition 5.1 (Separation of synchronized events). Let M be a marking of $\mathcal{W}[[t]]_\Gamma$ and let $s_1 = e_1 \cdot u_1$ and $s_2 = e_2 \cdot u_2$ be control-independent synchronized events of $\mathcal{W}[[t]]_\Gamma$. The *separation property* of s_1 and s_2 at M is defined as:

- (1) If $M \xrightarrow{s_1} M_1$ and $M \xrightarrow{s_2} M_2$ and s_1 and s_2 are not independent then either:
- s_1 and s_2 compete to allocate the same location: $e_1 = \text{alloc}_{(C_1, C'_1)}(\ell, v, \ell', v')$ and $e_2 = \text{alloc}_{(C_2, C'_2)}(k, w, \ell', w')$ for some $\ell, \ell', k, v, v', w'$;
 - s_1 and s_2 compete to make the same resource current: $e_1 = \text{decl}_{(C_1, C'_1)}(r)$ and $e_2 = \text{decl}_{(C_2, C'_2)}(r)$ for some r ; or
 - s_1 and s_2 compete to acquire the same resource: $e_1 = \text{acq}_{(C_1, C'_1)}(r)$ and $e_2 = \text{acq}_{(C_2, C'_2)}(r)$ for some r .

- (2) If $M \xrightarrow{s_1} M_1 \xrightarrow{s_2} M_2$ and s_1 and s_2 are not independent then either:
- s_1 deallocates a location that s_2 allocates: $e_1 = \text{dealloc}_{(C_1, C_2)}(\ell, v, \ell', v')$ and $e_2 = \text{alloc}_{(C_2, C'_2)}(k, w, \ell', w')$ for some $\ell, \ell', k, v, v', w'$;
 - s_1 makes a resource non-current that s_2 makes current: $e_1 = \text{end}_{(C_1, C'_1)}(r)$ and $e_2 = \text{decl}_{(C_2, C'_2)}(r)$ for some r ; or
 - s_1 releases a resource that s_2 takes: $e_1 = \text{rel}_{(C_1, C'_1)}(r)$ and $e_2 = \text{acq}_{(C_2, C'_2)}(r)$ for some r .
- (3) The symmetric statement for $M \xrightarrow{s_2} M_2 \xrightarrow{s_1} M_1$.

The first part of the property above tells us how the enabled events of parallel processes *conflict* with each other in a state: the way in which one parallel process can prevent the other acting in a particular way on the global state. The second part dictates how the event occurrences of parallel processes *causally depend* on each other: the way in which the ability of one process to affect the global state in a particular way is dependent on events of the other process.

Importantly, whenever the two events s_1 and s_2 arise from heap actions, they neither conflict nor causally depend on each other. This is our net analogue of race freedom. Theorem 5.4 shows that processes proved by the logic are race free when running from suitable initial states. We shall make use of the following rather technical lemmas in the proof.

For a synchronized event s and an interference event u , define the separation property for s and u at M similarly, recalling that any synchronized event is trivially control-independent from any interference event because ${}^c u^c = \emptyset$ for any interference event u . It is always the case that a synchronized event and an interference event satisfy the separation property in any consistent marking.

Lemma 5.2. *If M is a consistent marking of $\mathcal{W}[[t]]_\Gamma$ and s is a synchronized event and u is an interference event then s and u satisfy the separation property in M .*

Proof. A straightforward analysis of the many cases for s and u . □

The following lemma relates independence from an interference event to independence from any corresponding synchronized event. Recall that we write eIe' if e and e' are independent.

Lemma 5.3. *Let s be any synchronized event of $\mathcal{W}[[t]]_\Gamma$ and u be an interference event of $\mathcal{W}[[t]]_\Gamma$. Suppose that M is a consistent marking in which they both have concession. If e_1 is an event of $\mathcal{N}[[t]]$ that synchronizes with u and sIu and s is control-independent from e_1 then $sI(e_1 \cdot u)$.*

Proof. It is easy to see that the preconditions of $e_1 \cdot u$ are simply the preconditions of e_1 along with the pre-control conditions of e_1 apart from replacing $\omega_{\text{oth}}(\ell)$ with $\omega_{\text{proc}}(\ell)$ and replacing $\omega_{\text{oth}}(r)$ with $\omega_{\text{proc}}(r)$. The postconditions of $e_1 \cdot u$ are similar.

Suppose, for contradiction, that $\neg(sI(e_1 \cdot u))$. Since sIu and s is control-independent from e_1 , it follows that there must exist $z \in \text{Loc} \cup \text{Res}$ such that $\omega_{\text{proc}}(z) \in {}^\bullet s^\bullet \cap {}^\bullet (e_1 \cdot u)^\bullet$. From the definition of synchronization, we therefore have $\omega_{\text{oth}}(z) \in {}^\bullet u^\bullet$. The proof is completed by analysis of the cases for how $\omega_{\text{proc}}(z) \in {}^\bullet s^\bullet$; we shall show only one illustrative case, that where z is a location ℓ such that $\omega_{\text{proc}}(\ell) \in {}^\bullet s$ but $\omega_{\text{proc}}(\ell) \notin s^\bullet$.

In this case, the event s must either deallocate the location ℓ or must release a resource r with $r \in \text{dom}(\Gamma)$ and ℓ forms part of the heap used to satisfy the invariant for r . As the

event s has concession in M , we have $\omega_{\text{proc}}(\ell) \in M$. By assumption, u has concession in M and $\omega_{\text{oth}}(\ell) \in \bullet u \bullet$. We cannot have $\omega_{\text{oth}}(\ell) \in \bullet u$ since $\omega_{\text{proc}}(\ell) \in M$, so $\omega_{\text{oth}}(\ell) \in u \bullet$. Therefore, the event u is an interference event that either allocates the location ℓ or acquires an open resource r and ℓ is part of the heap that satisfies the invariant for r . If u is such an event, that acquires r , it must be the case that $\omega_{\text{inv}}(\ell) \in \bullet u$ so $\omega_{\text{inv}}(\ell) \in M$, contradicting that M is a consistent marking with $\omega_{\text{proc}}(\ell) \in M$. Consequently, u must in fact be an event that allocates the location ℓ , so therefore $\text{curr}(\ell) \notin M$. We then arrive at another contradiction since it must then be the case that $\omega_{\text{proc}}(\ell) \notin M$ because M is consistent. \square

We may now show that the separation property does indeed hold for any two events s_1 and s_2 in $\mathcal{W}[[t]]_\Gamma$ for any term t and environment Γ such that $\Gamma \vdash \{\varphi\} t \{\psi\}$ in any marking $M = (C, \sigma, W)$ reachable from an initial marking of t that satisfies φ in Γ . The proof is most interesting in the case where $t \equiv t_1 \parallel t_2$ and s_1 is an event of t_1 and s_2 is an event of t_2 . The case proceeds by establishing, as in Theorem 4.26, that there exists an ownership split W_1 and W_2 of W for which s_1 has concession in (C_1, σ, W_1) , where C_1 is the marking of control conditions in C for t_1 , and there exist e_2 and u_2 such that $s_2 = (2:e_2) \cdot u_2$ and u_2 also has concession in the marking (C_1, σ, W_1) of $\mathcal{W}[[t_1]]_\Gamma$. By Lemma 5.2, the separation property therefore holds for s_1 and u_2 in the marking (C_1, σ, W_1) . It follows that the separation property holds for s_1 and s_2 in M since, by Lemma 5.3, if the events s_1 and u_2 are independent then so are s_1 and s_2 .

Theorem 5.4 (Separation). *Suppose that $\Gamma \vdash \{\varphi\} t \{\psi\}$ and that $(\text{Ic}(t), \sigma_0, W_0)$ satisfies φ in Γ . For any events s_1 and s_2 in $\mathcal{W}[[t]]_\Gamma$ and any marking (C, σ, W) reachable from $(\text{Ic}(t), \sigma_0, W_0)$, the separation property holds for s_1 and s_2 at (C, σ, W) .*

Proof. By induction on the derivation of $\Gamma \vdash \{\varphi\} t \{\psi\}$. We shall show only one case:

(L-PAR): Assume that the marking $(\text{Ic}(t_1 \parallel t_2), \sigma_0, W_0)$ of $\mathcal{W}[[t_1 \parallel t_2]]_\Gamma$ satisfies $\varphi_1 * \varphi_2$ in Γ and that $M = (1:C_1 \cup 2:C_2, \sigma, W)$ is reachable from this marking. There exist W_{01} and W_{02} forming an ownership split of W_0 such that the marking $(\text{Ic}(t_1), \sigma_0, W_{01})$ of $\mathcal{W}[[t_1]]_\Gamma$ satisfies φ_1 in Γ and the marking $(\text{Ic}(t_2), \sigma_0, W_{02})$ of $\mathcal{W}[[t_2]]_\Gamma$ satisfies φ_2 in Γ . By assumption, $\Gamma \vdash \{\varphi_1\} t_1 \{\psi_1\}$ and $\Gamma \vdash \{\varphi_2\} t_2 \{\psi_2\}$, so according to Theorem 4.26 no violating marking is reachable from either of these markings.

Let s_1 and s_2 be synchronized events in $\mathcal{W}[[t_1 \parallel t_2]]_\Gamma$. If $s_1 = (1:e_1) \cdot u_1$ and $s_2 = (1:e_2) \cdot u_2$ for some $e_1, e_2 \in \text{Ev}(t_1)$ and interference events u_1 and u_2 in $\mathcal{W}[[t_1]]_\Gamma$, the result follows routinely from the induction hypothesis, and similarly if s_1 and s_2 both arise from events of $\mathcal{N}[[t_2]]$. Suppose instead that there exist $e_1 \in \text{Ev}(t_1)$, $e_2 \in \text{Ev}(t_2)$ and interference events u_1 and u_2 such that $s_1 = (1:e_1) \cdot u_1$ and $s_2 = (2:e_2) \cdot u_2$.

Suppose first that in the net $\mathcal{W}[[t_1 \parallel t_2]]_\Gamma$ we have

$$M \xrightarrow{s_1} (1:C'_1 \cup 2:C'_2, \sigma', W') \xrightarrow{s_2} (1:C''_1 \cup 2:C''_2, \sigma'', W'').$$

A simple induction applying the parallel decomposition lemma (Lemma 4.17) along the path to M shows that there exist W_1 and W_2 that form an ownership split of W such that

$$(C_1, \sigma, W_1) \xrightarrow{e_1 \cdot u_1} (C'_1, \sigma', W'_1) \xrightarrow{u_2} (C''_1, \sigma'', W''_1)$$

in $\mathcal{W}[[t_1]]_\Gamma$ for some W'_1, W''_1 . By Lemma 5.2, the separation property holds for $e_1 \cdot u_1$ and u_2 in (C_1, σ, W_1) ; consider how it might hold. If $e_1 \cdot u_1$ deallocates a location that u_2 allocates, then s_1 deallocates a location that s_2 allocates, so the separation property holds for s_1 and s_2 . The argument is similar for all the other cases where $e_1 \cdot u_1$ and u_2 are not

independent. Suppose instead that $e_1 \cdot u_1 I u_2$. The event u_2 has concession in the marking (C_1, σ, W_1) by virtue of the fact that the occurrence of independent events in a run can be interchanged (Proposition 3.4). Consider the marking $(1:C_1 \cup 2:C_2, \sigma, W_1)$ of $\mathcal{W}[[t_1 \parallel t_2]]_\Gamma$; this is straightforwardly seen to be consistent. The event s_1 is readily seen using Lemma 3.6 to have concession in this marking, as does u_2 . The event $2:e_2$ is control-independent from $1:e_1$, so by Lemma 5.3 we have $s_1 I s_2$, as required.

Now suppose that in the net $\mathcal{W}[[t_1 \parallel t_2]]_\Gamma$ we have

$$M \xrightarrow{s_1} (1:C'_1 \cup 2:C'_2, \sigma', W') \quad \text{and} \quad M \xrightarrow{s_2} (1:C''_1 \cup 2:C''_2, \sigma'', W'').$$

A simple induction applying the parallel decomposition lemma (Lemma 4.17) along the path to M shows that there exist W_1 and W_2 that form an ownership split of W such that

$$(C_1, \sigma, W_1) \xrightarrow{e_1 \cdot u_1} (C'_1, \sigma', W'_1) \quad \text{and} \quad (C_1, \sigma, W_1) \xrightarrow{u_2} (C''_1, \sigma'', W''_1)$$

in $\mathcal{W}[[t_1]]_\Gamma$ for some W'_1, W''_1 . By Lemma 5.2, the separation property holds for $e_1 \cdot u_1$ and u_2 in (C_1, σ, W_1) ; consider how it might hold. If $e_1 \cdot u_1$ allocates a location that u_2 also allocates, then s_1 allocates a location that s_2 allocates, so the separation property holds for s_1 and s_2 . The argument is similar for all the other cases where $e_1 \cdot u_1$ and u_2 are not independent. Suppose instead that $e_1 \cdot u_1 I u_2$. Consider the marking $(1:C_1 \cup 2:C_2, \sigma, W_1)$ of $\mathcal{W}[[t_1 \parallel t_2]]_\Gamma$; this is readily seen to be consistent. The event s_1 has concession in this marking as does u_2 . The event $2:e_2$ is control-independent from $1:e_1$, so by Lemma 5.3 we have $s_1 I s_2$, as required.

The remaining cases of the proof follow relatively straightforwardly by induction. The case for (L-RES) requires an observation along the lines of Lemma 4.25; that, for any term t with $\text{fv}(t) \subseteq \{w\}$ and resources $r, r' \notin \text{dom}(\Gamma)$, if the separation property holds for any two synchronized events of $\mathcal{W}[[r/w]t]]_\Gamma$ in any marking reachable from any initial marking satisfying φ in Γ then it also holds for $\mathcal{W}[[r'/w]t]]_\Gamma$.

The proof for the rule (L-SEQ) follows straightforwardly by induction using Lemma 3.15 except in the second (and symmetric third) cases of the definition of the separation property, where there are reachable markings M, M', M'' such that $M \xrightarrow{s_1} M' \xrightarrow{s_2} M''$ and there exist events $e_1 \in \text{Ev}(t_1)$ and $e_2 \in \text{Ev}(t_2)$ and interference events u_1, u_2 such that $s_1 = (P \triangleleft 1:e_1) \cdot u_1$ and $s_2 = (P \triangleright 2:e_2) \cdot u_2$ for $P = 1:\text{Tc}(t_1) \times 2:\text{Ic}(t_2)$. In this case, it follows from Lemma 3.15 and Lemma 3.12 that the events s_1 and s_2 are not control-independent. \square

The result can be applied, using Lemma 4.14 and the observation that $e_1 \cdot u_1 I e_2 \cdot u_2$ implies that $e_1 I e_2$, to obtain a similar result for the net semantics of terms without ownership.

Corollary 5.5. *Let t be a closed term. Suppose that $\emptyset \vdash \{\varphi\} t \{\psi\}$ and that $\sigma_0 = (D_0, L_0, \emptyset, \emptyset)$ is a state for which $D_0 \models \varphi$. If M is a marking reachable from $(\text{Ic}(t), \sigma_0)$ in $\mathcal{N}[[t]]$ and e_1 and e_2 are control-independent events then:*

- If $M \xrightarrow{e_1} M_1 \xrightarrow{e_2} M'$ then either e_1 and e_2 are independent or e_1 releases a resource or a location that e_2 correspondingly takes or allocates, or e_1 makes non-current a resource that e_2 makes current.
- If $M \xrightarrow{e_1} M_1$ and $M \xrightarrow{e_2} M_2$ then either e_1 and e_2 are independent or e_1 and e_2 compete either to make current the same resource, acquire the same resource or to allocate the same location. \square

5.1. Incompleteness. The separation result highlights an important form of possible interaction between concurrent processes. Observe that, although there is neither conflict nor causal dependence arising from heap events (and hence the processes are race-free in the sense of Brookes), there may be interaction through the occurrence of allocation and deallocation events. One may therefore give judgements for parallel processes that interact without using critical regions. Suppose, for example, that we have a heap

$$D = \{\ell_0 \mapsto \ell_1, \ell_1 \mapsto 1, \ell_2 \mapsto 2, \ell_3 \mapsto 3, \ell_4 \mapsto 4\}.$$

For any processes t_1 and t_2 such that t_1 does not deallocate ℓ_1 , if we place the process

$$t_1; \text{dealloc}(\ell_0)$$

in parallel with

$$\text{alloc}(\ell_2); \text{ while } [\ell_2] \neq \ell_1 \text{ do } \text{alloc}(\ell_2) \text{ od}; t_2,$$

the process t_2 only takes place once t_1 has terminated, and possibly never, even if t_1 terminates. This arises from the fact that the loop in the second process will only exit when location ℓ_1 is allocated by the command $\text{alloc}(\ell_2)$; this can only occur once $\text{dealloc}(\ell_0)$ makes ℓ_1 non-current and therefore available for allocation by $\text{alloc}(\ell_2)$. Denote this process $\text{seq}(t_1, t_2)$.

We can use this to show that concurrent separation logic is incomplete with respect to our definition of validity: Let t_1 be the assignment $[\ell_3] := 1$ and t_2 be $[\ell_3] := 2$. Define the formula

$$\delta \stackrel{\text{def}}{=} \ell_0 \mapsto \ell_1 * \ell_1 \mapsto - * \ell_2 \mapsto - * \ell_3 \mapsto - * \ell_4 \mapsto -.$$

We have $\emptyset \models \{\delta\} \text{seq}(t_1, t_2) \{\ell_3 \mapsto 2 * \top\}$ since, whenever $\text{seq}(t_1, t_2)$ terminates, the assignment $[\ell_3] := 2$ always occurs after the assignment $[\ell_3] := 1$. The separation property holds in any marking reachable from any heap initially satisfying δ . It can be shown that

$$\emptyset \not\models \{\delta\} \text{seq}(t_1, t_2) \{\ell_3 \mapsto 2 * \top\},$$

so the logic is incomplete, even for processes satisfying the separation property.

There are also examples of incompleteness where neither process accesses a common heap location along any run: Let

$$\begin{aligned} t'_1 &= \text{alloc}(\ell_3); \text{while } [\ell_3] \neq \ell_5 \text{ do } \text{alloc}(\ell_3) \text{ od} \\ t'_2 &= \text{alloc}(\ell_4); ([\ell_4] = \ell_5). \text{skip} + ([\ell_4] \neq \ell_5). \text{diverge}, \end{aligned}$$

for the previous definition of diverge and the obvious definition of skip , $\mathcal{A}[\text{skip}] = \{(\emptyset, \emptyset)\}$. Since the location ℓ_5 is always current following termination of t'_1 from D , process t'_2 always diverges. We have

$$\emptyset \models \{\delta\} \text{seq}(t'_1, t'_2) \{\perp\}.$$

However, there are no δ_1, δ_2 such that δ is logically equivalent to $\delta_1 * \delta_2$ and $\emptyset \models \{\delta_2\} t'_2 \{\perp\}$, which would be necessary if it were possible to prove $\emptyset \vdash \{\delta\} \text{seq}(t'_1, t'_2) \{\perp\}$.

6. REFINEMENT

As we remarked in the introduction, the atomicity assumed of primitive actions, also called their *granularity*, is of significance when considering parallel programs. For example, suppose that the concurrent program

$$\begin{aligned} & [\ell] := [\ell'] + 1 \\ \parallel & ([\ell]' \neq [\ell]).\mathbf{diverge} + ([\ell]' = [\ell]).\mathbf{skip} \end{aligned}$$

runs from the heap $\{\ell \mapsto 0, \ell' \mapsto 1\}$. Given the prior interpretations of **skip** and **diverge**, we might conclude that the program never terminates since the assignment $[\ell] := [\ell'] + 1$ maintains the property through execution that ℓ and ℓ' hold different values.

It may not, however, be reasonable to assume that the assignment is executed atomically. For instance, the processor on which the process runs might have primitive actions for copying the values held in memory locations and for incrementing them, but not for copying *and* incrementing in one clock step. The process $[\ell] := [\ell'] + 1$ might therefore be compiled to execute as $[\ell] := [\ell']; [\ell] := [\ell] + 1$. Quite clearly, the process

$$\begin{aligned} & [\ell] := [\ell']; [\ell] := [\ell] + 1 \\ \parallel & ([\ell]' \neq [\ell]).\mathbf{diverge} + ([\ell]' = [\ell]).\mathbf{skip} \end{aligned}$$

may terminate, so we failed to exhibit a proper degree of caution when asserting that it would fail to terminate.

In [Rey04], Reynolds proposes a form of trace semantics that regards the occurrence of uncontrolled interference between concurrent processes as ‘catastrophic’. The motivation behind this is the race freedom property arising from concurrent separation logic [Bro04]: in the semantics of a proved process running from a suitable initial state, no uncontrolled interference may occur. Reynolds’ observation is that, in this situation, judgements may be made that are insensitive to atomicity.

Within our net model we can provide a form of *refinement*, similar to that of [vGG89] but suited to processes executing in a shared environment, that begins to capture these ideas. Importantly, the property required to apply the refinement operation may be captured directly in terms of independence, with no changes to our semantics. We will relate the nets representing processes with different levels of atomicity by regarding them as alternative substitutions into a *context*. We will then give a condition on substitutions led by Theorem 5.4 to show that any partial correctness assertion made for one of the nets also holds for the other.

The treatment of substitution requires some restrictions to be placed on the nets we consider. In the remainder of this section and in Appendix A where we present the technical details of this section, we require that all embedded nets satisfy the structural properties described in Lemma 3.11 and Definition 3.13.

Definition 6.1 (Context). Define a *context* K to be an embedded net with a distinguished event $[-]$. The event $[-]$ is such that $\bullet[-]\bullet \subseteq \mathbf{C}$ and its pre- and postconditions form disjoint, nonempty sets.

We may now construct the net representing the substitution of a net N for the hole in a context K . We shall assume that, as in the semantics for terms, the two nets are formed with the same sets of conditions. As the nets are extensional (we regard an event simply as its set of preconditions paired with its set of postconditions), all that we need to specify is the events of the net and its initial and terminal markings of control conditions.

Definition 6.2 (Substitution). Let K be a context and N an embedded net. Define the sets

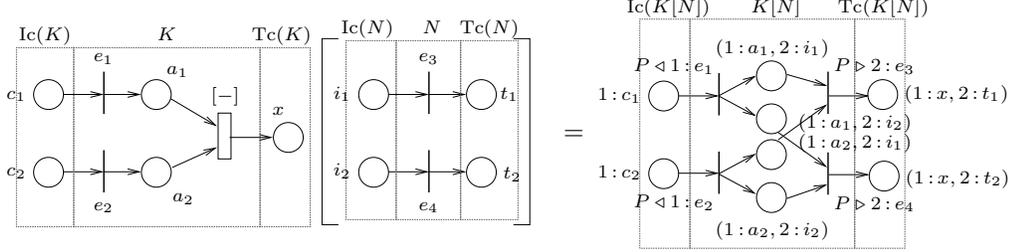
$$P_i \stackrel{\text{def}}{=} 1 : \bullet[-] \times 2 : \text{Ic}(N) \quad P_t \stackrel{\text{def}}{=} 1 : [-]^\bullet \times 2 : \text{Tc}(N).$$

The substitution $K[N]$ is defined to be the embedded net with:

$$\begin{aligned} \text{Ev}(K[N]) &\stackrel{\text{def}}{=} (P_i \cup P_t) \triangleleft 1 : (\text{Ev}(K) \setminus \{-\}) \cup (P_i \cup P_t) \triangleright 2 : \text{Ev}(N) \\ \text{Ic}(K[N]) &\stackrel{\text{def}}{=} (P_i \cup P_t) \triangleleft 1 : \text{Ic}(K) \\ \text{Tc}(K[N]) &\stackrel{\text{def}}{=} (P_i \cup P_t) \triangleleft 1 : \text{Tc}(K) \end{aligned}$$

To see the definition at work, consider the following example. We elide details of the action of events on state conditions, which is unaffected by the substitution operation.

Example 6.3. In the following example substitution, we depict the hole $[-]$ as a hollow rectangle.



Definition 6.4. Let π be a sequence of events of the net N . Sequence π is said to be *complete from σ to σ'* if

$$\pi : (\text{Ic}(N), \sigma) \longrightarrow^* (\text{Tc}(N), \sigma').$$

Write $N : \sigma \Downarrow \sigma'$ if there exists a complete sequence from σ to σ' in N .

Using this definition, we can define a notion of complete trace equivalence \simeq as:

$$N_1 \simeq N_2 \quad \text{iff} \quad (\forall \sigma, \sigma') N_1 : \sigma \Downarrow \sigma' \iff N_2 : \sigma \Downarrow \sigma'.$$

We wish to constrain K , N_1 and N_2 appropriately so that if $N_1 \simeq N_2$ then $K[N_1] \simeq K[N_2]$.

Example 6.5. Write, in the obvious way, $-$ for the action term that will be interpreted as forming the hole of a context. Define

$$\begin{aligned} K &\stackrel{\text{def}}{=} \mathcal{N}[- \parallel ([\ell]' \neq [\ell]).\text{diverge} + ([\ell]' = [\ell]).\text{skip}] \\ N_1 &\stackrel{\text{def}}{=} \mathcal{N}[[\ell] := [\ell]' + 1] \\ N_2 &\stackrel{\text{def}}{=} \mathcal{N}[[\ell] := [\ell]'; [\ell] := [\ell] + 1]. \end{aligned}$$

We clearly have $N_1 \simeq N_2$, but $K[N_1] \not\simeq K[N_2]$ since

$$K[N_1] : \{\ell \mapsto 0, \ell' \mapsto 1\} \not\Downarrow \{\ell \mapsto 2, \ell' \mapsto 1\}$$

but

$$K[N_2] : \{\ell \mapsto 0, \ell' \mapsto 1\} \Downarrow \{\ell \mapsto 2, \ell' \mapsto 1\}.$$

Return to the general case for a substitution $K[N]$. Intuitively, if the substituend N were an atomic event, it would start running only if the conditions P_i were marked and P_t were not. There are two distinct ways in which the context K can affect the execution of N . Firstly, it might affect the marking of conditions in P_i or P_t whilst N is running. Secondly,

it might change the marking of state conditions in a way that affects the execution of N . An instance of the latter form of interference is seen in the preceding example. We now define a form of constrained substitution, guided by Theorem 5.4, so that N is not subject to these forms of interference.

Say that a control condition c of $K[N]$ is *internal* to N if $c = 2 : c_2$ where c_2 is a pre- or a postcondition of an event of N that is not in $\text{Ic}(N)$ or $\text{Tc}(N)$. Given a marking M of $K[N]$, say that N is *active* if $P_i \subseteq M$ or there exists an internal condition of N in M .

Definition 6.6. For a given marking of state conditions σ , we say that $K[N]$ is a *non-interfering substitution* if, for all markings M reachable from $(\text{Ic}(K[N]), \sigma)$:

- (1) if $P_i \subseteq M$ then $P_t \cap M = \emptyset$, and
- (2) if N is active in M then no enabled event of K has a pre- or postcondition in P_i or P_t , and
- (3) if $M \xrightarrow{e_1} M_1 \xrightarrow{e_2} M'$, one of e_1 and e_2 is from N and the other is from K and N is active in M and M_1 , then e_1 and e_2 are independent.

Theorem 6.7. *If $N_1 \simeq N_2$ and $K[N_1]$ and $K[N_2]$ are non-interfering substitutions from state σ , then, for any σ' :*

$$K[N_1] : \sigma \Downarrow \sigma' \quad \text{iff} \quad K[N_2] : \sigma \Downarrow \sigma'.$$

Proof. Appendix A, Theorem A.12. □

The refinement operation defined in this section allows us to change the granularity of heap actions by substituting the occurrence of an action in the original net with a net representing the actual implementation of the action, but only once it has been shown that the noninterference property holds for both the original net and for the net formed. The operation might be a key to proving Reynolds' observation that an occurrence of an action α in the term t can be replaced by a term with the same overall behaviour as α without affecting the validity of the judgement $\Gamma \vdash \{\varphi\} t \{\psi\}$.

7. RELATED WORK AND CONCLUSIONS

The first component of this work provides an inductive definition of the semantics as a net of programs operating in a (shared) state. This is a relatively novel technique, but has in the past been applied to give the semantics of a language for investigating security protocols, SPL [CW01], though our language involves a richer collection of constructs. Other independence models for terms include the Box calculus [BDH92] and the event structure and net semantics of CCS [Stu80, Win82, WN95] ([Stu80] was, to our knowledge, the first Petri net denotational semantics of CCS), though these model interaction as synchronized communication rather than occurring through shared state. We hope that the novel Petri net semantics presented here and in [CW01] can be the start of *systematic and comprehensive* methods to attribute structural Petri net semantics to a full variety of programming languages, resulting in a Petri net companion to Plotkin's structural operational semantics (SOS) based on transition systems [Plo81]. Paralleling the (inductive) definitions of data and transitions of SOS would be (inductive) definitions of conditions and events of Petri nets.

The proof of soundness of separation logic here is led by Brookes' earlier work [Bro07]. There are a few minor differences in the syntax of processes, including that we allow the

dynamic binding of resource variables. Another minor difference between the programming language and logic considered here and that introduced by O’Hearn and proved sound by Brookes is that we do not distinguish *stack variables*. These may be seen as locations to which other locations may not point and are the only locations that terms can directly address. In Brookes’ model, as in [O’H07], interference of parallel processes through stack variables is constrained by the use of a side condition on the rule rather than using the concept of ownership (the area of current research on ‘permissions’ [BCOP05, BCY05, Bro06] promises a uniform approach). In particular, the rule allows the concurrent reading of stack locations. Though we have chosen not to include stack variables in our model in order to highlight the concept of ownership, our model and proofs could be easily extended to deal with them. Concurrent reading of memory would be at the cost of a more sophisticated notion of independence that allowed independent events to access the same condition providing that neither affects the marking of that condition.

More notably, at the core of Brookes’ work is a ‘local enabling relation’, which gives the semantics of programs over a restricted set of ‘owned’ locations. Our notion of validity involves maintaining a record of ownership and using this to constrain the occurrence of events in the interference net augmented to the process. This allows the intuition of ownership in O’Hearn’s introduction of concurrent separation logic [O’H07] to be seen directly as constraining interference. Though the relationship between our model and Brookes’ is fairly obvious, we believe that our approach leads to a clearer parallel decomposition lemma, upon which the proof of soundness of the logic critically stands.

The most significant difference between our work and Brookes’ is that the net model captures, as a primitive property, the independence of parallel processes enforced by the logic. We have used this property to define a straightforward, yet general, form of refinement suited to changing the atomicity of commands within the semantics of a term. This is in contrast to [Bro05], which gives a new semantics to race-free processes that abstracts entirely away from attaching any form of atomicity to the semantics of heap actions. As said at the end of the previous section, we hope to show that the refinement operation can be applied to change the atomicity of any action occurring within any process running from a suitable initial state proved using to the rules of concurrent separation logic.

Our characterization of ‘separation’ arising from the logic is much finer than that obtained from the existing proof of race freedom, for example showing that interaction between parallel processes may occur through allocation and deallocation. This is significant, as such interaction leads to examples of the incompleteness of concurrent separation logic.

There are a number of other areas for further research in addition to those mentioned above. One interesting consideration is the necessity (or otherwise) of precision in the proof of soundness of the logic. In forthcoming work, we hope to give a form of game semantics for the logic and a soundness proof without precision in the absence of the Hoare’s Law of Conjunction (L-CONJUNCTION). Another area of interest is whether symmetry present in our semantics for allocation and resource declaration might be exploited properly to obtain more compact nets to represent processes.

ACKNOWLEDGEMENTS

It is a pleasure to thank Peter O’Hearn and Matthew Parkinson for a number of helpful discussions during the development of this work. We would also like to thank the anonymous

referees of a draft of this paper and the anonymous referees of the conference version of this paper [HW06] for their constructive suggestions.

REFERENCES

- [BCOP05] Richard Bornat, Cristiano Calcagno, Peter O’Hearn, and Matthew Parkinson. Permission accounting in separation logic. In *Proc. POPL ’05*. ACM Press, 2005.
- [BCY05] Richard Bornat, Cristiano Calcagno, and Hongseok Yang. Variables as resource in separation logic. In *Proc. MFPS XXI*, ENTCS, 2005.
- [BDH92] Eike Best, Raymond Devillers, and Jon G. Hall. The box calculus: A new causal algebra with multi-label communication. In *Advances in Petri Nets*, volume 609 of *LNCS*. Springer-Verlag, 1992.
- [Bri72] Per Brinch Hansen. Structured multiprogramming. *Comm. ACM*, 15(7):574–578, 1972.
- [Bro04] Stephen Brookes. A semantics for concurrent separation logic. In *Proc. CONCUR ’04*, volume 3170 of *LNCS*. Springer-Verlag, 2004.
- [Bro05] Stephen Brookes. A grainless semantics for parallel programs with shared mutable data. In *Proc. MFPS XXI*, ENTCS, 2005.
- [Bro06] Stephen Brookes. Variables as resource for shared-memory programs: Semantics and soundness. In *Proc. MFPS XXII*, ENTCS, 2006.
- [Bro07] Stephen Brookes. A semantics for concurrent separation logic. *Theoretical Computer Science*, 375(1–3), 2007. Extended version of [Bro04].
- [BRR87] Wilfried Brauer, Wolfgang Reisig, and Grzegorz Rozenberg, editors. *Advances in Petri Nets 1986: Proceedings of an Advanced Course*, volume 254–255 of *LNCS*. Springer, 1987.
- [CGMP99] Edmund M. Clarke, Orna Grumberg, Marius Minea, and Doron Peled. State space reduction using partial order techniques. *Int. Journal on Software Tools for Technology Transfer*, 2(3), 1999.
- [CW01] Federico Crazzolaro and Glynn Winskel. Events in security protocols. In *Proc. CCS ’01*, New York, 2001. ACM Press.
- [Dij68] Edsger Dijkstra. Cooperating sequential processes. In F. Genuys, editor, *Programming Languages*. Academic Press, 1968.
- [Hoa72] C. A. R. Hoare. Towards a theory of parallel programming. In C. A. R. Hoare and R. H. Perrot, editors, *Operating Systems Techniques*. Academic Press, 1972.
- [HW06] Jonathan Hayman and Glynn Winskel. Independence and concurrent separation logic. In *Proc. LICS ’06*. IEEE Press, 2006.
- [IO01] Samin S. Ishtiaq and Peter W. O’Hearn. BI as an assertion language for mutable data structures. In *Proc. POPL ’01*. ACM Press, 2001.
- [JNW93] André Joyal, Mogens Nielsen, and Glynn Winskel. Bisimulation from open maps. In *Proc. LICS ’93*, volume 127(2) of *Information and Computation*. Elsevier, 1993.
- [Jon83] Cliff B. Jones. Specification and design of (parallel) programs. In R. E. A. Mason, editor, *Information Processing 83: Proc. IFIP Congress*, pages 321–332, 1983.
- [Lam86] Leslie Lamport. On interprocess communication. *Distributed Computing*, 1(2):77–101, June 1986.
- [NW96] Mogens Nielsen and Glynn Winskel. Petri nets and bisimulation. *Theoretical Computer Science*, 153(1–2), 1996.
- [OG76] Susan Owicki and David Gries. Verifying properties of parallel programs: An axiomatic approach. *Comm. ACM*, 19(5):279–285, 1976.
- [O’H04] Peter W. O’Hearn. Resources, concurrency and local reasoning. In *Proc. CONCUR ’04*, volume 3170. Springer-Verlag, 2004.
- [O’H07] Peter W. O’Hearn. Resources, concurrency and local reasoning. *Theoretical Computer Science*, 375(1–3):271–307, 2007. Extended version of [O’H04].
- [OP99] Peter W. O’Hearn and David J. Pym. The logic of bunched implications. *Bulletin of Symbolic Logic*, 5(2), 1999.
- [Pet62] Carl Adam Petri. *Kommunikation mit Automaten*. PhD thesis, Institut für Instrumentelle Mathematik, Bonn, 1962.
- [Plo81] Gordon D Plotkin. A structural approach to operational semantics. Technical Report DAIMI FN-19, University of Aarhus, 1981.

- [Pra86] Vaughan Pratt. Modeling concurrency with partial orders. *Int. Journal of Parallel Programming*, 15(1), 1986.
- [Rey00] John C. Reynolds. Intuitionistic reasoning about shared mutable data structure. In *Millennial Perspectives in Computer Science*, 2000.
- [Rey04] John C. Reynolds. Towards a grainless semantics for shared variable concurrency. In *Proc. FSTTCS '04*, volume 3328 of *LNCS*. Springer-Verlag, 2004.
- [Stu80] Students. Projects for the course “Models for Concurrency” given by G. Winskel and M. Nielsen. Dept. of Computer Science, University of Aarhus, Denmark, 1980.
- [vGG89] Rob J. van Glabbeek and Ursula Goltz. Equivalence notions for concurrent systems and refinement of actions. In *Proc. MFCS '89*, volume 379 of *LNCS*. Springer-Verlag, 1989.
- [vGV87] Robert van Glabbeek and Fritz Vaandrager. Petri net models for algebraic theories of concurrency. In *Proc. PARLE 87*, volume 259 of *LNCS*. Springer-Verlag, 1987.
- [Win82] Glynn Winskel. Event structure semantics for CCS and related languages. In *Proc. ICALP '82*, volume 140 of *LNCS*. Springer-Verlag, 1982.
- [Win86] Glynn Winskel. Event structures. In *Advances in Petri Nets, Part II*, volume 255 of *LNCS*. Springer-Verlag, 1986.
- [Win87] Glynn Winskel. Petri nets, algebras, morphisms and compositionality. *Information and Computation*, 7(3):197–238, 1987.
- [WN95] Glynn Winskel and Mogens Nielsen. Models for concurrency. In *Handbook of Logic and the Foundations of Computer Science*, volume 4, pages 1–148. OUP, 1995.

APPENDIX A. REFINEMENT

A sequence of events $\pi = (e_1, \dots, e_n)$ considered from a marking M can be thought of equivalently as a sequence $M \xrightarrow{e_1} M_1 \dots \xrightarrow{e_n} M_n$. To describe the structure of such sequences, we shall say that π from marking M is of form $\Pi_1 \cdot \Pi_2$ if there exist π_1 and π_2 such that $\pi = \pi_1 \cdot \pi_2$, where \cdot denotes the obvious concatenation of sequences, and π_1 is of form Π_1 from marking M and π_2 is of form Π_2 from the marking obtained by following π_1 from M . Sequence π is of form Π^* if it is the concatenation of a finite number of sequences, each of form Π .

Throughout this section, when we consider the substitution $K[N]$ let P_i and P_t be defined as in Definition 6.2:

$$\begin{aligned} P_i &\stackrel{\text{def}}{=} 1 : \bullet[-] \times 2 : \text{Ic}(N) \\ P_t &\stackrel{\text{def}}{=} 1 : [-]^\bullet \times 2 : \text{Tc}(N). \end{aligned}$$

Any reachable marking of conditions of the net can be partitioned into two sets: conditions that occur solely within K and conditions that are either N -internal or in P_i or P_t . Formally, a condition c is a K -condition if $c = 1 : c_1$ for some condition c_1 of K not in $\bullet[-]^\bullet$. A condition c is an N -condition if either $c \in P_i \cup P_t$ or $c = 2 : c_2$ for some condition c_2 of N not in $\text{Ic}(N) \cup \text{Tc}(N)$. Recall that we call $2 : c_2$ an N -internal condition. It is easy to see that, for any σ_0 , from the marking $(\text{Ic}(K[N]), \sigma_0)$ only K - or N -control conditions may be marked: If (C, σ) is a reachable marking of $K[N]$, we have $C = C_N \cup C_K$ for some marking C_N of N -conditions and some marking C_K of K -conditions. We shall frequently use the notation (C_N, C_K) for a marking of control conditions, where C_N comprises only N -conditions and C_K comprises only K -conditions.

Henceforth, when considering a substitution $K[N]$, we shall refer to an event e as being an N -event if it is equal to $(P_i \cup P_t) \triangleright 2 : e_2$ for some e_2 in N . Otherwise, it is a K -event. A little care is necessary since an event in the net $K[N]$ might arise from both K and N if there are events e and $e' \neq [-]$ of N and K , respectively, with the same effect on state conditions and:

$$\begin{aligned} c_e &= \text{Ic}(N), & e^c &= \text{Tc}(N) \\ c_{e'} &= \bullet[-], & e'^c &= [-]^\bullet. \end{aligned}$$

Throughout the remainder of this section, for simplicity we shall require that the substitution $K[N]$ has no such events. This restriction may be lifted with little effect on the development so-far by allowing the net formed to be non-extensional, or by considering this as a special case when demonstrating properties of the net $K[N]$

Lemma A.1. *In $K[N]$, no K event has as a either a pre- or a postcondition an N -internal condition.*

Proof. Immediate from the definition of substitution $K[N]$. □

Recall that a marking (C_N, C_K, σ) reachable from $(\text{Ic}(K[N]), \sigma_0)$ is N -active if either there is an N -internal condition in C_N or if $C_N = P_i$. It is useful to further classify the markings of conditions in C_N according to whether they support the occurrence of N - or K -events on the conditions P_i and P_t :

Definition A.2. A marking (C_N, C_K, σ) of $K[N]$ is an N -marking if for all $a, a' \in \bullet[-]$, $x, x' \in [-]^\bullet$, $i \in \text{Ic}(N)$ and $t \in \text{Tc}(N)$:

- if $(1 : a, 2 : i) \in C_N$ then $(1 : a', 2 : i) \in C_N$, and

- if $(1 : x, 2 : t) \in C_N$ then $(1 : x', 2 : t) \in C_N$.

A marking $(C_N \cup C_K, \sigma)$ of $K[N]$ is a K -marking if there is no N -internal condition marked, and furthermore, for all $a \in \bullet[-]$, $x \in [-]^\bullet$, $i, i' \in \text{Ic}(N)$ and $t, t' \in \text{Tc}(N)$:

- if $(1 : a, 2 : i) \in C_N$ then $(1 : a, 2 : i') \in C_N$, and
- if $(1 : x, 2 : t) \in C_N$ then $(1 : x, 2 : t') \in C_N$.

From a marking of control conditions (C_N, C_K) , we can extract markings of control conditions for the nets N and K . We define $\rho_N(C_N)$ to be the marking of N obtained from (C_N, C_K) , which is not dependent on the marking C_K of K -conditions, and $\rho_K(C_N, C_K)$ for the marking of K obtained from (C_N, C_K) , which is dependent on the marking of N -conditions (namely, the marking of N -conditions in $P_i \cup P_t$).

For a marking C of the context K , we define $\theta_K(C)$ to be the corresponding marking of $K[N]$. For a marking C' of the net N , we define $\theta_N(C')$ to be the marking of N -conditions in the net $K[N]$ corresponding to C' .

Definition A.3. Let $K[N]$ be any substitution. For any marking C_N of N -conditions and C_K of K -conditions, define

$$\begin{aligned} \rho_K(C_N, C_K) &\stackrel{\text{def}}{=} \left\{ \begin{array}{l} a \in \bullet[-] \quad | \quad \forall i \in \text{Ic}(N). (1 : a, 2 : i) \in C_N \\ \cup \{ \quad x \in [-]^\bullet \quad | \quad \forall t \in \text{Tc}(N). (1 : x, 2 : t) \in C_N \\ \cup \{ \quad c \notin \bullet[-] \cup [-]^\bullet \quad | \quad 1 : c \in C_K \end{array} \right\} \\ \rho_N(C_N) &\stackrel{\text{def}}{=} \left\{ \begin{array}{l} i \in \text{Ic}(N) \quad | \quad \forall a \in \bullet[-]. (1 : a, 2 : i) \in C_N \\ \cup \{ \quad t \in \text{Tc}(N) \quad | \quad \forall x \in [-]^\bullet. (1 : x, 2 : t) \in C_N \\ \cup \{ \quad c \notin \text{Ic}(N) \cup \text{Tc}(N) \quad | \quad 2 : c \in C_N \end{array} \right\}. \end{aligned}$$

For any marking C of control conditions of the net K and marking C' of control conditions of the net N , define

$$\begin{aligned} \theta_K(C) &\stackrel{\text{def}}{=} P \triangleleft 1 : C \\ \theta_N(C) &\stackrel{\text{def}}{=} P \triangleright 2 : C. \end{aligned}$$

For an event e of $K[N]$, define $\rho_N(e) = e'$ for the unique e' such that $e = (P_i \cup P_t) \triangleright 2 : e'$. For an event e of N , define $\theta_N(e) = (P_i \cup P_t) \triangleright 2 : e$. Define $\rho_K(e)$ and $\theta_K(e)$ similarly, apart from having $\theta_K([-])$ undefined.

Lemma A.4. *For any marking C of control conditions of K , the marking $\theta_K(C)$ is a K -marking in $K[N]$. For any marking C' of control conditions of N , the marking $\theta_N(C')$ is an N -marking in $K[N]$.*

Proof. Immediate from the definitions. □

It is clear that ρ_N and θ_N form a bijection between N -events and $\text{Ev}(N)$. It is also clear that ρ_K and θ_K form a bijection between K -events and $\text{Ev}(K) \setminus \{[-]\}$. On markings, the situation is a little more intricate:

Lemma A.5. *Let $K[N]$ be a substitution. For any marking of control conditions $C_K \cup C_N$ of $K[N]$ that is a K -marking and any marking C of control conditions of K :*

$$\theta_K(\rho_K(C_K \cup C_N)) = C_K \cup C_N \quad \text{and} \quad \rho_K(\theta_K(C)) = C.$$

For any marking of control conditions $C_K \cup C_N$ of $K[N]$ that is an N -marking and any marking C of control conditions of N :

$$\theta_N(\rho_N(C_N)) = C_N \quad \text{and} \quad \rho_N(\theta_N(C)) = C.$$

Proof. First, let C be any marking of control conditions of K . We shall show that $\rho_K(\theta_K(C)) = C$. Let c be any control condition of the net K . Since K is an embedded net, by the restrictions imposed in Lemma 3.11 there are three distinct cases: $c \notin \bullet[-]^\bullet$, $c \in \bullet[-]$ or $c \in [-]^\bullet$. The first case is straightforward since the operation of θ_K on such conditions is to add a ‘1:’-tag which is removed by ρ_K . Now consider $c \in \bullet[-]$; the case for $c \in [-]^\bullet$ will be similar. By the definition of θ_K , since $\text{Ic}(N)$ is nonempty (again by Lemma 3.11):

$$c \in C \quad \text{iff} \quad \forall i \in \text{Ic}(N). (1:c, 2:i) \in \theta_K(C).$$

From the definition of ρ_K , we have $\forall i \in \text{Ic}(N). (1:c, 2:i) \in \theta_K(C)$ iff $c \in \rho_K(\theta_K(C))$. So $c \in C$ iff $c \in \rho_K(\theta_K(C))$.

Now suppose that (C_K, C_N) is a K -marking of the substitution $K[N]$. Let c be any condition of the net $K[N]$. There are three distinct possible cases: $c \notin P_1 \cup P_t$, $c \in P_1$ or $c \in P_t$. First, suppose that $c \notin P_1 \cup P_t$:

$$\begin{aligned} c \in C_N \cup C_K & \text{ iff } c \in C_K && \text{(def. of } K\text{-marking)} \\ & \text{ iff } \exists c_1. (c_1 \in \rho_K(C_N, C_K) \text{ and } c = 1:c_1) && \text{(def. of } \rho_K) \\ & \text{ iff } c \in \theta_K(\rho_K(C_N \cup C_K)) && \text{(def. of } \theta_K) \end{aligned}$$

Now suppose that $c \in P_1$, so $c = (1:a, 2:i)$ for some $a \in \bullet[-]$ and $i \in \text{Ic}(N)$:

$$\begin{aligned} c \in C_N \cup C_K & \text{ iff } \forall i' \in \text{Ic}(N). (1:a, 2i') \in C_N \cup C_K && \text{(def. of } K\text{-marking)} \\ & \text{ iff } a \in \rho_K(C_N \cup C_K) && \text{(def. of } \rho_K) \\ & \text{ iff } c \in \theta_K(\rho_K(C_N \cup C_K)) && \text{(def. of } \theta_K) \end{aligned}$$

We have a similar analysis if $c \in P_t$. Hence $(C_K, C_N) = \theta_K(\rho_K(C_K, C_N))$.

For any marking of control conditions C of the net N and any N -marking (C_K, C_N) ,

$$\theta_N(\rho_N(C_N)) = C_N \quad \text{and} \quad \rho_N(\theta_N(C)) = C$$

are shown similarly, this time with the first analysis considering conditions in $\text{Ic}(N)$, $\text{Tc}(N)$ and conditions not in either set. \square

Lemma A.6. *Let (C_K, C_N, σ) and (C'_K, C'_N, σ') be markings of $K[N]$. Suppose that e is an event such that $(C_K, C_N, \sigma) \xrightarrow{e} (C'_K, C'_N, \sigma')$.*

(1) *If e is a K -event and (C_K, C_N) and (C'_K, C'_N) are K -markings then*

$$(\rho_K(C_K, C_N), \sigma) \xrightarrow{\rho_K(e)} (\rho_K(C'_K, C'_N), \sigma') \text{ in } K.$$

(2) *If e is an N -event and (C_K, C_N) and (C'_K, C'_N) are N -markings then $C_K = C'_K$ and*

$$(\rho_N(C_N), \sigma) \xrightarrow{\rho_N(e)} (\rho_N(C'_N), \sigma') \text{ in } N.$$

Proof. First consider (1). The event e is a K -event, so there is an event e_1 of K such that $e_1 \neq [-]$ and $e = (P_1 \cup P_t) \triangleleft 1:e_1$. We have

$$(C_N, C_K, \sigma) \xrightarrow{e} (C'_N, C'_K, \sigma')$$

in $K[N]$. By Lemma A.5, we have $\theta_K(\rho_K(C_N, C_K)) = (C_N, C_K)$ and $\theta_K(\rho_K(C'_N, C'_K)) = (C'_N, C'_K)$. From the definition of θ_K , we therefore have

$$((P_1 \cup P_t) \triangleleft 1: \rho_K(C_N, C_K), \sigma) \xrightarrow{(P_1 \cup P_t) \triangleleft 1: e_1} ((P_1 \cup P_t) \triangleleft 1: \rho_K(C'_N, C'_K), \sigma').$$

Using Lemma 3.6, we may therefore conclude that

$$(\rho_K(C_N, C_K), \sigma) \xrightarrow{e_1} (\rho_K(C'_N, C'_K), \sigma')$$

in K . The proof of (2) is similar. \square

Lemma A.7. (1) Let C and C' be markings of control conditions of K . If $(C, \sigma) \xrightarrow{e} (C', \sigma')$ in K then $(\theta_K(C), \sigma) \xrightarrow{\theta_K(e)} (\theta_K(C'), \sigma')$ in $K[N]$.
(2) Now let C and C' be markings of control conditions of N . If $(C, \sigma) \xrightarrow{e} (C', \sigma')$ in N then $(\theta_N(C), C_K, \sigma) \xrightarrow{\theta_N(e)} (\theta_N(C'), C_K, \sigma')$ in $K[N]$ for any marking C_K of K -conditions.

Proof. First consider (1). Suppose that $(C, \sigma) \xrightarrow{e} (C', \sigma')$ in K for some event $e \neq [-]$. By Lemma 3.6, we have

$$((P_i \cup P_t) \triangleleft 1 : C, \sigma) \xrightarrow{(P_i \cup P_t) \triangleleft 1 : e} ((P_i \cup P_t) \triangleleft 1 : C', \sigma')$$

in $K[N]$. Since $\theta_K(C) = (P_i \cup P_t) \triangleleft 1 : C$, and similarly for C' and e , we therefore have

$$(\theta_K(C), \sigma) \xrightarrow{\theta_K(e)} (\theta_K(C'), \sigma'),$$

as required. The proof of (2) is similar. \square

We are now able to characterize the runs of the net $K[N]$ when a non-interfering substitution is formed.

Lemma A.8. Let $K[N]$ be a non-interfering substitution from σ_0 . Any complete sequence π from $(\text{Ic}(K[N]), \sigma_0)$ is of the form $\Pi_0 \cdot (\Pi_1 \cdot \Pi_0)^*$, where:

- Π_0 ranges over sequences consisting of K -events between K -markings.
- Π_1 ranges over nonempty sequences π_1 of any events between N -markings, where no K -event uses any condition in P_i or P_t . If $(C_N \cup C_K, \sigma)$ and $(C'_N \cup C'_K, \sigma')$ are the initial and final markings of π_1 , respectively, then $C_N = P_i$ and $C'_N = P_t$. The first event of π_1 is an N -event and the final event of π_1 is also an N -event.

Proof. We first show that any sequence π in N from $(\text{Ic}(K[N]), \sigma_0)$ is of the form $\Pi_0 \cdot (\Pi_1 \cdot \Pi_0)^*$ or $\Pi_0 \cdot (\Pi_1 \cdot \Pi_0)^* \cdot \Pi'_1$ by induction on the length of sequence, where a sequence is of form Π'_1 if:

- it is a sequence of K - and N -events between N -markings where no K -event uses any condition in P_i or P_t , and
- if (C_N, C_K, σ) is the initial marking of π_1 then $C_N = P_i$, and the first event of π_1 is an N -event.

We shall simultaneously show that if $\pi : (\text{Ic}(K[N]), \sigma_0) \xrightarrow{*} (C_N, C_K, \sigma)$ and (C_N, C_K, σ) is an N -marking then either it is N -active or $C_N = P_t$. Furthermore, if $P_t \subseteq C_N$ then $P_t = C_N$.

The base case for the induction is straightforward. Suppose that $\pi : (\text{Ic}(K[N]), \sigma_0) \xrightarrow{*} M$ where $M = (C_N, C_K, \sigma)$ and that e is an event such that $M \xrightarrow{e} M'$. Let $M' = (C'_N, C'_K, \sigma')$. We shall show that $\pi \cdot e$ from marking $(\text{Ic}(K[N]), \sigma_0)$ is of the correct form and that M' satisfies the required properties.

Suppose that M' is an N -marking but $C'_N \neq P_t$ and M' is not N -active. As M' is not N -active, we must have $C'_N \neq P_i$. From the induction hypothesis, there must exist a path π' and markings C''_K and σ'' such that

$$\pi' : (P_i, C''_K, \sigma'') \xrightarrow{*} (C'_N, C'_K, \sigma')$$

and (P_i, C''_K, σ'') is reachable from $(\text{Ic}(K[N]), \sigma_0)$. Furthermore, from (P_i, C''_K, σ'') the path π' is between N -active markings. Since $K[N]$ is a non-interfering substitution from state

σ_0 , it follows from the requirement that consecutive K - and N -events must be independent that there must exist paths π_1 and π_2 made exclusively of N - and K -events, respectively, such that $\pi_1 \cdot \pi_2 : (P_i, C''_K, \sigma'') \longrightarrow^* (C'_N, C'_K, \sigma')$. Since N -events do not affect the marking of K -conditions and from the requirement that K -events do not affect the marking of N -conditions along the path π_2 because $K[N]$ is a non-interfering substitution from σ_0 , there exists a state σ_1 such that $\pi_1 : (P_i, C''_K, \sigma'') \longrightarrow^* (C'_N, C''_K, \sigma_1)$. Since $\rho_N(P_i) = \text{Ic}(N)$, a simple induction on the length of this sequence using Lemma A.6 shows that the marking $(\rho_N(C'_N), \sigma_1)$ is reachable from $(\text{Ic}(N), \sigma'')$ in N . Consider the ways in which the N -marking (C'_N, C'_K, σ') may fail to be N -active: Firstly, if $C'_N \subsetneq P_i$, it follows that $\rho_N(C'_N) \subsetneq \text{Ic}(N)$. Since $(\rho_N(C'_N), \sigma_1)$ is reachable from $(\text{Ic}(N), \sigma')$, this contradicts the requirement of Definition 3.13. The proof is similar in the other cases, $C'_N \subsetneq P_{\dagger}$ and $P_{\dagger} \subsetneq C'_N$, which may cause the marking to fail to be N -active without $C'_N = P_{\dagger}$.

To complete the proof, it suffices to show the following properties:

- (1) K -events preserve K -markings: If M is a K -marking and e is a K -event and $M \xrightarrow{e} M'$ then M' is a K -marking.
- (2) N -events preserve N -markings: If M is an N -marking and e is an N -event and $M \xrightarrow{e} M'$ then M' is an N -marking.
- (3) If e is a K -event with no pre- or postcondition inside $P_i \cup P_{\dagger}$ and M is an N -marking and $M \xrightarrow{e} M'$ then M' is an N -marking.
- (4) The only markings that are both N - and K -markings are of the form (P_i, C_K, σ) or $(P_{\dagger}, C_K, \sigma)$ or $(P_i \cup P_{\dagger}, C_K, \sigma)$ for some C_K and σ .
- (5) No N -event has concession in any reachable marking that is not N -active.

Properties (1) and (2) are straightforward calculations using Lemmas A.4, A.5 and A.6. Property (3) follows immediately from Lemma A.1. Property (4) is obvious from the definitions of N - and K -markings. Property (5) is straightforward from the induction hypotheses and the fact that no event has concession in the terminal marking of N according to the requirements of Lemma 3.11.

Finally, to see that any complete run is of the form $\Pi_0 \cdot (\Pi_1 \cdot \Pi_0)^*$, observe that the terminal marking of control conditions $\text{Tc}(K[N])$ is a K -marking. There are no C_K and σ such that the marking (P_i, C_K, σ) is terminal since then $\bullet[-] \cap \text{Tc}(K) \neq \emptyset$, contradicting the requirement that K should be an embedded net satisfying the requirements of Lemma 3.11. Hence the terminal marking is not N -active. \square

Having now dealt with the control structure of contexts, we return to the idea that, given a net $K[N_1]$ which is a non-interfering substitution from state σ , the events in any sequence may be reordered in a way that ensures that events of N_1 occur consecutively and form a “complete run” of the net N_1 . As $N_1 \simeq N_2$, the net $K[N_2]$ will therefore have a path between the same sets of state conditions.

To formalize this, let π be any sequential run of a non-interfering substitution $K[N]$ from marking M . The set $\mathcal{P}_{K[N]}(\pi, M)$ is defined to be the least set of sequences from marking M of $K[N]$ closed under the operation of swapping consecutive independent events that contains the sequence π . It is easy to see that if $\pi : M \longrightarrow^* M'$ and $\pi' \in \mathcal{P}_{K[N]}(\pi, M)$ then $\pi' : M \longrightarrow^* M'$ for any paths π and π' . Define the order \prec on $\mathcal{P}_{K[N]}(\pi, M)$ as follows:

Definition A.9. Let $\pi, \pi' \in \mathcal{P}_{K[N]}(\pi_0, M)$. Define \prec to be the transitive closure of \prec_1 , where $\pi \prec_1 \pi'$ iff there exist sequences π_1 and π_2 , an N -event e and a K -event e' such that eIe' and $\pi = \pi_1 \cdot e \cdot e' \cdot \pi_2$ and $\pi' = \pi_1 \cdot e' \cdot e \cdot \pi_2$.

It is clear that the order \prec is well-founded since any path is, by definition, of finite length.

Definition A.10. Say that a sequence π of $K[N]$ from marking M is N -complete if $M = (P_i, C_K, \sigma)$ for some C_K and σ , every event of π is an N -event, and

$$\pi : (P_i, C_K, \sigma) \longrightarrow^* (P_t, C_K, \sigma').$$

Lemma A.11. Let $K[N]$ be a non-interfering substitution from state σ_0 and let $M_0 = (\text{Ic}(K[N]), \sigma_0)$. Suppose that π_0 is a complete sequence of $K[N]$ from M_0 . The \prec -minimal elements of $\mathcal{P}_{K[N]}(\pi_0, M_0)$ are of the form

$$\Pi_0 \cdot (\Pi_N \cdot \Pi_0)^*,$$

where Π_N matches N -complete paths and Π_0 is as in Lemma A.8.

Proof. Suppose that π is a \prec -minimal element of $\mathcal{P}_{K[N]}(\pi_0, M_0)$ but not of the form above. The sequence π is of the form of Lemma A.8 because π is a complete path of $K[N]$. Consequently, there are π_1, π_2 and π_3 such that $\pi = \pi_1 \cdot \pi_2 \cdot \pi_3$ and $\pi_2 = (e \cdot e')$ where e is a K -event and e' is an N -event. Furthermore, the marking M_1 such that $\pi_1 : M_0 \longrightarrow^* M_1$ is N -active. Now, from the definition of non-interfering substitution, the events e and e' are independent. Hence the sequence $\pi_1 \cdot e' \cdot e \cdot \pi_3$ is in $\mathcal{P}_{K[N]}(\pi_0, M_0)$ and is beneath π , contradicting its minimality. \square

This gives us the ability to prove Theorem 6.7 by induction on paths of $K[N_1]$.

Theorem A.12. If $K[N_1]$ and $K[N_2]$ are non-interfering substitutions from σ_0 and $N_1 \simeq N_2$ then, for all states σ :

$$K[N_1] : \sigma_0 \Downarrow \sigma \quad \text{iff} \quad K[N_2] : \sigma_0 \Downarrow \sigma.$$

Proof. Suppose that π is a complete sequence of $K[N_1]$ from σ_0 to σ' . We shall show that, for all $\pi_1 \in \mathcal{P}_{K[N_1]}(\pi, (\text{Ic}(K[N_1]), \sigma_0))$, if π_1 is a complete sequence from σ_0 to σ' then there exists a complete sequence π_2 of $K[N_2]$ from σ_0 to σ' . The proof shall proceed by induction on the well-founded order \prec . In particular $\pi \in \mathcal{P}_{K[N_1]}(\pi, (\text{Ic}(K[N]), \sigma_0))$, so, with the symmetric proof for the other direction, this will complete the proof of the required property.

(π_1 MINIMAL) The sequence π_1 is minimal within $\mathcal{P}_{K[N_1]}(\pi, \sigma_0)$, so, by Lemma A.11, there exists an $n \in \mathbb{N}$ such that there exist sequences $\pi_0, \pi_{01}, \pi_{11}, \dots, \pi_{0n}, \pi_{1n}$ with

$$\pi_1 = \pi_0 \cdot \pi_{11} \cdot \pi_{01} \dots \pi_{1n} \cdot \pi_{0n}.$$

Furthermore, for each $i \leq n$, the sequence π_{0i} is of the form Π_0 defined in Lemma A.8, as is the sequence π_0 ; and, for each $i \leq n$, the sequence π_{1i} is of the form Π_{N_1} , which matches N_1 -complete subpaths of $K[N_1]$ as defined in Definition A.10. Define:

$$\begin{array}{ll} P_i^{(1)} \stackrel{\text{def}}{=} 1 : \bullet[-] \times 2 : \text{Ic}(N_1) & P_i^{(2)} \stackrel{\text{def}}{=} 1 : \bullet[-] \times 2 : \text{Ic}(N_2) \\ P_t^{(1)} \stackrel{\text{def}}{=} 1 : [-] \bullet \times 2 : \text{Tc}(N_1) & P_t^{(2)} \stackrel{\text{def}}{=} 1 : [-] \bullet \times 2 : \text{Tc}(N_2). \end{array}$$

Let $\rho_K^{(1)}$ be ρ_K from Definition A.3 for $K[N_1]$ and let $\rho_K^{(2)}$ be ρ_K from Definition A.3 for $K[N_2]$, and similarly for $\rho_{N_1}^{(1)}, \rho_{N_2}^{(2)}, \theta_K^{(1)}$, etc. We shall show, by induction on n , that if π_1 is a sequence of this form in $K[N_1]$ from $(\text{Ic}(K[N_1]), \sigma_0)$ to the marking (C'_1, σ') then there exists a path π_2 from $(\text{Ic}(K[N_2]), \sigma_0)$ to (C'_2, σ') for some C'_2 such that $\rho_K^{(1)}(C'_1) = \rho_K^{(2)}(C'_2)$.

- $n = 0$: Assume that π_1 is of the form Π_0 . Let $\pi_1 = (e_1 \cdot \dots \cdot e_m)$ and suppose that in $K[N_1]$ we have

$$(\text{Ic}(K[N_1]), \sigma_0) \xrightarrow{e_1} (C_1, \sigma_1) \xrightarrow{e_2} \dots \xrightarrow{e_m} (C_m, \sigma_m).$$

By assumption, π_1 is a path from $(\text{Ic}(K[N_1]), \sigma_0)$ to (C'_1, σ') , so $C'_1 = C_m$ and $\sigma' = \sigma_m$. Now, $\text{Ic}(K[N_1])$ is a K -marking, and, since π_1 is of the form Π_0 , for every i such that $0 < i \leq m$, the marking C_i is a K -marking and e_i is a K -event. By Lemma A.6, in the net K we have

$$(\rho_K^{(1)}(\text{Ic}(K[N_1])), \sigma) \xrightarrow{\rho_K^{(1)}(e_1)} (\rho_K^{(1)}(C_1), \sigma_1) \xrightarrow{\rho_K^{(1)}(e_2)} \dots \xrightarrow{\rho_K^{(1)}(e_m)} (\rho_K^{(1)}(C_m), \sigma_m).$$

In the net $K[N_2]$, by Lemma A.7, we therefore have

$$\begin{array}{ccc} (\theta_K^{(2)} \rho_K^{(1)}(\text{Ic}(K[N_1])), \sigma) & \xrightarrow{\theta_K^{(2)} \rho_K^{(1)}(e_1)} & (\theta_K^{(2)} \rho_K^{(1)}(C_1), \sigma_1) \\ & \xrightarrow{\theta_K^{(2)} \rho_K^{(1)}(e_2)} & \dots \\ & \xrightarrow{\theta_K^{(2)} \rho_K^{(1)}(e_m)} & (\theta_K^{(2)} \rho_K^{(1)}(C_m), \sigma_m). \end{array}$$

Let $C'_2 = \theta_K^{(2)} \rho_K^{(1)}(C_m)$. From Lemma A.4, $\theta_K^{(2)}$ generates K -markings of $K[N_2]$ from markings of K . By Lemma A.5, we therefore have $\rho_K^{(2)}(C'_2) = \rho_K^{(1)}(C'_1)$ since $C'_1 = C_m$, which is a K -marking. It is an easy calculation to show that $\rho_K^{(1)}(\text{Ic}(K[N_1])) = \text{Ic}(K)$ and $\theta_K^{(2)}(\text{Ic}(K)) = \text{Ic}(K[N_2])$. There therefore exists a path from $(\text{Ic}(K[N_2]), \sigma_0)$ to (C'_2, σ_m) in $K[N_2]$ and $\rho_K^{(2)}(C'_2) = \rho_K^{(1)}(C'_1)$, which is all that is required since $\sigma_m = \sigma'$.

- $n > 0$: Assume that $\pi_1 = \pi_{11} \cdot \pi_{12} \cdot \pi_{13}$ for some sequence π_{11} of form $\Pi_0 \cdot (\Pi_{N_1} \cdot \Pi_0)^{n-1}$, some sequence π_{12} of form Π_{N_1} and some sequence π_{13} of form Π_0 . Let (C'_1, σ') be the marking obtained by following π_1 from $(\text{Ic}(K[N_1]), \sigma_0)$ in $K[N_1]$. We wish to show that there is a path π_2 of $K[N_2]$ from $(\text{Ic}(K[N_2]), \sigma_0)$ to (C'_2, σ') for some C'_2 such that $\rho_K^{(1)}(C'_1) = \rho_K^{(2)}(C'_2)$.

Let (C_{11}, σ_1) be the marking obtained by following path π_{11} from $(\text{Ic}(K[N_1]), \sigma_0)$. Since π_{12} follows π_{11} and π_{12} is of form Π_{N_1} , it must be the case that $C_{11} = (P_i^{(1)}, C_K)$ for some marking C_K of K -conditions.

By induction, there is a path π_{21} in $K[N_2]$ from $(\text{Ic}(K[N_2]), \sigma_0)$ to (C_{21}, σ_1) for some C_{21} such that $\rho_K^{(1)}(C_{11}) = \rho_K^{(2)}(C_{21})$. Now, $C_{11} = (P_i^{(1)}, C_K)$, so $\rho_K^{(1)}(C_{11}) = \bullet[-] \cup \{c \mid 1:c \in C_K\}$. From the definition of $\rho_K^{(2)}$, we must therefore have $C_{21} = (P_i^{(2)}, C_K)$. Hence

$$\pi_{21} : (\text{Ic}(K[N_2]), \sigma_0) \xrightarrow{*} (P_i^{(2)}, C_K, \sigma_1).$$

Suppose that in $K[N_1]$ we have $\pi_{12} : (P_i^{(1)}, C_K, \sigma_1) \xrightarrow{*} (C'_{N_1}, C'_K, \sigma_2)$. Since π_{12} is of the form Π_{N_1} , it is an N_1 -complete path, so $C'_{N_1} = P_t^{(1)}$. The events of π_{12} are all N -events. Using Lemma A.6, a simple induction shows that $C_K = C'_K$ and that there is a path from $(\rho_{N_1}^{(1)}(P_i^{(1)}), \sigma_1)$ to $(\rho_{N_1}^{(1)}(P_t^{(1)}), \sigma_2)$ in N_1 . Observe that $\rho_{N_1}^{(1)}(P_i^{(1)}) = \text{Ic}(N_1)$ and $\rho_{N_1}^{(1)}(P_t^{(1)}) = \text{Tc}(N_1)$, so $N_1 : \sigma_1 \Downarrow \sigma_2$. As $N_1 \simeq N_2$, there is therefore a path of N_2 from $(\text{Ic}(N_2), \sigma_1)$ to $(\text{Tc}(N_2), \sigma_2)$. By Lemma A.7,

a simple induction on the length of this sequence shows that there is a sequence π_{22} from $(\theta_{N_2}^{(2)}(\text{Ic}(N_2)), C_K, \sigma_1)$ to $(\theta_{N_2}^{(2)}(\text{Tc}(N_2)), C_K, \sigma_2)$ in $K[N_2]$. Observe that $\theta_{N_2}^{(2)}(\text{Ic}(N_2)) = P_i^{(2)}$ and $\theta_{N_2}^{(2)}(\text{Tc}(N_2)) = P_t^{(2)}$, so

$$\pi_{22} : (P_i^{(2)}, C_K, \sigma_1) \longrightarrow^* (P_t^{(2)}, C_K, \sigma_2).$$

As π_{13} follows path π_{12} in π_1 , the sequence π_{13} is from $(P_t^{(1)}, C_K, \sigma_2)$ to (C'_1, σ') and contains only K -events. Using Lemma A.6, a simple induction on the length of π_{13} shows that there is a path from $(\rho_K^{(1)}(P_t^{(1)}, C_K), \sigma_2)$ to $(\rho_K^{(1)}(C'_1), \sigma')$ in K . A simple induction on the length of this path, using Lemma A.7 shows that there is a path π_{23} of $K[N_2]$ such that $\pi_{23} : (\theta_K^{(2)} \rho_K^{(1)}(P_t^{(1)}, C_K), \sigma_2) \longrightarrow^* (\theta_K^{(2)} \rho_K^{(1)}(C'_1), \sigma')$. From the definition of $\rho_K^{(1)}$, we have $\rho_K^{(1)}(P_t^{(1)}, C_K) = [-]^\bullet \cup \{c \mid 1 : c \in C_K\}$. From the definition of $\theta_K^{(2)}$, we have $\theta_K^{(2)}([-]^\bullet \cup \{c \mid 1 : c \in C_K\}) = P_t^{(2)} \cup C_K$. Hence

$$\pi_{23} : (P_t^{(2)}, C_K, \sigma_2) \longrightarrow^* (\theta_K^{(2)} \rho_K^{(1)}(C'_1), \sigma').$$

Take $C'_2 = (\theta_K^{(2)} \rho_K^{(1)}(C'_1), \sigma')$. By Lemma A.5, we have $\rho_K^{(2)}(C'_2) = \rho_K^{(1)}(C'_1)$. Consequently, the path $\pi_2 = \pi_{21} \cdot \pi_{22} \cdot \pi_{23}$ satisfies

$$\pi_2 : (\text{Ic}(K[N_2]), \sigma_0) \longrightarrow^* (C'_2, \sigma'),$$

for some C'_2 such that $\rho_K(C'_1) = \rho_K(C'_2)$, which is all that is required to complete this inner induction.

Now, recall that π_1 is a complete sequence of $K[N_1]$, so

$$\pi_1 : (\text{Ic}(K[N_1]), \sigma_0) \longrightarrow^* (\text{Tc}(K[N_1]), \sigma').$$

From the immediately preceding induction, there exists a path π_2 of $K[N_2]$ such that $\pi_2 : (\text{Ic}(K[N_2]), \sigma_0) \longrightarrow^* (C'_2, \sigma')$ for some C'_2 such that $\rho_K^{(1)}(\text{Tc}(K[N_1])) = \rho_K^{(2)}(C'_2)$. Now, clearly $\rho_K^{(1)}(\text{Tc}(K[N_1])) = \text{Tc}(K)$ by the definitions of ρ and $K[N_1]$. Hence $\text{Tc}(K) = \rho_K^{(2)}(C'_2)$, so by Lemma A.5 we have $\theta_K^{(2)}(\text{Tc}(K)) = C'_2$. The definition of $K[N_2]$ and θ gives $\theta_K^{(2)}(\text{Tc}(K)) = \text{Tc}(K[N_2])$. Hence

$$\pi_2 : (\text{Ic}(K[N_2]), \sigma_0) \longrightarrow^* (\text{Tc}(K[N_2]), \sigma'),$$

as required.

(π_1 NOT MINIMAL) Suppose that the path π_1 is not minimal and that π_1 is a complete path of $K[N_1]$ with $\pi_1 : (\text{Ic}(K[N_1]), \sigma_0) \longrightarrow^* (\text{Tc}(K[N_1]), \sigma')$. It is easy to see that the order \prec is irreflexive, so there exists a path π'_1 such that $\pi'_1 \prec_1 \pi_1$. Hence there exist paths π_2 and π_3 and a K -event e and an N -event e' such that $\pi_1 = \pi_2 \cdot e \cdot e' \cdot \pi_3$ and $\pi'_1 = \pi_2 \cdot e' \cdot e \cdot \pi_3$. Furthermore, the events e and e' are independent, so π'_1 must also be a path $\pi'_1 : (\text{Ic}(K[N_1]), \sigma_0) \longrightarrow^* (\text{Tc}(K[N_1]), \sigma')$. By induction, there exists a path $\pi'_2 : (\text{Ic}(K[N_1]), \sigma_0) \longrightarrow^* (\text{Tc}(K[N_1]), \sigma')$, as required to complete the case.

Hence, if $K[N_1] : \sigma_0 \Downarrow \sigma'$, there exists a path $\pi : (\text{Ic}(K[N_1]), \sigma_0) \longrightarrow^* (\text{Tc}(K[N_1]), \sigma')$ in $K[N_1]$. Since $\pi \in \mathcal{P}_{K[N_1]}(\pi, \sigma_0)$, we have a path $\pi_2 : (\text{Ic}(K[N_2]), \sigma_0) \longrightarrow^* (\text{Tc}(K[N_2]), \sigma')$ in $K[N_2]$, so $K[N_2] : \sigma_0 \Downarrow \sigma'$. The proof for the reverse implication is symmetric. \square