

Reactive Temporal Logic

Rob van Glabbeek

Data61, CSIRO, Sydney, Australia

School of Computer Science and Engineering, University of New South Wales, Sydney, Australia

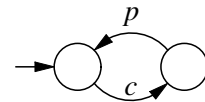
rvg@cs.stanford.edu

Whereas standard treatments of temporal logic are adequate for *closed systems*, having no run-time interactions with their environment, they fall short for *reactive systems*, interacting with their environments through synchronisation of actions. This paper introduces *reactive temporal logic*, a form of temporal logic adapted for the study of reactive systems. I illustrate its use by applying it to formulate definitions of a fair scheduler, and of a correct mutual exclusion protocol. Previous definitions of these concepts were conceptually much more involved or less precise, leading to debates on whether or not a given protocol satisfies the implicit requirements.

1 Introduction

Labelled transition systems are a common model of distributed systems. They consist of sets of states, also called *processes*, and transitions—each transition going from a source state to a target state. A given distributed system \mathcal{D} corresponds to a state P in a transition system \mathbb{T} —the initial state of \mathcal{D} . The other states of \mathcal{D} are the processes in \mathbb{T} that are reachable from P by following the transitions. The transitions are labelled by *actions*, either visible ones or the invisible action τ . Whereas a τ -labelled transition represents a state-change that can be made spontaneously by the represented system, a -labelled transitions for $a \neq \tau$ merely represent potential activities of \mathcal{D} , for they require cooperation from the *environment* in which \mathcal{D} will be running, sometimes identified with the *user* of system \mathcal{D} . A typical example is the acceptance of a coin by a vending machine. For this transition to occur, the vending machine should be in a state where it is enabled, i.e., the opening for inserting coins should not be closed off, but also the user of the system should partake by inserting the coin.

Consider a vending machine that alternately accepts a coin (c) and produces a pretzel (p). Its labelled transition system is depicted on the right. In standard temporal logic one can express that each action c is followed by p : whenever a



coin is inserted, a pretzel will be produced. Aligned with intuition, this formula is valid for the depicted system. However, by symmetry one obtains the validity of a formula saying that each p is followed by a c : whenever a pretzel is produced, eventually a new coin will be inserted. But that clashes with intuition.

In this paper I enrich temporal logic judgements $P \models \varphi$, saying that system P satisfies formula φ , with a third argument B , telling which actions can be blocked by the environment (by failing to act as a synchronisation partner) and which cannot. When stipulating that the coin needs cooperation from a user, but producing the pretzel does not, the two temporal judgements can be distinguished, and only one of them holds. I also introduce a fourth argument CC —a completeness criterion—that incorporates progress, justness and fairness assumptions employed when making a temporal judgement. This yields statements of the form $P \models_B^{CC} \varphi$.

Then I use the so obtained formalism to formalise the correctness requirements of mutual exclusion protocols and of fair schedulers. Making these requirements precise helps in stating negative results on the possibilities to render such protocols in a given setting. In the case of fair schedulers, reactive

temporal logic leads to a much easier to understand formalisation than the one in the literature. In the case of mutual exclusion protocols it leads to more precise and less ambiguous requirements, that may help to settle debates on whether or not some formalisation of a mutual exclusion protocol is correct.

2 Kripke Structures and Linear-time Temporal Logic

Definition 1 Let AP be a set of *atomic predicates*. A *Kripke structure* over AP is tuple $(S, \rightarrow, \models)$ with S a set (of *states*), $\rightarrow \subseteq S \times S$, the *transition relation*, and $\models \subseteq S \times AP$. $s \models p$ says that predicate $p \in AP$ holds in state $s \in S$.

Here I generalise the standard definition [14] by dropping the condition of *totality*, requiring that for each state $s \in S$ there is a transition $(s, s') \in \rightarrow$. A *path* in a Kripke structure is a nonempty finite or infinite sequence s_0, s_1, \dots of states, such that $(s_i, s_{i+1}) \in \rightarrow$ for each adjacent pair of states s_i, s_{i+1} in that sequence. A *suffix* π' of a path π is any path obtained from π by removing an initial segment. Write $\pi \Rightarrow \pi'$ if π' is a suffix of π ; this relation is reflexive and transitive.

A distributed system \mathcal{D} can be modelled as a state s in a Kripke structure K . A run of \mathcal{D} then corresponds with a path in K starting in s . Whereas each finite path in K starting from s models a *partial run* of \mathcal{D} , i.e., an initial segment of a (complete) run, typically not each path models a run. Therefore a Kripke structure constitutes a good model of distributed systems only in combination with a *completeness criterion* [9]: a selection of a set of paths as *complete paths*, modelling runs of the represented system.

The default completeness criterion, implicitly used in almost all work on temporal logic, classifies a path as complete iff it is infinite. In other words, only the infinite paths, and all of them, model (complete) runs of the represented system. This applies when adopting the condition of totality, so that each finite path is a prefix of an infinite path. Naturally, in this setting there is no reason to use the word “complete”, as “infinite” will do. As I plan to discuss alternative completeness criteria in Section 4, I will here already refer to paths satisfying a completeness criterion as “complete” rather than “infinite”. Moreover, when dropping totality, the default completeness criterion is adapted to declare a path complete iff it either is infinite or ends in a state without outgoing transitions [1].

Linear-time temporal logic (LTL) [23, 14] is a formalism explicitly designed to formulate properties such as the safety and liveness requirements of mutual exclusion protocols. Its syntax is

$$\varphi, \psi ::= p \mid \neg\varphi \mid \varphi \wedge \psi \mid \mathbf{X}\varphi \mid \mathbf{F}\varphi \mid \mathbf{G}\varphi \mid \psi\mathbf{U}\varphi$$

with $p \in AP$ an atomic predicate. The propositional connectives \Rightarrow and \vee can be added as syntactic sugar. It is interpreted on the paths in a Kripke structure. The relation \models between paths and LTL formulae, with $\pi \models \varphi$ saying that the path π *satisfies* the formula φ , or that φ is *valid* on π , is inductively defined by

- $\pi \models p$, with $p \in AP$, iff $s \models p$, where s is the first state of π ,
- $\pi \models \neg\varphi$ iff $\pi \not\models \varphi$,
- $\pi \models \varphi \wedge \psi$ iff $\pi \models \varphi$ and $\pi \models \psi$,
- $\pi \models \mathbf{X}\varphi$ iff $\pi' \models \varphi$, where π' is the suffix of π obtained by omitting the first state,
- $\pi \models \mathbf{F}\varphi$ iff $\pi' \models \varphi$ for some suffix π' of π ,
- $\pi \models \mathbf{G}\varphi$ iff $\pi' \models \varphi$ for each suffix π' of π , and
- $\pi \models \psi\mathbf{U}\varphi$ iff $\pi' \models \varphi$ for some suffix π' of π , and $\pi'' \models \psi$ for each path $\pi'' \neq \pi'$ with $\pi \Rightarrow \pi'' \Rightarrow \pi'$.

In [18], Lamport argues against the use of the next-state operator \mathbf{X} , as it is incompatible with abstraction from irrelevant details in system descriptions. Following this advice, I here restrict attention to LTL without the next-state modality, $\text{LTL}_{\mathbf{X}}$.

In the standard treatment of LTL [23, 14], judgements $\pi \models \varphi$ are pronounced only for infinite paths π . Here I apply the same definitions verbatim to finite paths as well. At this point I benefit from the exclusion of the next-state operator \mathbf{X} . In its presence I would have to decide what is the meaning of a judgement $\pi \models \mathbf{X}\varphi$ when π is a path consisting of a single state.¹

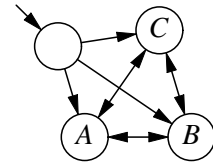
Having given meaning to judgements $\pi \models \varphi$, as a derived concept one defines when an $\text{LTL}_{\mathbf{X}}$ formula φ holds for a state s in a Kripke structure, modelling a distributed system \mathcal{D} , notation $s \models \varphi$ or $\mathcal{D} \models \varphi$. This is the case iff φ holds for all runs of \mathcal{D} .

Definition 2 $s \models \varphi$ iff $\pi \models \varphi$ for all complete paths π starting in state s .

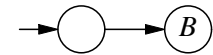
Note that this definition depends on the underlying completeness criterion, telling which paths model actual system runs. In situations where I consider different completeness criteria, I make this explicit by writing $s \models^{CC} \varphi$, with CC the name of the completeness criterion used. When leaving out the superscript CC I here refer to the default completeness criterion, defined above.

Example 1 Alice, Bart and Cameron stand behind a bar, continuously ordering and drinking beer. Assume they do not know each other and order individually. As there is only one barman, they are served sequentially. Also assume that none of them is served twice in a row, but as it takes no longer to drink a beer than to pour it, each of them is ready for the next beer as soon as another person is served.

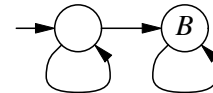
A Kripke structure of this distributed system \mathcal{D} is drawn on the right. The initial state of \mathcal{D} is indicated by a short arrow. The other three states are labelled with the atomic predicates A , B and C , indicating that Alice, Bart or Cameron, respectively, has just acquired a beer. When assuming the default completeness criterion, valid $\text{LTL}_{\mathbf{X}}$ formulae are $\mathbf{F}(A \vee C)$, saying that eventually either Alice or Cameron will get a beer, or $\mathbf{G}(A \Rightarrow \mathbf{F}\neg A)$, saying that each time Alice got a beer is followed eventually by someone else getting one. However, it is not guaranteed that Bart will ever get a beer: $\mathcal{D} \not\models \mathbf{F}B$. A counterexample for this formula is the infinite run in which Alice and Cameron get a beer alternatingly.



Example 2 Bart is the only customer in a bar in London, with a single barman. He only wants one beer. A Kripke structure of this system \mathcal{E} is drawn on the right. When assuming the default completeness criterion, this time Bart gets his beer: $\mathcal{E} \models \mathbf{F}B$.



Example 3 Bart is the only customer in a bar in London, with a single barman. He only wants one beer. At the same time, Alice and Cameron are in a bar in Tokyo. They drink a lot of beer. Bart is not in contact with Alice and Cameron, nor is there any connection between the two bars. Yet, one may choose to model the drinking in these two bars as a single distributed system. A Kripke structure of this system \mathcal{F} is drawn on the right, collapsing the orders of Alice and Cameron, which can occur before or after Bart gets a beer, into self-loops. When assuming the default completeness criterion, Bart cannot count on a beer: $\mathcal{F} \not\models \mathbf{F}B$.



¹One possibility would be to declare this judgement to be false, regardless of φ . However, this would invalidate the self-duality of the \mathbf{X} modality, stating that $\neg\mathbf{X}\varphi$ holds for the same paths as $\mathbf{X}\neg\varphi$.

3 Labelled Transition Systems, Process Algebra and Petri Nets

The most common formalisms in which to present reactive distributed systems are pseudocode, process algebra and Petri nets. The semantics of these formalisms is often given by translation into labelled transition systems (LTSs), and these in turn can be translated into Kripke structures, on which temporal formulae from languages such as LTL are interpreted. These translations make the validity relation \models for temporal formulae applicable to all these formalisms. A state in an LTS, for example, is defined to satisfy an $LTL_{\mathbf{x}}$ formula φ iff its translation into a state in a Kripke structure satisfies this formula.

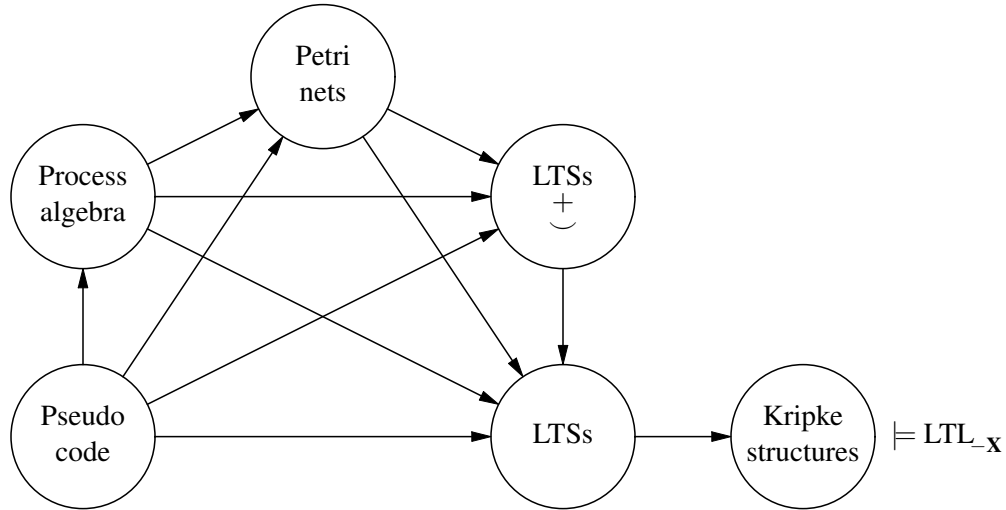


Figure 1: *Formalisms for modelling mutual exclusion protocols*

Figure 1 shows a commuting diagram of semantic translations found in the literature, from pseudocode, process algebra and Petri nets via LTSs to Kripke structures. Each step in the translation abstracts from certain features of the formalism at its source. Some useful requirements on distributed systems can be adequately formalised in process algebra or Petri nets, and informally described for pseudocode, whereas LTSs and Kripke structures have already abstracted from the relevant information. An example will be FS 1 on page 64. I also consider LTSs upgraded with a concurrency relation \smile between transitions; these will be expressive enough to formalise some of these requirements.

3.1 Labelled Transition Systems

Definition 3 Let A be a set of *observable actions*, and let $Act := A \cup \{\tau\}$, with $\tau \notin A$ the *hidden action*. A *labelled transition system* (LTS) over Act is tuple $(\mathbb{P}, Tr, source, target, \ell)$ with \mathbb{P} a set (of *states* or *processes*), Tr a set (of *transitions*), $source, target : Tr \rightarrow \mathbb{P}$ and $\ell : Tr \rightarrow Act$.

Write $s \xrightarrow{\alpha} s'$ if there exists a transition t with $source(t) = s \in \mathbb{P}$, $\ell(t) = \alpha \in Act$ and $target(t) = s' \in \mathbb{P}$. In this case t goes from s to s' , and is an *outgoing transition* of s . A *path* in an LTS is a finite or infinite alternating sequence of states and transitions, starting with a state, such that each transition goes from the state before it to the state after it (if any). A *completeness criterion* on an LTS is a set of its paths.

As for Kripke structures, a distributed system \mathcal{D} can be modelled as a state s in an LTS upgraded with a completeness criterion. A (complete) run of \mathcal{D} is then modelled by a complete path starting in s . As for Kripke structures, the default completeness criterion deems a path complete iff it either is infinite or ends

in a *deadlock*, a state without outgoing transitions. An alternative completeness criterion could declare some infinite paths incomplete, saying that they do not model runs that can actually occur, and/or declare some finite paths that do not end in deadlock complete. A complete path π ending in a state models a run of the represented system that follows the path until its last state, and then stays in that state forever, without taking any of its outgoing transitions. A complete path that ends in a transition models a run in which the action represented by this last transition starts occurring but never finishes. It is often assumed that transitions are instantaneous, or at least of finite duration. This assumption is formalised through the adoption of a completeness criterion that holds all paths ending in a transition to be incomplete.

The most prominent translation from LTSs to Kripke structures is from De Nicola & Vaandrager [1]. Its purpose is merely to efficiently lift the validity relation \models from Kripke structures to LTSs. It simply creates a new state halfway along any transition labelled by a visible action, and moves the transition label to that state.

Definition 4 Let $(\mathbb{P}, Tr, source, target, \ell)$ be an LTS over $Act = A \cup \{\tau\}$. The associated Kripke structure $(S, \rightarrow, \models)$ over A is given by

- $S := \mathbb{P} \cup \{t \in Tr \mid \ell(t) \neq \tau\}$,
- $\rightarrow := \{(source(t), t), (t, target(t)) \mid t \in Tr \wedge \ell(t) \neq \tau\} \cup \{(source(t), target(t)) \mid t \in Tr \wedge \ell(t) = \tau\}$
- and $\models := \{(t, \ell(t)) \mid t \in Tr \wedge \ell(t) \neq \tau\}$.

Ignoring paths ending within a τ -transition, which are never deemed complete anyway, this translation yields a bijective correspondence between the paths in an LTS and the path in its associated Kripke structure. Consequently, any completeness criterion on the LTS induces a completeness criterion on the Kripke structure. Hence it is now well-defined when $s \models^{CC} \varphi$, with s a state in an LTS, CC a completeness criterion on this LTS and φ an $LTL_{\neg X}$ formula.

3.2 Petri Nets

Definition 5 A (labelled) Petri net over Act is a tuple $N = (S, T, F, M_0, \ell)$ where

- S and T are disjoint sets (of *places* and *transitions*),
- $F : (S \times T \cup T \times S) \rightarrow \mathbb{N}$ (the *flow relation* including *arc weights*) such that $\forall t \in T \exists s \in S. F(s, t) > 0$,
- $M_0 : S \rightarrow \mathbb{N}$ (the *initial marking*), and
- $\ell : T \rightarrow Act$ (the *labelling function*).

Petri nets are depicted by drawing the places as circles and the transitions as boxes, containing their label. For $x, y \in S \cup T$ there are $F(x, y)$ arrows (*arcs*) from x to y . When a Petri net represents a distributed system, a global state of this system is given as a *marking*, a multiset of places, depicted by placing $M(s)$ dots (*tokens*) in each place s . The initial state is M_0 . The behaviour of a Petri net is defined by the possible moves between markings M and M' , which take place when a finite multiset G of transitions *fires*. In that case, each occurrence of a transition t in G consumes $F(s, t)$ tokens from each place s . Naturally, this can happen only if M makes all these tokens available in the first place. Next, each t produces $F(t, s)$ tokens in each s . Definition 7 formalises this notion of behaviour.

A *multiset* over a set X is a function $A : X \rightarrow \mathbb{N}$, i.e. $A \in \mathbb{N}^X$. Object $x \in X$ is an *element of A* iff $A(x) > 0$. A multiset is *empty* iff it has no elements, and *finite* iff the set of its elements is finite. For multisets A and B over X I write $A \leq B$ iff $A(x) \leq B(x)$ for all $x \in X$; $A + B$ denotes the multiset over X with $(A + B)(x) := A(x) + B(x)$, $A - B$ is given by $(A - B)(x) := A(x) \dot{-} B(x) = \max(A(x) - B(x), 0)$, and for $k \in \mathbb{N}$ the multiset $k \cdot A$ is given by $(k \cdot A)(x) := k \cdot A(x)$. With $\{x, x, y\}$ I denote a multiset A with $A(x) = 2$ and $A(y) = 1$, rather than the set $\{x, y\}$ itself.

Definition 6 Let $N = (S, T, F, M_0, \ell)$ be a Petri net and $t \in T$. The multisets $\bullet t, t^\bullet : S \rightarrow \mathbb{N}$ are given by $\bullet t(s) = F(s, t)$ and $t^\bullet(s) = F(t, s)$ for all $s \in S$. The elements of $\bullet t$ and t^\bullet are called *pre-* and *postplaces* of t , respectively. These functions extend to finite multisets $G : T \rightarrow \mathbb{N}$ as usual, by $\bullet G := \sum_{t \in T} G(t) \cdot \bullet t$ and $G^\bullet := \sum_{t \in T} G(t) \cdot t^\bullet$.

Definition 7 Let $N = (S, T, F, M_0, \ell)$ be a Petri net, $G \in \mathbb{N}^T$, G non-empty and finite, and $M, M' \in \mathbb{N}^S$. G is a *step* from M to M' , written $M \xrightarrow{G}_N M'$, iff $\bullet G \leq M$ (G is *enabled*) and $M' = (M - \bullet G) + G^\bullet$.

Write $M_0 \twoheadrightarrow_N M$ iff there are transitions $t_i \in T$ and markings $M_i \in \mathbb{N}^S$ for $i = 1, \dots, k$, such that $M_k = M$ and $M_{i-1} \xrightarrow{\{t_i\}}_N M_i$ for $i = 1, \dots, k$. Moreover, $M_0 \twoheadrightarrow \xrightarrow{G}$ means that $M_0 \twoheadrightarrow_N M \xrightarrow{G}_N M'$ for some M and M' .

Definition 8 ([10]) $N = (S, T, F, M_0, \ell)$ is a *structural conflict net* iff $\forall t, u. (M_0 \twoheadrightarrow \xrightarrow{\{t, u\}}) \Rightarrow \bullet t \cap \bullet u = \emptyset$.

Here I restrict myself to structural conflict nets, henceforth simply called *nets*, a class of Petri nets containing the *safe* Petri nets that are normally used to give semantics to process algebras.

Given a net $N = (S, T, F, M_0, \ell)$, its associated LTS $(\mathbb{P}, Tr, source, target, \ell)$ is given by $\mathbb{P} := \mathbb{N}^S$, $Tr := \{(M, t) \in \mathbb{N}^S \times Tr \mid \bullet t \leq M\}$, $source(M, t) := M$, $target(M, t) := (M - \bullet t) + t^\bullet$ and $\ell(M, t) := \ell(t)$. The net N maps to the state M_0 in this LTS. A *completeness criterion* on a net is a completeness criterion on its associated LTS. Now $N \models^{CC} \varphi$ is defined to hold iff $M_0 \models^{CC} \varphi$ in the associated LTS.

3.3 CCS

CCS [19] is parametrised with sets \mathcal{K} of *agent identifiers* and \mathcal{A} of *names*; each $X \in \mathcal{K}$ comes with a defining equation $X \stackrel{def}{=} P$ with P being a CCS expression as defined below. $Act := \mathcal{A} \dot{\cup} \bar{\mathcal{A}} \dot{\cup} \{\tau\}$ is the set of *actions*, where τ is a special *internal action* and $\bar{\mathcal{A}} := \{\bar{a} \mid a \in \mathcal{A}\}$ is the set of *co-names*. Complementation is extended to $\bar{\mathcal{A}}$ by setting $\bar{\bar{a}} = a$. Below, a ranges over $\mathcal{A} \cup \bar{\mathcal{A}}$, α over Act , and X, Y over \mathcal{K} . A *relabelling* is a function $f : \mathcal{A} \rightarrow \mathcal{A}$; it extends to Act by $f(\bar{a}) = \overline{f(a)}$ and $f(\tau) := \tau$. The set T_{CCS} of CCS expressions or *processes* is the smallest set including:

$\sum_{i \in I} \alpha_i.P_i$	for I an index set, $\alpha_i \in Act$ and $P_i \in T_{CCS}$	<i>guarded choice</i>
$P Q$	for $P, Q \in T_{CCS}$	<i>parallel composition</i>
$P \setminus L$	for $L \subseteq \mathcal{A}$ and $P \in T_{CCS}$	<i>restriction</i>
$P[f]$	for f a relabelling and $P \in T_{CCS}$	<i>relabelling</i>
X	for $X \in \mathcal{K}$	<i>agent identifier</i>

The process $\sum_{i \in \{1,2\}} \alpha_i.P_i$ is often written as $\alpha_1.P_1 + \alpha_2.P_2$, and $\sum_{i \in \emptyset} \alpha_i.P_i$ as $\mathbf{0}$. The semantics of CCS is given by the transition relation $\rightarrow \subseteq T_{CCS} \times Act \times \mathcal{P}(\mathcal{C}) \times T_{CCS}$, where transitions $P \xrightarrow{\alpha, \mathcal{C}} Q$ are derived from the rules of Table 1. Ignoring the labels $\mathcal{C} \in \mathcal{P}(\mathcal{C})$ for now, such a transition indicates that process P can perform the action $\alpha \in Act$ and transform into process Q . The process $\sum_{i \in I} \alpha_i.P_i$ performs one of the actions α_j for $j \in I$ and subsequently acts as P_j . The parallel composition $P|Q$ executes an action from P , an action from Q , or a synchronisation between complementary actions c and \bar{c} performed by P and Q , resulting in an internal action τ . The restriction operator $P \setminus L$ inhibits execution of the actions from L and their complements. The relabelling $P[f]$ acts like process P with all labels α replaced by $f(\alpha)$. Finally, the rule for agent identifiers says that an agent X has the same transitions as the body P of its defining equation. The standard version of CCS [19] features a *choice* operator $\sum_{i \in I} P_i$; here I use the fragment of CCS that merely features guarded choice.

The second label of a transition indicates the set of (parallel) *components* involved in executing this transition. The set \mathcal{C} of components is defined as $\{L, R\}^*$, that is, the set of strings over the indicators Left and Right, with $\varepsilon \in \mathcal{C}$ denoting the empty string and $D \cdot C := \{D\sigma \mid \sigma \in C\}$ for $D \in \{L, R\}$ and $C \subseteq \mathcal{C}$.

Table 1: Structural operational semantics of CCS

$\sum_{i \in I} \alpha_i.P_i \xrightarrow{\alpha_i, \{\varepsilon\}} P_j \quad (j \in I)$		
$\frac{P \xrightarrow{\alpha, C} P'}{P Q \xrightarrow{\alpha, L.C} P' Q}$	$\frac{P \xrightarrow{a, C} P', Q \xrightarrow{\bar{a}, D} Q'}{P Q \xrightarrow{\tau, L.C \cup R.D} P' Q'}$	$\frac{Q \xrightarrow{\alpha, D} Q'}{P Q \xrightarrow{\alpha, R.D} P Q'}$
$\frac{P \xrightarrow{\alpha, C} P'}{P \setminus L \xrightarrow{\alpha, C} P' \setminus L} \quad (\alpha, \bar{\alpha} \notin L)$	$\frac{P \xrightarrow{\alpha, C} P'}{P[f] \xrightarrow{f(\alpha), C} P'[f]}$	$\frac{P \xrightarrow{\alpha, C} P'}{X \xrightarrow{\alpha, C} P'} \quad (X \stackrel{def}{=} P)$

Example 4 The CCS process $P := (X|\bar{a}.0)|\bar{a}.b.0$ with $X \stackrel{def}{=} a.X$ has as outgoing transitions $P \xrightarrow{a, \{LL\}} P$, $P \xrightarrow{\tau, \{LL, LR\}} (X|0)|\bar{a}.b.0$, $P \xrightarrow{\bar{a}, \{LR\}} (X|0)|\bar{a}.b.0$, $P \xrightarrow{\tau, \{LL, R\}} (X|\bar{a}.0)|b.0$ and $P \xrightarrow{\bar{a}, \{R\}} (X|\bar{a}.0)|b.0$. These components stem from Victor Dyseryn [personal communication] and were introduced in [8]. They were not part of the standard semantics of CCS [19], which can be retrieved by ignoring them.

The LTS of CCS is $(\mathbb{T}, Tr, source, target, \ell)$, with $Tr = \{(P, \alpha, C, Q) \mid P \xrightarrow{\alpha, C} Q\}$, $\ell(P, \alpha, C, Q) = \alpha$, $source(P, \alpha, C, Q) = P$ and $target(P, \alpha, C, Q) = Q$. Employing this interpretation of CCS, one can pronounce judgements $P \models^{CC} \varphi$ for CCS processes P .

3.4 Labelled Transition Systems with Concurrency

Definition 9 A *labelled transition system with concurrency* (LTSC) is a tuple $(\mathbb{P}, Tr, source, target, \ell, \smile)$ consisting of a LTS $(\mathbb{P}, Tr, source, target, \ell)$ and a *concurrency relation* $\smile \subseteq Tr \times Tr$, such that:

$$t \not\smile t \text{ for all } t \in Tr, \quad (1)$$

$$\text{if } t \in Tr \text{ and } \pi \text{ is a path from } source(t) \text{ to } s \in \mathbb{P} \text{ such that } t \smile v \text{ for all transitions } v \text{ occurring in } \pi, \text{ then there is a } u \in Tr \text{ such that } source(u) = s, \ell(u) = \ell(t) \text{ and } t \not\smile u. \quad (2)$$

Informally, $t \smile v$ means that the transition v does not interfere with t , in the sense that it does not affect any resources that are needed by t , so that in a state where t and v are both possible, after doing v one can still do a future variant u of t .

LTSCs were introduced in [9], although there the model is more general on various counts. I do not need this generality in the present paper. In particular, I only need symmetric concurrency relations \smile ; in [9] \smile is not always symmetric, and denoted \smile .

The LTS associated with CCS can be turned into an LTSC by defining $(P, \alpha, C, P') \smile (Q, \beta, D, Q')$ iff $C \cap D = \emptyset$, that is, two transitions are concurrent iff they stem from disjoint sets of components [13, 8].

Example 5 Let the 5 transitions from Example 4 be t, u, v, w and x , respectively. Then $t \not\smile w$ because these transitions share the component LL. Yet $v \smile w$.

The LTS associated with a Petri net can be turned into an LTSC by defining $(M, t) \smile (M', u)$ iff $\bullet t \cap \bullet u = \emptyset$, i.e., the two LTS-transitions stem from net-transitions that have no preplaces in common.

Naturally, an LTSC can be turned into a LTS, and further into a Kripke structure, by forgetting \smile .

4 Progress, Justness and Fairness

In this section I define completeness criteria $CC \in \{SF(\mathcal{T}), WF(\mathcal{T}), J, Pr, \top \mid \mathcal{T} \in \mathcal{P}(\mathcal{P}(Tr))\}$ on LTSs $(\mathbb{P}, Tr, source, target, \ell)$, to be used in judgements $P \models^{CC} \varphi$, for $P \in \mathbb{P}$ and φ an $LTL_{\mathbf{X}}$ formula. These

criteria are called *strong fairness* (SF), *weak fairness* (SF), both parametrised with a set $\mathcal{T} \subseteq \mathcal{P}(Tr)$ of tasks, *justness* (J), *progress* (Pr) and the *trivial* completeness criterion (\top). Justness is merely defined on LTSCs. I confine myself to criteria that hold finite paths ending within a transition to be incomplete.

Reading Example 1, one could find it unfair that Bart might never get a beer. Strong and weak *fairness* are completeness criteria that postulate that Bart will get a beer, namely by ruling out as incomplete the infinite paths in which he does not. They can be formalised by introducing a set \mathcal{T} of tasks, each being a set of transitions (in an LTS or Kripke structure).

Definition 10 ([13]) A task $T \in \mathcal{T}$ is *enabled* in a state s iff s has an outgoing transition from T . It is *perpetually enabled* on a path π iff it is enabled in every state of π . It is *relentlessly enabled* on π , if each suffix of π contains a state in which it is enabled.² It *occurs* in π if π contains a transition $t \in T$.

A path π is *weakly fair* if, for every suffix π' of π , each task that is perpetually enabled on π' , occurs in π' . It is *strongly fair* if, for every suffix π' of π , each task that is relentlessly enabled on π' , occurs in π' .

As completeness criteria, these notions take only the fair paths to be complete. In Example 1 it suffices to have a task “Bart gets a beer”, consisting of the three transitions leading to the B state. Now in any path in which Bart never gets a beer this task is perpetually enabled, yet never taken. Hence weak fairness suffices to rule out such paths. We have $\mathcal{D} \models^{WF(\mathcal{T})} \mathbf{FB}$.

Local fairness [13] allows the tasks \mathcal{T} to be declared on an ad hoc basis for the application at hand. On this basis one can call it unfair if Bart doesn’t get a beer, without requiring that Cameron should get a beer as well. *Global fairness*, on the other hand, distils the tasks of an LTS in a systematic way out of the structure of a formalism, such as pseudocode, process algebra or Petri nets, that gave rise to the LTS. A classification of many ways to do this, and thus of many notions of strong and weak fairness, appears in [13]. In *fairness of directions* [5], for instance, each transition in an LTS is assumed to stem from a particular *direction*, or *instruction*, in the pseudocode that generated the LTS; now each direction represents a task, consisting of all transitions derived from that direction.

In [13] the assumption that a system will never stop when there are transitions to proceed is called *progress*. In Example 2 it takes a progress assumption to conclude that Bart will get his beer. Progress fits the default completeness criterion introduced before, i.e., \models^{Pr} is the same as \models . Not (even) assuming progress can be formalised by the trivial completeness criterion \top that declares all paths to be complete. Naturally, $\mathcal{E} \not\models^{\top} \mathbf{FB}$.

Completeness criterion D is called *stronger* than criterion C if it rules out more paths as incomplete. So \top is the weakest of all criteria, and, for any given collection \mathcal{T} , strong fairness is stronger than weak fairness. When assuming that each transition occurs in at least one task—which can be ensured by incorporating a default task consisting of all transitions—progress is weaker than weak fairness.

Justness [13] is a strong form of progress, defined on LTSCs.

Definition 11 A path π is *just* if for each transition t with its source state $s := source(t)$ occurring on π , the suffix of π starting at s contains a transition u with $t \not\prec u$.

Example 6 The infinite path π that only ever takes transition t in Example 4/5 is unjust. Namely with transition v in the rôle of the t from Definition 11, π contains no transition y with $v \not\prec y$.

Informally, the only reason for an enabled transition not to occur, is that one of its resources is eventually used for some other transition. In Example 3 for instance, the orders of Alice and Cameron are clearly concurrent with the one of Bart, in the sense that they do not compete for shared resources. Taking t to be the transition in which Bart gets his beer, any path in which t does not occur is unjust. Thus $\mathcal{F} \models^J \mathbf{FB}$.

²This is the case if the task is enabled in infinitely many states of π , in a state that occurs infinitely often in π , or in the last state of a finite π .

For most choices of \mathcal{T} found in the literature, weak fairness is a strictly stronger completeness criterion than justness. In Example 1, for instance, the path in which Bart does not get a beer is just. Namely, any transition u giving Alice or Cameron a beer competes for the same resource as the transition t giving Bart a beer, namely the attention of the barman. Thus $t \not\prec u$, and consequently $\mathcal{D} \not\models^J \mathbf{FB}$.

5 Reactive Temporal Logic

Standard treatments of temporal logic [23, 14] are adequate for *closed systems*, having no run-time interactions with their environment. However, they fall short for *reactive systems*, interacting with their environments through synchronisation of actions.

Example 7 Consider a vending machine that accepts a coin c and produces a pretzel p . We assume that accepting the coin requires cooperation from the user/environment, but producing the pretzel does not. A CCS specification is

$$VM = c.p.VM.$$

In standard $LTL_{\mathbf{X}}$ (assuming progress) we have $VM \models \mathbf{G}(c \Rightarrow \mathbf{F}p)$. This formula says that whenever a coin is inserted, eventually a pretzel is produced. This formula is intuitively true indeed. But we also have $VM \models \mathbf{G}(p \Rightarrow \mathbf{F}c)$. This formula says that whenever a pretzel is produced, eventually a new coin will be inserted. This formula is intuitively false. This example shows that standard $LTL_{\mathbf{X}}$ is not suitable to correctly describe the behaviour of this vending machine.

For this reason I here introduce *reactive* $LTL_{\mathbf{X}}$. The syntax and semantics are unchanged, except that I use a validity relation \models_B that is parametrised with a set $B \subseteq A$ of *blockable* actions. Here A is the set of all observable actions of the LTS on which $LTL_{\mathbf{X}}$ is interpreted. The intuition is that actions $b \in B$ may be blocked by the environment, but actions $a \in A \setminus B$ may not. The relation \models_B can be used to formalise the assumption that the actions in $A \setminus B$ are not under the control of the user of the modelled system, or that there is an agreement with the user not to block them. Either way, it is a disclaimer on the wrapping of our temporal judgement, that it is valid only when applying the involved distributed system in an environment that may block actions from B only. The hidden action τ may never be blocked.

The subscript B modifies the default completeness criterion, to call a path complete iff it is either infinite or ends in a state of which all outgoing transitions have a label from B . Note that the standard $LTL_{\mathbf{X}}$ interpretation \models is simply \models_{\emptyset} , obtained by taking the empty set of blocking actions.

In Example 7 one takes $B = \{c\}$. This choice of B says that the environment may block the action c , namely by not inserting a coin; however, the environment may not block p . As intuitively expected, we have $VM \models_B \mathbf{G}(c \Rightarrow \mathbf{F}p)$ but $VM \not\models_B \mathbf{G}(p \Rightarrow \mathbf{F}c)$.

Naturally, reactive $LTL_{\mathbf{X}}$ can also be combined with a non-default completeness criterion, as discussed in Sections 2–4. When writing $P \models_B^{CC} \varphi$ the modifier B adapts the default completeness criterion by declaring certain finite paths complete, and the modifier $CC \neq \top$ adapts it by declaring some infinite paths incomplete. In the presence of the modifier B , Definition 11 and the first sentence of Definition 10 are adapted as follows:

Definition 12 A path π is *just* (or *B-just*) if for each transition $t \in Tr$ with $\ell(t) \notin B$ and its source state $s := source(t)$ occurring on π , the suffix of π starting at s contains a transition u with $t \not\prec u$.

Note that it doesn't matter whether $\ell(u) \in B$ or not.

Definition 13 A task $T \in \mathcal{T}$ is *enabled* in a state s iff s has an outgoing transition $t \in T$ with $\ell(t) \notin B$.

The above completes the formal definition of the validity of temporal judgements $P \models_B^{CC} \varphi$ with φ an LTL_x formula, $B \subseteq A$, and either

- $CC = Pr$ and P a state in an LTS, a CCS expression or a Petri net,
- $CC = J$ and P a state in an LTSC, a CCS expression or a Petri net,
- $CC = WF(\mathcal{T})$ or $SF(\mathcal{T})$ and P a state in an LTS $(\mathbb{P}, Tr, source, target, \ell)$ with $\mathcal{T} \in \mathcal{P}(\mathcal{P}(Tr))$, or P a CCS expression or Petri net with associated LTS $(\mathbb{P}, Tr, source, target, \ell)$ and $\mathcal{T} \in \mathcal{P}(\mathcal{P}(Tr))$.

Namely, in case P is a state in an LTS, it is also a state in the associated Kripke structure K . Moreover, B and CC combine into a single completeness criterion BC on that LTS, which translates as a completeness criterion BC on K . Now Definition 2 tells whether $P \models^{BC} \varphi$ holds.

In case $CC = J$ and P a state in an LTSC, B and J combine into a single completeness criterion BJ on that LTSC, which is also a completeness criterion on the associated LTS; now proceed as above.

In case P is a Petri net or CCS expression, first translate it into a state in an LTS or LTSC, using the translations at the end of Sections 3.2 or 3.3, respectively, and proceed as above.

Temporal judgements $P \models_B^{CC} \varphi$, as introduced above, are not limited to the case that φ is an LTL formula. In Section 10 I will show that allowing φ to be a CTL formula instead poses no additional complications, and I expect the same to hold for other temporal logics.

Judgements $P \models_B^{CC} \varphi$ get stronger (= less likely true) when the completeness criterion CC is weaker, and the set B of blockable actions larger.

Most concepts of reactive temporal logic introduced above stem from [12]. The main novelty contributed here is the annotated satisfaction relation \models_B^{CC} . In [12] we simply wrote \models , expecting CC and B to be determined once and for all in a given paper or application. Requirement specifications in which different values for B are combined, such as FS 1–2 in Section 8, were not foreseen there.

6 The Mutual Exclusion Problem and its History

The mutual exclusion problem was presented by Dijkstra in [3] and formulated as follows:

“To begin, consider N computers, each engaged in a process which, for our aims, can be regarded as cyclic. In each of the cycles a so-called “critical section” occurs and the computers have to be programmed in such a way that at any moment only one of these N cyclic processes is in its critical section. In order to effectuate this mutual exclusion of critical-section execution the computers can communicate with each other via a common store. Writing a word into or nondestructively reading a word from this store are undividable operations; i.e., when two or more computers try to communicate (either for reading or for writing) simultaneously with the same common location, these communications will take place one after the other, but in an unknown order.”

Dijkstra proceeds to formulate a number of requirements that a solution to this problem must satisfy, and then presents a solution that satisfies those requirements. The most central of these are:

- (*Safety*) “no two computers can be in their critical section simultaneously”, and
- (*Dijkstra’s Liveness*) If at least one computer intends to enter its critical section, then at least one “will be allowed to enter its critical section in due time”.

Two other important requirements formulated by Dijkstra are

- (*Speed independence*) “(b) Nothing may be assumed about the relative speeds of the N computers”,
- and (*Optionality*) “(c) If any of the computers is stopped well outside its critical section, this is not allowed to lead to potential blocking of the others.”

A crucial assumption is that each computer, in each cycle, spends only a finite amount of time in its critical section. This is necessary for the correctness of any mutual exclusion protocol.

For the purpose of the last requirement one can partition each cycle into a *critical section*, a *non-critical section* (in which the process starts), an *entry protocol* between the noncritical and the critical section, during which a process prepares for entry in negotiation with the competing processes, and an *exit protocol*, that comes right after the critical section and before return to the noncritical section. Now “well outside its critical section” means in the noncritical section. Requirement (c) can equivalently be stated as admitting the possibility that a process chooses to remain forever in its noncritical section, without applying for entry in the critical section ever again.

Knuth [16] proposes a strengthening of Dijkstra’s liveness requirement, namely

- (*Liveness*) If a computer intends to enter its critical section, then it will be allowed to enter in due time.

He also presents a solution that is shown to satisfy this requirement, as well as Dijkstra’s requirements.³ Henceforth I define a correct solution of the mutual exclusion problem as one that satisfies both safety and liveness, as formulated above, as well as optionality. I sometimes speak of “speed independent mutual exclusion” when also insisting on requirement (b) above.

The special case of the mutual exclusion problem for two processes ($N = 2$) was presented by Dijkstra in [2], three years prior to [3]. There Dijkstra presented a solution found by T.J. Dekker in 1959, and shows that it satisfies all requirements of [3]. Although not explicitly stated in [2], the arguments given therein imply straightforwardly that Dekker’s solution also satisfy the liveness requirement above.

Peterson [22] presented a considerable simplification of Dekker’s algorithm that satisfies the same correctness requirements. Many other mutual exclusion protocols appear in the literature, the most prominent being Lamport’s bakery algorithm [17] and Szymański’s mutual exclusion algorithm [24]. These guarantee some additional correctness criteria besides the ones discussed above.

7 Fair Schedulers

In [11] a *fair scheduler* is defined as

FS1 “a reactive system with two input channels: one on which it can receive requests r_1 from its environment and one on which it can receive requests r_2 . We allow the scheduler to be too busy shortly after receiving a request r_i to accept another request r_i on the same channel. However, the system will always return to a state where it remains ready to accept the next request r_i until r_i arrives. In case no request arrives it remains ready forever. The environment is under no obligation to issue requests, or to ever stop issuing requests. Hence for any numbers n_1 and $n_2 \in \mathbb{N} \cup \{\infty\}$ there is at least one run of the system in which exactly that many requests of type r_1 and r_2 are received.

FS2 Every request r_i asks for a task t_i to be executed. The crucial property of the fair scheduler is that it will eventually grant any such request. Thus, we require that in any run of the system each occurrence of r_i will be followed by an occurrence of t_i .”

FS3 “We require that in any partial run of the scheduler there may not be more occurrences of t_i than of r_i , for $i = 1, 2$.”

³It can be argued, however, that Knuth’s mutual exclusion protocol is correct only when making certain assumptions on the hardware on which it will be running [7]; the same applies to all other mutual exclusion protocols mentioned in this section. This matter is not addressed in the present paper. However, the material presented in Section 9 paves the way for discussing it.

FS4 The last requirement is that between each two occurrences of t_i and t_j for $i, j \in \{1, 2\}$ an intermittent activity e is scheduled.”

This fair scheduler serves two clients, but the concept generalises smoothly to N clients.

The intended applications of fair schedulers are for instance in operating systems, where multiple application processes compete for processing on a single core, or radio broadcasting stations, where the station manager needs to schedule multiple parties competing for airtime. In such cases each applicant must get a turn eventually. The event e signals the end of the time slot allocated to an application process on the single core, or to a broadcast on the radio station.

Fair schedulers occur (in suitable variations) in many distributed systems. Examples are *First in First out*⁴, *Round Robin*, and *Fair Queueing* scheduling algorithms⁵ as used in network routers [20, 21] and operating systems [15], or the *Completely Fair Scheduler*,⁶ which is the default scheduler of the Linux kernel since version 2.6.23.

Each action r_i , t_i and e can be seen as a communication between the fair scheduler and one of its clients. In a reactive system such communications will take place only if both the fair scheduler and its client are ready for it. Requirement FS1 of a fair scheduler quoted above effectively shifts the responsibility for executing r_i to the client. The actions t_i and e , on the other hand, are seen as the responsibility of the fair scheduler. We do not consider the possibility that the fair scheduler fails to execute t_i merely because the client does not collaborate. Hence [11] assumes that the client cannot prevent the actions t_i and e from occurring. It is furthermore assumed that executing the actions r_i , t_i and e takes a finite amount of time only.

A fair scheduler closely resembles a mutual exclusion protocol. However, its goal is not to achieve mutual exclusion. In most applications, mutual exclusion can be taken for granted, as it is physically impossible to allocate the single core to multiple applications at the same time, or the (single frequency) radio sender to multiple simultaneous broadcasts. Instead, its goal is to ensure that no applicant is passed over forever.

It is not hard to obtain a fair scheduler from a mutual exclusion protocol. For suppose we have a mutual exclusion protocol M , serving two processes P_i ($i = 1, 2$). I instantiate the non-critical section of process P_i as patiently awaiting the request r_i . As soon as this request arrives, P_i leaves the noncritical section and starts the entry protocol to get access to the critical section. The liveness property for mutual exclusion guarantees that P_i will reach its critical section. Now the critical section consists of scheduling task t_i , followed by the intermittent activity e . Trivially, the composition of the two process P_i , in combination with protocol M , constitutes a fair scheduler, in that it meets the above four requirements.

One can not quite construct a mutual exclusion protocol from a fair scheduler, due to fact that in a mutual exclusion protocol leaving the critical section is controlled by the client process. For this purpose one would need to adapt the assumption that the client of a fair scheduler cannot block the intermittent activity e into the assumption that the client can postpone this action, but for a finite amount of time only. In this setting one can build a mutual exclusion protocol, serving two processes P_i ($i = 1, 2$), from a fair scheduler F . Process i simply issues request r_i at F as soon as it has left the non-critical section, and when F communicates the action t_i , Process i enters its critical section. Upon leaving its critical section, which is assumed to happen after a finite amount of time, it participates in the synchronisation e with F . Trivially, this yields a correct mutual exclusion protocol.

⁴Also known as First Come First Served (FCFS)

⁵[http://en.wikipedia.org/wiki/Scheduling_\(computing\)](http://en.wikipedia.org/wiki/Scheduling_(computing))

⁶http://en.wikipedia.org/wiki/Completely_Fair_Scheduler

8 Formalising the Requirements for Fair Schedulers in Reactive $LTL_{\mathbf{X}}$

The main reason fair schedulers were defined in [11] was to serve as an example of a realistic class of systems of which no representative can be correctly specified in CCS, or similar process algebras, or in Petri nets. Proving this impossibility result necessitated a precise formalisation of the four requirements quoted in Section 7. Through the provided translations of CCS and Petri nets into LTSs, a fair scheduler rendered in CCS or Petri nets can be seen as a state F in an LTS over the set $\{r_i, t_i, e \mid i = 1, 2\}$ of visible actions; all other actions can be considered internal and renamed into τ .

Let a *partial trace* of a state s in an LTS be the sequence of visible actions encountered on a path starting in s [6]. Now the last two requirements (FS3) and (FS4) of a fair scheduler are simple properties that should be satisfied by all partial traces σ of state F :

(FS3) σ contains no more occurrences of t_i than of r_i , for $i = 1, 2$,

(FS4) σ contains an occurrence of e between each two occurrences of t_i and t_j for $i, j \in \{1, 2\}$.

FS4 can be conveniently rendered in $LTL_{\mathbf{X}}$:

(FS4) $F \models \mathbf{G}(t_i \Rightarrow (t_i \mathbf{U}((\neg t_1 \wedge \neg t_2) \mathbf{W} e)))$

for $i \in \{1, 2\}$. Here the *weak until* modality $\psi \mathbf{W} \varphi$ is syntactic sugar for $\mathbf{G}\psi \vee (\psi \mathbf{U} \varphi)$. If I hadn't lost the \mathbf{X} modality, I could write \mathbf{X} for $t_i \mathbf{U}$ in the above formula; on Kripke structures distilled from LTSs the meaning is the same. The formula in FS4 is of a kind where the meaning of \models_B^{CC} is independent of B and CC . This follows from the fact that FS4 merely formulates a property that should hold for all partial runs. Hence one need not worry about which B and CC to employ here.

Unfortunately, FS3 cannot be formulated in $LTL_{\mathbf{X}}$, due to the need to keep count of the difference in the number of r_i and t_i actions encountered on a path. However, one could strengthen FS3 into

(FS3') σ contains an occurrence of r_i between each two occurrences of t_i , and prior to the first occurrence of t_i , for $i \in \{1, 2\}$.

This would restrict the class of acceptable fair schedulers, but keep the most interesting examples. Consequently, the impossibility result from [11] applies to this modified class as well. FS3 can be rendered in $LTL_{\mathbf{X}}$ in the same style as FS4:

(FS3') $F \models ((\neg t_i) \mathbf{W} r_i) \wedge \mathbf{G}(t_i \Rightarrow (t_i \mathbf{U}((\neg t_i) \mathbf{W} r_i)))$

for $i \in \{1, 2\}$.

Requirement FS2 involves a quantification over all complete runs of the system, and thus depends on the completeness criterion CC employed. It can be formalised as

(FS2) $F \models_B^{CC} \mathbf{G}(r_i \Rightarrow \mathbf{F}t_i)$

for $i \in \{1, 2\}$, where $B = \{r_1, r_2\}$. The set B should contain r_1 and r_2 , as these actions are supposed to be under the control of the users of a fair scheduler. However, actions t_1 , t_2 and e should not be in B , as they are under the control of the scheduler itself. In [11], the completeness criterion employed is justness, so the above formula with $CC := J$ captures the requirement on the fair schedulers that are shown in [11] not to exist in CCS or Petri nets. However, keeping CC a variable allows one to pose to the question under which completeness criterion a fair scheduler *can* be rendered in CCS. Naturally, it needs to be a stronger criterion than justness. In [11] it is shown that weak fairness suffices.

FS2 is a good example of a requirement that can *not* be rendered correctly in standard LTL. Writing $F \models^{CC} \mathbf{G}(r_i \Rightarrow \mathbf{F}t_i)$ would rule out the complete runs of F that end because the user of F never supplies the input $r_j \in B$. The CCS process

$$F \stackrel{def}{=} r_1.r_2.t_1.e.t_2.e.F$$

for instance satisfies this formula, as well as FS3 and 4; yet it does not satisfy requirement FS2. Namely, the path consisting of the r_1 -transition only is complete, since it ends in a state of which the only outgoing transition has the label $r_2 \in B$. Yet on this path r_1 is not followed by t_1 .

Requirement FS1 is by far the hardest to formalise. In [11] two formalisations are shown to be equivalent: one involving a coinductive definition of B -just paths that exploits the syntax of CCS, and the other requiring that requirements FS2–4 are preserved under putting an input interface around process F . The latter demands that also $\widehat{F} := (I_1 | F[f] | I_2) \setminus \{c_1, c_2\}$ should satisfy FS2–4; here f is a relabelling with $f(r_i) = c_i$, $f(t_i) = t_i$ and $f(e) = e$ for $i = 1, 2$, and $I_i \stackrel{\text{def}}{=} r_i.\bar{c}_i.I_i$ for $i \in \{1, 2\}$.

A formalisation of FS1 on Petri nets also appears in [11]: each complete path π with only finitely many occurrences of r_i should contain a state (= marking) M , such that there is a transition v with $\ell(v) = r_i$ and $\bullet v \leq M$, and for each transition u that occurs in π past M one has $\bullet v \cap \bullet u = \emptyset$.

When discussing proposals for fair schedulers by others, FS1 is the requirement that is most often violated, and explaining why is not always easy.

In reactive LTL_X, this requirement is formalised as

$$(FS1) \quad F \models_{B \setminus \{r_i\}}^J \mathbf{GF}r_i$$

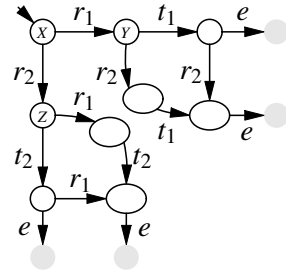
for $i \in \{1, 2\}$, or $F \models_{B \setminus \{r_i\}}^{CC} \mathbf{GF}r_i$ if one wants to discuss the completeness criterion CC as a parameter. The surprising element in this temporal judgement is the subscript $B \setminus \{r_i\} = \{r_{3-i}\}$, which contrasts with the assumption that requests are under the control of the environment. FS1 says that, although we know that there is no guarantee that user i of F will ever issue request r_i , under the assumption that the user *does* want to make such a request, making the request should certainly succeed. This means that the protocol itself does not sit in the way of making this request.

The combination of requirements FS1 and 2, which use different sets of blockable actions as a parameter, is enabled by reactive LTL_X as presented here.

The following examples, taken from [11], show that all the above requirements are necessary for the result from [11] that fair schedulers cannot be rendered in CCS.

- The CCS process $F_1 | F_2$ with $F_i \stackrel{\text{def}}{=} r_i.t_i.e.F_i$ satisfies FS1, FS2 and FS3'. In FS1 and 2 one needs to take $CC := J$, as progress is not a strong enough assumption here.
- The process $E_1 | G | E_2$ with $E_i \stackrel{\text{def}}{=} r_i.E_i$ and $G \stackrel{\text{def}}{=} t_1.e.t_2.e.G$ satisfies FS1, 2 and 4, again with $CC := J$.
- The process $E_1 | E_2$ satisfies FS1, 3' and 4, again with $CC := J$ in FS1.
- The process F_0 with $F_0 \stackrel{\text{def}}{=} r_1.t_1.e.F_0 + r_2.t_2.e.F_0$ satisfies FS2–4. Here FS2 merely needs $CC := Pr$, that is, the assumption of progress. Furthermore, it satisfies FS1 with $CC := SF(\mathcal{T})$, as long as $R_1, R_2 \in \mathcal{T}$. Here R_i is the set of transitions with label r_i .

The process X given by $X \stackrel{\text{def}}{=} r_1.Y + r_2.Z$, $Y \stackrel{\text{def}}{=} r_2.t_1.e.Z + t_1.(r_2.e.Z + e.X)$ and $Z \stackrel{\text{def}}{=} r_1.t_2.e.Y + t_2.(r_1.e.Y + e.X)$, the *gatekeeper*, is depicted on the right. The grey shadows represent copies of the states at the opposite end of the diagram, so the transitions on the far right and bottom loop around. This process satisfies FS3' and 4, FS2 with $CC := Pr$, and FS1 with $CC := WF(\mathcal{T})$, thereby improving process F_0 , and constituting the best CCS approximation of a fair scheduler seen so far. Yet, intuitively FS1 is not ensured at all, meaning that weak fairness is too strong an assumption. Nothing really prevents all the choices between r_2 and any other action a to be made in favour of a .



9 Formalising Requirements for Mutual Exclusion in Reactive LTL_{-X}

Define a process i participating in a mutual exclusion protocol to cycle through the stages *noncritical section*, *entry protocol*, *critical section*, and *exit protocol*, in that order, as explained in Section 6. Modelled as an LTS, its visible actions will be en_i , ln_i , ec_i and lc_i , of entering and leaving its (non)critical section. Put ln_i in B to make leaving the critical section a blockable action. The environment blocking it is my way of allowing the client process to stay in its noncritical section forever. This is the manner in which the requirement *Optionality* is captured in reactive temporal logic. On the other hand, ec_i should not be in B , for one does not consider the liveness property of a mutual exclusion protocol to be violated simply because the client process refuses to enter the critical section when allowed by the protocol. Likewise, en_i is not in B . Although exiting the critical section is in fact under control of the client process, it is assumed that it will not stay in the critical section forever. In the models of this paper this can be simply achieved by leaving lc_i outside B . Hence $B := \{ln_i \mid i = 1, \dots, N\}$.

My first requirement on mutual exclusion protocols P simply says that the actions en_i , ln_i , ec_i and lc_i have to occur in the right order:

$$(ME\ 1) \quad P \models \left((\neg act_i) \mathbf{W} ln_i \right) \wedge \mathbf{G} \left(ln_i \Rightarrow \left(ln_i \mathbf{U} \left((\neg act_i) \mathbf{W} ec_i \right) \right) \right) \wedge \mathbf{G} \left(ec_i \Rightarrow \left(ec_i \mathbf{U} \left((\neg act_i) \mathbf{W} lc_i \right) \right) \right) \\ \wedge \mathbf{G} \left(lc_i \Rightarrow \left(lc_i \mathbf{U} \left((\neg act_i) \mathbf{W} en_i \right) \right) \right) \wedge \mathbf{G} \left(en_i \Rightarrow \left(en_i \mathbf{U} \left((\neg act_i) \mathbf{W} ln_i \right) \right) \right)$$

for $i = 1, \dots, N$. Here $act_i := (ln_i \vee ec_i \vee lc_i \vee en_i)$.

The second is a formalisation of *Safety*, saying that only one process can be in its critical section at the same time:

$$(ME\ 2) \quad P \models \mathbf{G} \left(ec_i \Rightarrow \left((\neg ec_j) \mathbf{W} lc_i \right) \right)$$

for $i, j = 1, \dots, N$ with $i \neq j$. Both ME 1 and ME 2 would be unaffected by changing \models into \models^{CC} or \models_B^{CC} .

Requirement *Liveness* of Section 6 can be formalised as

$$(ME\ 3) \quad P \models_B^{CC} \mathbf{G} \left(ln_i \Rightarrow \mathbf{F} ec_i \right)$$

Here the choice of a completeness criterion is important. Finally, the following requirements are similar to *Liveness*, and state that from each section in the cycle of a Process i , the next section will in fact be reached. In regards to reaching the end of the noncritical section, this should be guaranteed only when assuming that the process wants to leave it critical section; hence ln_i is excepted from B .

$$(ME\ 4) \quad P \models_B^{CC} \mathbf{G} \left(ec_i \Rightarrow \mathbf{F} lc_i \right)$$

$$(ME\ 5) \quad P \models_B^{CC} \mathbf{G} \left(lc_i \Rightarrow \mathbf{F} en_i \right)$$

$$(ME\ 6) \quad P \models_{B \setminus \{ln_i\}}^{CC} \mathbf{F} ln_i \wedge \mathbf{G} \left(en_i \Rightarrow \mathbf{F} ln_i \right)$$

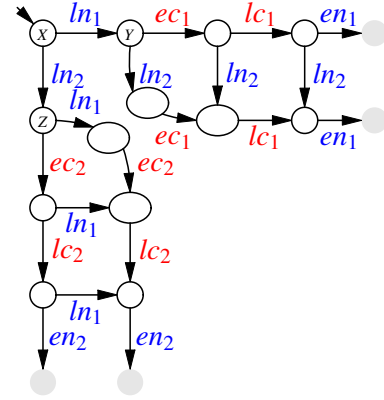
for $i = 1, \dots, N$.

The requirement *Speed independence* is automatically satisfied for models of mutual exclusion protocols rendered in any of the formalisms discussed in this paper, as these formalisms lack the expressiveness to make anything dependent on speed.

The following examples show that none of the above requirements are redundant.

- The CCS process $F_1 | F_2 | \dots | F_N$ with $F_i \stackrel{def}{=} ln_i . ec_i . lc_i . en_i . F_i$ satisfies all requirements, with $CC := J$, except for ME 2.
- The process $R_1 | R_2 | \dots | R_N$ with $R_i \stackrel{def}{=} ln_i . \mathbf{0}$ satisfies all requirements except for ME 3.
- In case $N = 2$, the process $ln_1 . ec_1 . ln_2 . ec_2 . \mathbf{0} + ln_2 . ec_2 . ln_1 . ec_1 . \mathbf{0}$ satisfies all requirements except for ME 4. The case $N > 2$ is only notationally more cumbersome. In the same spirit one finds counterexamples failing only on ME 5, or on the second conjunct of ME 6.
- The process $\mathbf{0}$ satisfies all requirements except for the first conjunct of ME 6.
- In case $N = 1$, the process X with $X \stackrel{def}{=} lc_1 . ec_1 . lc_1 . en_1 . ln_1 . X$ satisfies all requirements but ME 1.

The process X , a gatekeeper variant, given by $X \stackrel{\text{def}}{=} ln_1.Y + ln_2.Z$,
 $Y \stackrel{\text{def}}{=} ln_2.ec_1.lc_1.en_1.Z + ec_1.(ln_2.lc_1.en_1.Z + lc_1.(ln_2.en_1.Z + en_1.X))$
 $Z \stackrel{\text{def}}{=} ln_1.ec_2.lc_2.en_2.Y + ec_2.(ln_1.lc_2.en_2.Y + lc_2.(ln_1.en_2.Y + en_2.X))$
 is depicted on the right. It satisfies ME 1–5 with $CC := Pr$ and ME 6 with $CC := WF(\mathcal{T})$, where $LN_1, LN_2 \in \mathcal{T}$. It could be seen as a mediator that synchronises, on the actions ln_i , ec_i , lc_i and en_i , with the actual processes that need to exclusively enter their critical sections. Yet, it would not be commonly accepted as a valid mutual exclusion protocol, since nothing prevents it to never choose ln_2 when an alternative is available. This means that merely requiring weak fairness in ME 6 makes this requirement unacceptably weak. The problem with this protocol is that it ensures *Liveness* by making it hard for processes to leave their noncritical sections.



10 Reactive CTL

This section presents a reactive version of *Computation Tree Logic* (CTL) [4]. This shows that the ideas presented here are not specific to a linear-time logic. The syntax of CTL is

$$\varphi, \psi ::= p \mid \neg\varphi \mid \varphi \wedge \psi \mid \mathbf{EX}\varphi \mid \mathbf{AX}\varphi \mid \mathbf{EF}\varphi \mid \mathbf{AF}\varphi \mid \mathbf{EG}\varphi \mid \mathbf{AG}\varphi \mid \mathbf{E}\psi\mathbf{U}\varphi \mid \mathbf{A}\psi\mathbf{U}\varphi$$

with $p \in AP$ an atomic predicate. The relation \models between states s in a Kripke structure, CTL formulae φ and completeness criteria CC is inductively defined by

- $s \models^{CC} p$, with $p \in AP$, iff $(s, p) \in \models$,
- $s \models^{CC} \neg\varphi$ iff $s \not\models^{CC} \varphi$,
- $s \models^{CC} \varphi \wedge \psi$ iff $s \models^{CC} \varphi$ and $s \models^{CC} \psi$,
- $s \models^{CC} \mathbf{EX}\varphi$ iff there is a state s' with $s \rightarrow s'$ and $s' \models^{CC} \varphi$,
- $s \models^{CC} \mathbf{AX}\varphi$ iff for each state s' with $s \rightarrow s'$ one has $s' \models^{CC} \varphi$,
- $s \models^{CC} \mathbf{EF}\varphi$ iff some complete path starting in s contains a state s' with $s' \models^{CC} \varphi$,
- $s \models^{CC} \mathbf{AF}\varphi$ iff each complete path starting in s contains a state s' with $s' \models^{CC} \varphi$,
- $s \models^{CC} \mathbf{EG}\varphi$ iff all states s' on some complete path starting in s satisfy $s' \models^{CC} \varphi$,
- $s \models^{CC} \mathbf{AG}\varphi$ iff all states s' on all complete paths starting in s satisfy $s' \models^{CC} \varphi$,
- $s \models^{CC} \mathbf{E}\psi\mathbf{U}\varphi$ iff some complete path π starting in s contains a state s' with $s' \models^{CC} \varphi$, and each state s'' on π prior to s' satisfies $s'' \models^{CC} \psi$,
- $s \models^{CC} \mathbf{A}\psi\mathbf{U}\varphi$ iff each complete path π starting in s contains a state s' with $s' \models^{CC} \varphi$, and each state s'' on π prior to s' satisfies $s'' \models^{CC} \psi$.

Exactly as for $LTL_{\neg X}$, this allows the formulation of CTL judgements $s \models_B^{CC} \varphi$.

11 Conclusion

I proposed a formalism for making temporal judgements $P \models_B^{CC} \varphi$, with P a process specified in any formalism that admits a translation into LTSs, φ a temporal formula from a logic like LTL or CTL, CC a completeness criterion, stating which paths in the LTS model complete system runs, and B the set of actions that may be blocked by the user or environment of a system. I applied this formalism to unambiguously express the requirements defining fair schedulers and mutual exclusion protocols.

References

- [1] R. De Nicola & F.W. Vaandrager (1995): *Three Logics for Branching Bisimulation*. *Journal of the ACM* 42(2), pp. 458–487, doi:10.1145/201019.201032.
- [2] E.W. Dijkstra (1962 or 1963): *Over de sequentialiteit van processbeschrijvingen*. Available at <http://www.cs.utexas.edu/users/EWD/ewd00xx/EWD35.PDF>.
- [3] E.W. Dijkstra (1965): *Solution of a problem in concurrent programming control*. *Communications of the ACM* 8(9), p. 569, doi:10.1145/365559.365617.
- [4] E. Allen Emerson & Edmund M. Clarke (1982): *Using Branching Time Temporal Logic to Synthesize Synchronization Skeletons*. *Sci. Comput. Program.* 2(3), pp. 241–266, doi:10.1016/0167-6423(83)90017-5.
- [5] N. Francez (1986): *Fairness*. Springer, New York, doi:10.1007/978-1-4612-4886-6.
- [6] R.J. van Glabbeek (1993): *The Linear Time – Branching Time Spectrum II; The semantics of sequential systems with silent moves*. In E. Best, editor: *Proceedings CONCUR’93, 4th International Conference on Concurrency Theory*, Hildesheim, Germany, August 1993, LNCS 715, Springer, pp. 66–81, doi:10.1007/3-540-57208-2_6.
- [7] R.J. van Glabbeek (2018): *Is Speed-Independent Mutual Exclusion Implementable?* In S. Schewe & L. Zhang, editors: *Proceedings 29th International Conference on Concurrency Theory (CONCUR’18)*, Beijing, China, September 2018, *Leibniz International Proceedings in Informatics (LIPIcs)* 118, Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany, doi:10.4230/LIPIcs.CONCUR.2018.3.
- [8] R.J. van Glabbeek (2019): *Ensuring liveness properties of distributed systems: Open problems*. *Journal of Logical and Algebraic Methods in Programming* 109:100480, doi:10.1016/j.jlamp.2019.100480. Available at <http://arxiv.org/abs/1912.05616>.
- [9] R.J. van Glabbeek (2019): *Justness: A Completeness Criterion for Capturing Liveness Properties (extended abstract)*. In M. Bojańczyk & A. Simpson, editors: *Proceedings 22st International Conference on Foundations of Software Science and Computation Structures (FoSSaCS’19)*; held as part of the *European Joint Conferences on Theory and Practice of Software (ETAPS’19)*, Prague, Czech Republic, April 2019, LNCS 11425, Springer, pp. 505–522, doi:10.1007/978-3-030-17127-8_29.
- [10] R.J. van Glabbeek, U. Goltz & J.-W. Schicke (2011): *Abstract Processes of Place/Transition Systems*. *Information Processing Letters* 111(13), pp. 626–633, doi:10.1016/j.ipl.2011.03.013. Available at <http://arxiv.org/abs/1103.5916>.
- [11] R.J. van Glabbeek & P. Höfner (2015): *CCS: It’s not fair! - Fair schedulers cannot be implemented in CCS-like languages even under progress and certain fairness assumptions*. *Acta Informatica* 52(2-3), pp. 175–205, doi:10.1007/s00236-015-0221-6. Available at <http://arxiv.org/abs/1505.05964>.
- [12] R.J. van Glabbeek & P. Höfner (2015): *Progress, Fairness and Justness in Process Algebra*. Technical Report 8501, NICTA, Sydney, Australia. Available at <http://arxiv.org/abs/1501.03268>.
- [13] R.J. van Glabbeek & P. Höfner (2019): *Progress, Justness and Fairness*. *ACM Computing Surveys* 52(4):69, doi:10.1145/3329125. Available at <https://arxiv.org/abs/1810.07414>.
- [14] M. Huth & M.D. Ryan (2004): *Logic in Computer Science — Modelling and Reasoning about Systems*, 2nd edition. Cambridge University Press, doi:10.1017/CBO9780511810275.
- [15] L. Kleinrock (1964): *Analysis of A Time-Shared Processor*. *Naval Research Logistics Quarterly* 11(1), pp. 59–73, doi:10.1002/nav.3800110105.
- [16] D.E. Knuth (1966): *Additional comments on a problem in concurrent programming control*. *Communications of the ACM* 9(5), pp. 321–322, doi:10.1145/355592.365595.
- [17] L. Lamport (1974): *A New Solution of Dijkstra’s Concurrent Programming Problem*. *Communications of the ACM* 17(8), pp. 453–455, doi:10.1145/361082.361093.
- [18] L. Lamport (1983): *What good is temporal logic?* In R.E. Mason, editor: *Information Processing 83*, North-Holland, pp. 657–668.

- [19] R. Milner (1990): *Operational and algebraic semantics of concurrent processes*. In J. van Leeuwen, editor: *Handbook of Theoretical Computer Science*, chapter 19, Elsevier Science Publishers B.V. (North-Holland), pp. 1201–1242. Alternatively see *Communication and Concurrency*, Prentice-Hall, Englewood Cliffs, 1989, of which an earlier version appeared as *A Calculus of Communicating Systems*, LNCS 92, Springer, 1980, doi:10.1007/3-540-10235-3.
- [20] J. Nagle (1985): *On Packet Switches with Infinite Storage*. RFC 970, Network Working Group. Available at <http://tools.ietf.org/rfc/rfc970.txt>.
- [21] J. Nagle (1987): *On Packet Switches with Infinite Storage*. *IEEE Trans. Communications* 35(4), pp. 435–438, doi:10.1109/TCOM.1987.1096782.
- [22] G.L. Peterson (1981): *Myths About the Mutual Exclusion Problem*. *Information Processing Letters* 12(3), pp. 115–116, doi:10.1016/0020-0190(81)90106-X.
- [23] Amir Pnueli (1977): *The Temporal Logic of Programs*. In: *Foundations of Computer Science (FOCS '77)*, IEEE, pp. 46–57, doi:10.1109/SFCS.1977.32.
- [24] B.K. Szymański (1988): *A simple solution to Lamport's concurrent programming problem with linear wait*. In J. Lenfant, editor: *Proceedings of the 2nd international conference on Supercomputing, ICS 1988, Saint Malo, France, July 4-8, 1988*, ACM, pp. 621–626, doi:10.1145/55364.55425.