Computation Tree Logic
with Deadlock Detection

Rob van Glabbeek [1,2], Bas Luttik [3,4] & Nikola Trčka [3]

[1] National ICT Australia, Sydney, Australia
[2] School of Computer Science and Engineering, Univ. of New South Wales, Sydney
[3] Dept. of Math. & Comp. Sc., Technische Universiteit Eindhoven, The Netherlands
[4] CWI, The Netherlands

Abstract. We study the equivalence relation on states of labelled transition systems of satisfying the same formulas in Computation Tree Logic without the next state modality ($\mathsf{CTL}_{-\mathsf{X}}$). This relation is obtained by De Nicola & Vaandrager by translating labelled transition systems to Kripke structures, while lifting the totality restriction on the latter. They characterised it as divergence sensitive branching bisimulation equivalence. We find that this equivalence fails to be a congruence for interleaving parallel composition. The reason is that the proposed application of $\mathsf{CTL}_{-\mathsf{X}}$ to non-total Kripke structures lacks the expressiveness to cope with deadlock properties that are important in the context of parallel composition. We propose an extension of $\mathsf{CTL}_{-\mathsf{X}}$, or an alternative treatment of non-totality, that fills this hiatus. The equivalence induced by our extension is characterised as branching bisimulation equivalence with explicit divergence, which is, moreover, shown to be the coarsest congruence contained in divergence sensitive branching bisimulation equivalence.

1 Introduction

CTL^* [5] is a powerful temporal logic combining linear time and branching time modalities; it generalises the branching time temporal logic CTL [4]. CTL^* is interpreted in terms of Kripke structures. As the *next state* modality X is incompatible with abstraction of the notion of state, it is often excluded in high-level specifications. By $\mathsf{CTL}^*_{-\mathsf{X}}$ we denote CTL^* without this modality. To characterise the equivalence induced on states of Kripke structures by validity of $\mathsf{CTL}^*_{-\mathsf{X}}$ formulas, Browne, Clarke & Grumberg [2] defined the notion of *stuttering equivalence.* They proved that two states in a finite Kripke structure are stuttering equivalent if and only if they satisfy the same $\mathsf{CTL}^*_{-\mathsf{X}}$ formulas, and moreover, they established that this is already the case if and only if the two states satisfy the same $\mathsf{CTL}_{-\mathsf{X}}$ formulas.

There is an intuitive correspondence between the notions of stuttering equivalence on Kripke structures and *branching bisimulation equivalence* on labelled transition systems (LTSs) as defined in [8]. De Nicola & Vaandrager [3] have provided a framework for constructing natural translations between LTSs and Kripke structures in which this correspondence can be formalised. Stuttering equivalence corresponds in their framework to a *divergence sensitive* variant

of branching bisimulation equivalence, and conversely, branching bisimulation equivalence corresponds to a *divergence blind* variant of stuttering equivalence. The latter characterises the equivalence induced on states of Kripke structures by a divergence blind variant of validity of $\mathsf{CTL}^*_{-\mathsf{X}}$ formulas.

In [4, 5, 2] and other work on CTL^*, Kripke structures are required to be *total*, meaning that every state has an outgoing transition. These correspond with LTSs that are *deadlock-free*. In the world of LTSs requiring deadlock-freeness is considered a serious limitation, as deadlock is introduced by useful process algebraic operators like the *restriction* of CCS and the *synchronous parallel composition* of CSP. Conceptually, a deadlock may arise as the result of an unsuccessful synchronisation attempt between parallel components, and often one wants to verify that the result of a parallel composition is deadlock-free. This is possible only when working in a model of concurrency where deadlocks can be expressed.

Through the translations of [3] it is possible to define the validity of $\mathsf{CTL}^*_{-\mathsf{X}}$ formulas on states of LTSs. To apply $\mathsf{CTL}^*_{-\mathsf{X}}$-formulas to LTSs that may contain deadlocks we consider Kripke structures with deadlocks as well, and hence lift the requirement of totality. Following [3], we do so by using maximal paths instead of infinite paths in the definition of validity of $\mathsf{CTL}^*_{-\mathsf{X}}$ formulas. Barring further changes, this amounts to the addition of a self loop to every deadlock state. As a consequence, $\mathsf{CTL}^*_{-\mathsf{X}}$ formulas cannot see the difference between a state without outgoing transitions (a *deadlock*) and one whose only outgoing transition constitutes a self loop (a *livelock*), and accordingly a deadlock state is stuttering equivalent to a livelock state that satisfies the same atomic propositions. This paper will challenge the wisdom of this set-up.

We observe that for systems with deadlock, the divergence sensitive branching bisimulation equivalence of [3] fails to be a congruence for interleaving operators. We characterise the coarsest congruence contained in divergence sensitive branching bisimulation equivalence as the *branching bisimulation equivalence with explicit divergence*, introduced in [8]. This equivalence differs from divergence sensitive branching bisimulation equivalence in that it distinguishes deadlock and livelock. For deadlock-free systems the equivalences coincide.

That divergence sensitive branching bisimulation equivalence is not a congruence for interleaving operators means that there are properties of concurrent systems, pertaining to their deadlock behaviour, that cannot be expressed in $\mathsf{CTL}^*_{-\mathsf{X}}$, but that can be expressed in terms of the validity of a $\mathsf{CTL}^*_{-\mathsf{X}}$ formula on the result of putting these systems in a given context involving an interleaving operator. We find this unsatisfactory, and therefore propose an extension of $\mathsf{CTL}^*_{-\mathsf{X}}$ in which this type of property can be expressed directly. We obtain that two states are branching bisimulation equivalent with explicit divergence if and only if they satisfy the same formulas in the resulting logic.

Treating $\mathsf{CTL}_{-\mathsf{X}}$ in the same way leads either to an extension of $\mathsf{CTL}_{-\mathsf{X}}$ or, equivalently, to a modification of its semantics. The new semantics we propose for $\mathsf{CTL}_{-\mathsf{X}}$ is a valid extension of the original semantics [4] to non-total Kripke structures. It slightly differs from the semantics of [3] and it is arguably better suited to deal with deadlock behaviour.

Instead of extending $\mathsf{CTL}^*_{-\mathsf{X}}$ or modifying $\mathsf{CTL}_{-\mathsf{X}}$ we also achieve the same effect by amending the translation from LTSs to Kripke structures in such a way that every LTS maps to a total Kripke structure.

2 $\mathsf{CTL}^*_{-\mathsf{X}}$ and stuttering equivalence

We presuppose a set $\mathbf{AP}$ of *atomic propositions*. A *Kripke structure* is a tuple $(S, \mathscr{L}, \to)$ consisting of a set of states S, a *labelling function* $\mathscr{L}: S \to 2^{\mathbf{AP}}$ and a *transition relation* $\to \subseteq S \times S$. For the remainder of the section we fix a Kripke structure $(S, \mathscr{L}, \to)$.

A *finite path* from s is a finite sequence of states $s_0, \ldots, s_n$ such that $s = s_0$ and $s_k \longrightarrow s_{k+1}$ for all $0 \leq k < n$. An *infinite path* from s is an infinite sequence of states $s_0, s_1, s_2, \ldots$ such that $s = s_0$ and $s_k \longrightarrow s_{k+1}$ for all $k \in \omega$. A *path* is a finite or infinite path. A *maximal path* is an infinite path or a finite path $s_0, \ldots, s_n$ such that $\neg \exists s'.\ s_n \longrightarrow s'$. We write $\pi \trianglerighteq \pi'$ if the path π' is a suffix of the path π, and $\pi \triangleright \pi'$ if $\pi \trianglerighteq \pi'$ and $\pi \neq \pi'$.

Temporal properties of states in S are defined using $\mathsf{CTL}^*_{-\mathsf{X}}$ formulas.

Definition 1. The classes Φ of $\mathsf{CTL}^*_{-\mathsf{X}}$ *state formulas* and Ψ of $\mathsf{CTL}^*_{-\mathsf{X}}$ *path formulas* are generated by the following grammar:

$$\varphi ::= p \mid \neg\varphi \mid \bigwedge \Phi' \mid \exists\psi \qquad \psi ::= \varphi \mid \neg\psi \mid \bigwedge \Psi' \mid \psi \,\mathsf{U}\, \psi$$

with $p \in \mathbf{AP}$, $\varphi \in \Phi$, $\Phi' \subseteq \Phi$, $\psi \in \Psi$ and $\Psi' \subseteq \Psi$.

In case the cardinality of the set of states of our Kripke structure is less than some infinite cardinal κ,[1] we may require that $|\Phi'| < \kappa$ and $|\Psi'| < \kappa$ in conjunctions, thus obtaining a *set* of formulas rather than a proper class. Normally, S is required to be finite, and accordingly $\mathsf{CTL}^*_{-\mathsf{X}}$ admits finite conjunctions only.

Definition 2. We define when a $\mathsf{CTL}^*_{-\mathsf{X}}$ state formula φ is *valid* in a state s (notation: $s \models \varphi$) and when a $\mathsf{CTL}^*_{-\mathsf{X}}$ path formula ψ is *valid* on a maximal path π (notation: $\pi \models \psi$) by simultaneous induction as follows:

- $s \models p$ iff $p \in \mathscr{L}(s)$;
- $s \models \neg\varphi$ iff $s \not\models \varphi$;
- $s \models \bigwedge \Phi'$ iff $s \models \varphi$ for all $\varphi \in \Phi'$;
- $s \models \exists\psi$ iff there exists a maximal path π from s such that $\pi \models \psi$;
- $\pi \models \varphi$ iff s is the first state of π and $s \models \varphi$;
- $\pi \models \neg\psi$ iff $\pi \not\models \psi$;
- $\pi \models \bigwedge \Psi'$ iff $\pi \models \psi$ for all $\psi \in \Psi'$; and
- $\pi \models \psi \,\mathsf{U}\, \psi'$ iff there exists a suffix π' of π such that $\pi' \models \psi'$, and $\pi'' \models \psi$ for all $\pi \trianglerighteq \pi'' \triangleright \pi'$.

A formula $\psi \,\mathsf{U}\, \psi'$ says that, along a given path, ψ holds *until* ψ' holds. One writes $\top$ for the empty conjunction (which is always valid), $\mathsf{F}\psi$ for $\top \,\mathsf{U}\, \psi$ ("ψ will hold *eventually*") and $\mathsf{G}\psi$ for $\neg\mathsf{F}\neg\psi$ ("ψ holds *always* (along a path)").

---

[1] In fact it suffices to require that for every state s the cardinality of the set of states reachable from s is less than κ.

The above is the standard, *divergence sensitive* interpretation of $\mathsf{CTL}^*_{-\mathsf{X}}$ [5, 2], but extended to Kripke structures that are not required to be total. Following [3], this is achieved by using maximal paths in the definition of validity of $\mathsf{CTL}^*_{-\mathsf{X}}$ formulas, instead of the traditional use of infinite paths [5, 2]. The resulting definition generalises the traditional one, because for total Kripke structures a path is maximal iff it is infinite.

An equivalent way of thinking of this generalisation of $\mathsf{CTL}^*_{-\mathsf{X}}$ to non-total Kripke structures is by means of a transformation that makes a Kripke structure K total by the addition of a self loop $s \longrightarrow s$ to every deadlock state s, together with the convention that a formula is valid in a state of K iff it is valid in the same state of the total Kripke structure obtained by this transformation. It is not hard to check that this yields the same notion of validity as our Definition 2.

The *divergence blind* interpretation of [3] (notation: $s \models_{db} \varphi$ and $\pi \models_{db} \psi$) is obtained by dropping the word "maximal" in the fourth clause of Definition 2.

Definition 3. A *colouring* is a function $\mathcal{C} : S \to \mathbf{C}$, for $\mathbf{C}$ any set of *colours*.

Given a colouring $\mathcal{C}$ and a (finite or infinite) path $\pi = s_0, s_1, s_2, \ldots$ from s, let $\mathcal{C}(\pi)$ be the sequence of colours obtained from $\mathcal{C}(s_0), \mathcal{C}(s_1), \mathcal{C}(s_2), \ldots$ by contracting all its (finite or infinite) maximal consecutive subsequences $C, C, C, \ldots$ to C. The sequence $\mathcal{C}(\pi)$ is called a $\mathcal{C}$-*coloured trace* of s; it is *complete* if π is maximal.

A colouring $\mathcal{C}$ is *[fully] consistent* if two states of the same colour always satisfy the same atomic propositions and have the same [complete] $\mathcal{C}$-coloured traces. Two states s and t are *divergence blind stuttering equivalent*, notation $s \approx_{dbs} t$, if there exists a consistent colouring $\mathcal{C}$ such that $\mathcal{C}(s) = \mathcal{C}(t)$. They are *(divergence sensitive) stuttering equivalent*, notation $s \approx_s t$, if there exists a fully consistent colouring $\mathcal{C}$ such that $\mathcal{C}(s) = \mathcal{C}(t)$. The difference between $\approx_{dbs}$ and $\approx_s$ is illustrated in the following example.

Example 1. Consider the Kripke structure and its colouring depicted in Figure 1a. This colouring is consistent, implying $s \approx_{dbs} t \approx_{dbs} u$ and $x \approx_{dbs} y$, but it is not fully consistent because state t has a complete trace ⬤ while u does not. Note that t has, due to the self loop, a complete coloured trace that consists of just the colour of a p-labelled state, whereas the unique complete coloured trace of u contains the colour of a q-labelled state too. Since a consistent colouring assigns different colours to states with different labels, every fully consistent colouring must assign different colours to states t and u, i.e. it must be that $t \not\approx_s u$. One such colouring is given in Figure 1b. This colouring shows that $x \approx_s y$.

Lemma 1. *Let $\mathcal{C}$ be a colouring such that two states with the same colour satisfy the same atomic propositions and have the same $\mathcal{C}$-coloured traces of length two. Then $\mathcal{C}$ is consistent.*

Proof. Suppose $\mathcal{C}(s_0) = \mathcal{C}(t_0)$ and $C_0, C_1, C_2, \ldots$ is an infinite coloured trace of s_0. Then, for $i > 0$, there are states s_i and finite paths π_i from s_{i-1} to s_i, such that $\mathcal{C}(\pi_i) = C_{i-1}, C_i$. With induction on i, for $i > 0$ we find states t_i with
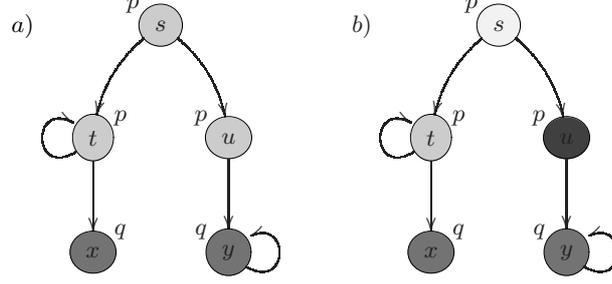
Fig. 1. Difference between a) $\approx_{dbs}$ and b) $\approx_s$.

$C(s_i) = C(t_i)$ and finite paths ρ_i from t_{i-1} to t_i such that $\mathcal{C}(\rho_i) = C_{i-1}, C_i$. Namely, the assumption about $\mathcal{C}$ allows us to find ρ_i given t_{i-1}, and then t_i is defined as the last state of ρ_i. Concatenating all the paths ρ_i yields an infinite path ρ from t_0 with $\mathcal{C}(\rho) = C_0, C_1, C_2, \ldots$.

The case that $\mathcal{C}(s_0) = \mathcal{C}(t_0)$ and $C_0, \ldots, C_n$ is a finite coloured trace of s_0 goes likewise. □

Lemma 2. *Let $\mathcal{C}$ be a colouring such that two states with the same colour satisfy the same atomic propositions and have the same $\mathcal{C}$-coloured traces of length two, and the same complete $\mathcal{C}$-coloured traces of length one. Then $\mathcal{C}$ is fully consistent.*

Proof. Suppose $\mathcal{C}(s) = \mathcal{C}(t)$ and σ is a complete $\mathcal{C}$-coloured trace of s. Then $\sigma = \mathcal{C}(\pi)$ for a maximal path π from s. By Lemma 1, σ is also a $\mathcal{C}$-coloured trace of t. It remains to show that it is a *complete* $\mathcal{C}$-coloured trace of t. Let ρ be a path from t with $\mathcal{C}(\rho) = \sigma$. If ρ is infinite, we are done. Otherwise, let t' be the last state of ρ. Then $\mathcal{C}(t')$ is the last colour of σ. Therefore, there is a state s' on π such that the suffix π' of π starting from s' is a maximal path with $\mathcal{C}(\pi') = \mathcal{C}(s') = \mathcal{C}(t')$. By the assumption about $\mathcal{C}$, $\mathcal{C}(t')$ must also be a complete $\mathcal{C}$-coloured trace of t', i.e. there is a maximal path ρ' from t' with $\mathcal{C}(\rho') = \mathcal{C}(t')$. Concatenating ρ and ρ' yields a maximal path ρ'' from t with $\mathcal{C}(\rho'') = \sigma$. □

The following two theorems were proved in [3] and [2], respectively, for states s and t in a finite Kripke structure. Here we drop the finiteness restriction.

Theorem 1. $s \approx_{dbs} t$ *iff* $s \models_{db} \varphi \Leftrightarrow t \models_{db} \varphi$ *for all* CTL^*_{-X} *state formulas* φ.

Proof. "Only if": Let $\mathcal{C}$ be a consistent colouring. With structural induction on φ and ψ we show that

$$\mathcal{C}(s) = \mathcal{C}(t) \Rightarrow (s \models_{db} \varphi \Leftrightarrow t \models_{db} \varphi) \ \text{ and } \ \mathcal{C}(\pi) = \mathcal{C}(\rho) \Rightarrow (\pi \models_{db} \psi \Leftrightarrow \rho \models_{db} \psi).$$

The case $\varphi = p$ for $p \in \mathbf{AP}$ follows immediately from Definition 3. The cases $\varphi = \neg \varphi'$ and $\varphi = \bigwedge \Phi'$ follow immediately from the induction hypothesis.

Suppose $\mathcal{C}(s) = \mathcal{C}(t)$ and $s \models_{db} \exists \psi$. Then there exists a path π from s such that $\pi \models_{db} \psi$. $\mathcal{C}(\pi)$ is a coloured trace of s, and hence of t. Thus there must be a path ρ from t with $\mathcal{C}(\pi) = \mathcal{C}(\rho)$. By induction, $\rho \models_{db} \psi$. Hence, $t \models_{db} \exists \psi$.

The case $\varphi \in \Phi$ follows since the first states of two paths with the same colour also have the same colour. The cases $\psi = \neg\psi'$ and $\psi = \bigwedge \Psi'$ follow immediately from the induction hypothesis.

Finally, suppose $\mathcal{C}(\pi) = \mathcal{C}(\rho)$ and $\pi \models_{db} \psi \, \mathsf{U} \, \psi'$. Then there exists a suffix π' of π such that $\pi' \models_{db} \psi'$ and $\pi'' \models_{db} \psi$ for all $\pi \trianglerighteq \pi'' \triangleright \pi'$. As $\mathcal{C}(\pi) = \mathcal{C}(\rho)$, there must be a suffix ρ' of ρ such that $\mathcal{C}(\pi') = \mathcal{C}(\rho')$ and for every path ρ'' such that $\rho \trianglerighteq \rho'' \triangleright \rho'$ there exists a path π'' with $\pi \trianglerighteq \pi'' \triangleright \pi'$ such that $\mathcal{C}(\pi'') = \mathcal{C}(\rho'')$. By induction, this implies $\rho' \models_{db} \psi'$ and $\rho'' \models_{db} \psi$ for all $\rho \trianglerighteq \rho'' \triangleright \rho'$. Hence $\rho \models_{db} \psi \, \mathsf{U} \, \psi'$.

"If": Let $\mathcal{C}$ be the colouring given by $\mathcal{C}(s) = \{\varphi \in \Phi \mid s \models_{db} \varphi\}$. It suffices to show that $\mathcal{C}$ is consistent. So suppose $\mathcal{C}(s) = \mathcal{C}(t)$. Trivially, s and t satisfy the same atomic propositions. By Lemma 1 it remains to show that s and t have the same coloured traces of length two. Suppose s has a coloured trace C, D. Let $s_0, \ldots, s_k$ be a path from s such that $\mathcal{C}(s_i) = C$ for $0 \le i < k$ and $\mathcal{C}(s_k) = D \ne C$. Let

$$\mathcal{U} = \{u \mid \text{there is a path from } t \text{ to } u \text{ and } \mathcal{C}(u) \ne C\},$$
$$\mathcal{V} = \{v \mid \text{there is a path from } t \text{ to } v \text{ and } \mathcal{C}(v) \ne D\}.$$

For every $u \in \mathcal{U}$ pick a $\mathsf{CTL}^*_{-\mathsf{X}}$ formula $\varphi_u \in C - \mathcal{C}(u)$ (using negation on a formula in $\mathcal{C}(u) - C$ if needed), and for every $v \in \mathcal{V}$ pick a $\mathsf{CTL}^*_{-\mathsf{X}}$ formula $\varphi'_v \in D - \mathcal{C}(v)$. Now $s \models_{db} \exists(\bigwedge_{u \in \mathcal{U}} \varphi_u) \, \mathsf{U} \, (\bigwedge_{v \in \mathcal{V}} \varphi'_v)$ and, as $\mathcal{C}(s) = \mathcal{C}(t)$, also $t \models_{db} \exists(\bigwedge_{u \in \mathcal{U}} \varphi_u) \, \mathsf{U} \, (\bigwedge_{v \in \mathcal{V}} \varphi'_v)$. Thus, there is a path $t_0, \ldots, t_\ell$ from t such that $t_\ell \models_{db} \bigwedge_{v \in \mathcal{V}} \varphi'_v$ and $t_j \models_{db} \bigwedge_{u \in \mathcal{U}} \varphi_u$ for $0 \le j < \ell$. It follows that $t_\ell \notin \mathcal{V}$ and $t_j \notin \mathcal{U}$ for $0 \le j < \ell$. Hence $\mathcal{C}(t_\ell) = D$ and $\mathcal{C}(t_j) = C$ for $0 \le j < \ell$, so C, D is also a coloured trace of t.   □

Theorem 2. $s \approx_s t$ iff $s \models \varphi \Leftrightarrow t \models \varphi$ for all $\mathsf{CTL}^*_{-\mathsf{X}}$ state formulas φ.

Proof. "Only if" goes exactly as in the previous proof, reading $\models$ for $\models_{db}$, but requiring $\mathcal{C}$ to be *fully* consistent and, in the second paragraph, the paths π and ρ to be maximal and $\mathcal{C}(\pi)$ to be a *complete* coloured trace of s and t.

"If" goes as in the previous proof, but this time we have to show that $\mathcal{C}$ is *fully* consistent. Thus, applying Lemma 2, and assuming $\mathcal{C}(s) = \mathcal{C}(t)$, we additionally have to show that s and t have the same complete coloured traces of length one. Let π be a maximal path from s with $\mathcal{C}(\pi) = C$. Let

$$\mathcal{U} = \{u \mid \text{there is a path from } t \text{ to } u \text{ and } \mathcal{C}(u) \ne C\}.$$

For every $u \in \mathcal{U}$ pick a $\mathsf{CTL}^*_{-\mathsf{X}}$ formula $\varphi_u \in C - \mathcal{C}(u)$. Now $s \models \exists\mathsf{G}(\bigwedge_{u \in \mathcal{U}} \varphi_u)$ and, as $\mathcal{C}(s) = \mathcal{C}(t)$, also $t \models \exists\mathsf{G}(\bigwedge_{u \in \mathcal{U}} \varphi_u)$. Thus, there is a maximal path ρ from t such that $t' \models \bigwedge_{u \in \mathcal{U}} \varphi_u$ for all states t' in ρ. It follows that $t' \notin \mathcal{U}$. Hence $\mathcal{C}(t') = C$ and thus $\mathcal{C}(\rho) = C$.   □

Corollary 1. $\approx_{dbs}$ and $\approx_s$ are equivalence relations.   □

Note that, for any colouring $\mathcal{C}$, the $\mathcal{C}$-coloured traces of a state s are completely determined by the complete $\mathcal{C}$-coloured traces of s, namely as their prefixes.

Hence, any colouring that is fully consistent is certainly consistent, and thus $\approx_s$ is a finer (i.e. smaller, more discriminating) equivalence relation than $\approx_{dbs}$.

Above, the divergence blind interpretation of $\mathsf{CTL}^*_{-\mathsf{X}}$ is defined by using paths instead of maximal paths. It can equivalently be defined in terms of a transformation on Kripke structures, namely the addition of a self loop $s \longrightarrow s$ for every state s. Now $s \approx_{dbs} t$ holds in a certain Kripke structure iff $s \approx_s t$ holds in the Kripke structure obtained by adding all these self loops. This is because the colour of a path doesn't change when self-loops are added to it, and up to self loops any path is maximal. Likewise, $s \models_{db} \varphi$ in the original Kripke structure iff $s \models \varphi$ in the modified one.

Just like $\approx_{dbs}$ can be expressed in terms of $\approx_s$ by means of a transformation on Kripke structures, by means of a different transformation, at least for finite Kripke structures, $\approx_s$ can be expressed in terms of $\approx_{dbs}$. This is done in [3], Definitions 3.2.6 and 3.2.7.

3   Branching bisimulation in terms of coloured traces

We presuppose a set A of *actions* with a special element $\tau \in A$. A *labelled transition system* (LTS) is a structure $(S, \rightarrow)$ consisting of a set of states S and a *transition relation* $\rightarrow \subseteq S \times A \times S$. For the remainder of the section we fix an LTS $(S, \rightarrow)$. We write $s \xrightarrow{a} s'$ for $(s, a, s') \in \rightarrow$.

A *path* from s is an alternating sequence $s_0, a_1, s_1, a_2, s_2, \ldots$ of states and actions, ending with a state if the sequence is finite, such that $s = s_0$ and $s_{k-1} \xrightarrow{a_k} s_k$ for all relevant $k > 0$. A *maximal path* is an infinite path or a finite path $s_0, a_1, s_1, a_2, \ldots, a_n, s_n$ such that $\neg \exists a, s'. s_n \xrightarrow{a} s'$. We write $\pi \trianglerighteq \pi'$ if the path π' is a suffix of the path π, and $\pi \triangleright \pi'$ if $\pi \trianglerighteq \pi'$ and $\pi \neq \pi'$.

Definition 4. A *colouring* is a function $\mathcal{C} : S \rightarrow \mathbf{C}$, for $\mathbf{C}$ any set of *colours*.

For $\pi = s_0, a_1, s_1, a_2, s_2, \ldots$ a path from s, let $\mathcal{C}(\pi)$ be the alternating sequence of colours and actions obtained from $\mathcal{C}(s_0), a_1, \mathcal{C}(s_1), a_2, \mathcal{C}(s_2), \ldots$ by contracting all (finite or infinite) maximal consecutive subsequences $C, \tau, C, \tau, C, \ldots$ to C. The sequence $\mathcal{C}(\pi)$ is called a $\mathcal{C}$-*coloured trace* of s; it is *complete* if π is maximal; it is *divergent* if it is finite whilst π is infinite.

A colouring $\mathcal{C}$ is *[fully] consistent* if two states of the same colour always have the same [complete] $\mathcal{C}$-coloured traces. Two states s and t are *(divergence blind) branching bisimulation equivalent*, notation $s \underline{\leftrightarrow}_b t$, if there exists a consistent colouring $\mathcal{C}$ such that $\mathcal{C}(s) = \mathcal{C}(t)$. They are *divergence sensitive branching bisimulation equivalent*, notation $s \underline{\leftrightarrow}_b^\lambda t$, if there exists a fully consistent colouring $\mathcal{C}$ such that $\mathcal{C}(s) = \mathcal{C}(t)$.

A consistent colouring *preserves divergence* if two states of the same colour always have the same divergent $\mathcal{C}$-coloured traces. Two states s and t are *branching bisimulation equivalent with explicit divergence*, notation $s \underline{\leftrightarrow}_b^\Delta t$, if there exists a consistent, divergence preserving colouring $\mathcal{C}$ with $\mathcal{C}(s) = \mathcal{C}(t)$.

The difference between $\underline{\leftrightarrow}_b$, $\underline{\leftrightarrow}_b^\lambda$, and $\underline{\leftrightarrow}_b^\Delta$ is illustrated in the following example.
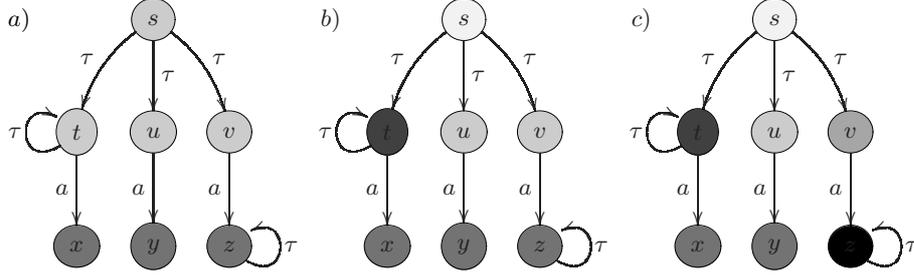
Fig. 2. Difference between a) $\underline{\leftrightarrow}_b$, b) $\underline{\leftrightarrow}_b^{\lambda}$, and c) $\underline{\leftrightarrow}_b^{\Delta}$.

Example 2. Consider first the LTS and its colouring depicted in Figure 2a. This colouring is consistent and we have $s \underline{\leftrightarrow}_b t \underline{\leftrightarrow}_b u \underline{\leftrightarrow}_b v$ and $x \underline{\leftrightarrow}_b y \underline{\leftrightarrow}_b z$. It is not fully consistent because state s has a complete trace ⬤ whereas t has not. It is easy to see that every fully consistent colouring must assign different colours to states t and u, and so that $t \not\underline{\leftrightarrow}_b^{\lambda} u$. One such colouring is given in Figure 2b and it shows that $u \underline{\leftrightarrow}_b^{\lambda} v$ and $x \underline{\leftrightarrow}_b^{\lambda} y \underline{\leftrightarrow}_b^{\lambda} z$. Note, however, that this colouring, although fully consistent, does not preserve divergence. State v has a divergent trace ⬤ a ⬤ whereas u has not, and similarly state z has a divergent trace ⬤ whereas y has not. Any colouring that preserves divergence must assign different colours to states u and v and to states y and z, meaning that $u \not\underline{\leftrightarrow}_b^{\Delta} v$ and $y \not\underline{\leftrightarrow}_b^{\Delta} z$. One such colouring is given in Figure 2c. It shows that $x \underline{\leftrightarrow}_b^{\Delta} y$. In fact, these are the only two (different) states that are branching bisimulation equivalent with explicit divergence.

In the definition of $\underline{\leftrightarrow}_b^{\Delta}$ above, consistency and preservation of divergence appear as two separate properties of colourings. Instead we could have integrated them by adding an extra bit (Δ) at the end of those finite coloured traces that stem from infinite paths. Likewise, $\underline{\leftrightarrow}_b^{\lambda}$ could have been defined by adding such an extra bit at the end of those finite coloured traces that stem from maximal paths.

Lemmas 1 and 2 about colourings on Kripke structures apply to labelled transition systems as well. The proofs are essentially the same.

Lemma 3. *Let $\mathcal{C}$ be a colouring such that two states with the same colour have the same $\mathcal{C}$-coloured traces of length three (i.e. colour - action - colour). Then $\mathcal{C}$ is consistent.* □

Lemma 4. *Let $\mathcal{C}$ be a consistent colouring such that two states with the same colour have the same complete $\mathcal{C}$-coloured traces of length one. Then $\mathcal{C}$ is fully consistent.* □

Lemma 5. *Let $\mathcal{C}$ be a consistent colouring such that two states with the same colour have the same divergent $\mathcal{C}$-coloured traces of length one. Then $\mathcal{C}$ preserves divergence.*

Proof. Exactly like the proof of Lemma 2, but letting σ be a *divergent* $\mathcal{C}$-coloured trace of s; π, π' *infinite* paths; $\mathcal{C}(t')$ a *divergent* $\mathcal{C}$-coloured trace of t'; and ρ', ρ'' *infinite* paths.   □

Branching bisimulation equivalence and branching bisimulation equivalence with explicit divergence were originally defined in Van Glabbeek & Weijland [8]. There, only *finite* coloured traces are considered, and a consistent colouring was defined as a colouring with the property that two states have the same colour only if they have the same finite coloured traces. By Lemma 3 this yields the same concept of consistent colouring as Definition 4 above.

In [8], a consistent colouring is said to *preserve divergence* if no divergent state has the same colour as a nondivergent state. Here a state s is *divergent* if it is the starting point of an infinite path of which all nodes have the same colour. This is the case if s has a divergent coloured trace of length one. Now Lemma 5 says that the definition of preservation of divergence from [8] agrees with the one proposed above. Hence the concepts of branching bisimulation and branching bisimulation with explicit divergence of [8] agree with ours.

Theorem 3. $\underline{\leftrightarrow}_b$, $\underline{\leftrightarrow}_b^\lambda$ *and* $\underline{\leftrightarrow}_b^\Delta$ *are equivalence relations.*

Proof. We show the proof for $\underline{\leftrightarrow}_b$; the other two cases proceed likewise.

We will regard any equivalence relation on S as a colouring, the colour of a state being its equivalence class. Conversely, any colouring can be considered as an equivalence relation on states.

The diagonal on S (i.e., the binary relation $\{(s,s) \mid s \in S\}$) is a consistent colouring, so $\underline{\leftrightarrow}_b$ is reflexive. That $\underline{\leftrightarrow}_b$ is symmetric is immediate from the required symmetry of colourings.

To prove that $\underline{\leftrightarrow}_b$ is transitive, suppose $s \underline{\leftrightarrow}_b t$ and $t \underline{\leftrightarrow}_b u$. So there exist consistent colourings $\mathcal{C}$ and $\mathcal{D}$ with $\mathcal{C}(s) = \mathcal{C}(t)$ and $\mathcal{D}(t) = \mathcal{D}(u)$. Let $\mathcal{E}$ be the finest equivalence relation containing $\mathcal{C}$ and $\mathcal{D}$. Then $\mathcal{E}(s) = \mathcal{E}(t) = \mathcal{E}(u)$. It suffices to show that $\mathcal{E}$ is consistent.

First of all note that the $\mathcal{E}$-colour of a state is completely determined by its $\mathcal{C}$-colour, as well as by its $\mathcal{D}$-colour: $\mathcal{C}(p) = \mathcal{C}(q) \Rightarrow \mathcal{E}(p) = \mathcal{E}(q)$ and $\mathcal{D}(p) = \mathcal{D}(q) \Rightarrow \mathcal{E}(p) = \mathcal{E}(q)$ for all $p, q \in S$. Thus, if two states have the same sets of $\mathcal{C}$-coloured traces or the same sets of $\mathcal{D}$-coloured traces, they must also have the same sets of $\mathcal{E}$-coloured traces.

Suppose $\mathcal{E}(p) = \mathcal{E}(q)$. Then there must be a sequence of states $(p_i)_{0 \leq i \leq n}$ such that $p = p_0$, $q = p_n$ and for $0 \leq i < n$ we have either $\mathcal{C}(p_i) = \mathcal{C}(p_{i+1})$ or $\mathcal{D}(p_i) = \mathcal{D}(p_{i+1})$. As $\mathcal{C}$ and $\mathcal{D}$ are consistent colourings, p_i and p_{i+1} have the same $\mathcal{C}$-coloured traces or the same $\mathcal{D}$-coloured traces. In either case they also have the same $\mathcal{E}$-coloured traces. This holds for $0 \leq i < n$, so p and q have the same $\mathcal{E}$-coloured traces. Thus $\mathcal{E}$ is consistent.   □

Lemma 6. *Let $\mathcal{C}$ be a consistent colouring and $s \in S$. Then the complete $\mathcal{C}$-coloured traces of s consist of the $\mathcal{C}$-coloured traces of s that are infinite, divergent, or maximal, in the sense that they cannot be extended.*

Proof. By definition, infinite and divergent $\mathcal{C}$-coloured traces of s are complete. Let σ be a maximal $\mathcal{C}$-coloured trace of s, and let π be a path from s such that $\mathcal{C}(\pi) = \sigma$. Let π' be an extension of π to a maximal path. As σ is a maximal $\mathcal{C}$-coloured trace, in the sense that it cannot be extended, we have $\mathcal{C}(\pi') = \sigma$. Hence σ is a complete $\mathcal{C}$-coloured trace of s.

Now let σ be a complete $\mathcal{C}$-coloured trace of s that is not infinite, nor a divergent $\mathcal{C}$-coloured trace of s. In that case $\sigma = \mathcal{C}(\pi)$ for π a finite maximal path from s. Let t be the last state of π. We have $\neg\exists a, t'.\ t \xrightarrow{a} t'$. Suppose, towards a contradiction, that σ is not maximal, i.e. there is a path π' from s such that $\mathcal{C}(\pi')$ is a proper extension of σ. Then there must be a state u on π' with $\mathcal{C}(u) = \mathcal{C}(t)$, such that u has a coloured trace σ' of length > 1, which is a suffix of $\mathcal{C}(\pi')$. As $\mathcal{C}$ is consistent, σ' is also a coloured trace of t, contradicting $\neg\exists a, t'.\ t \xrightarrow{a} t'$. $\qquad\square$

As for Kripke structures, for any colouring $\mathcal{C}$, the $\mathcal{C}$-coloured traces of a state s are the prefixes of the complete $\mathcal{C}$-coloured traces of s. Moreover, Lemma 6 says that the complete $\mathcal{C}$-coloured traces of a state s are completely determined by the $\mathcal{C}$-coloured traces of s together with the divergent $\mathcal{C}$-coloured traces of s. Hence, any colouring that is consistent and preserves divergence is also fully consistent. Therefore, $\underline{\leftrightarrow}_b^\Delta$ is finer than $\underline{\leftrightarrow}_b^\lambda$, which is finer than $\underline{\leftrightarrow}_b$.

The difference between $\underline{\leftrightarrow}_b^\lambda$ and $\underline{\leftrightarrow}_b^\Delta$ is that only the latter sees the difference between those maximal finite coloured traces that stem from finite paths (ending in *deadlock*) and those that stem from infinite paths (ending in *livelock*). For *deadlock-free* LTSs (having no states s with $\neg\exists a, s'.\ s \xrightarrow{a} s'$) the equivalences $\underline{\leftrightarrow}_b^\lambda$ and $\underline{\leftrightarrow}_b^\Delta$ coincide.

4   Translating between Kripke structures and labelled transition systems

We presuppose a set A of *actions* with a special element $\tau \in A$, and a set $\mathbf{AP}$ of *atomic propositions*. A *doubly labelled transition system* ($\mathrm{L}^2\mathrm{TS}$) is a structure $(S, \mathscr{L}, \rightarrow)$ that consists of a set of states S, a *labelling function* $\mathscr{L} : S \rightarrow 2^{\mathbf{AP}}$ and a *labelled transition relation* $\rightarrow \subseteq S \times A \times S$. From an $\mathrm{L}^2\mathrm{TS}$ one naturally obtains an LTS by omitting the labelling function $\mathscr{L}$, and a Kripke structure by replacing the labelled transition relation by one from which the labels are omitted. We call these the LTS or Kripke structure *associated* to the $\mathrm{L}^2\mathrm{TS}$. An $\mathrm{L}^2\mathrm{TS}$ $(S, \mathscr{L}, \rightarrow)$ is *consistent* if it satisfies the following three conditions:

(i)   if $s \xrightarrow{a} t$, then $(\mathscr{L}(s) = \mathscr{L}(t)$ iff $a = \tau)$;
(ii) if $s \xrightarrow{a} t$, $s' \xrightarrow{a} t'$ and $\mathscr{L}(s) = \mathscr{L}(s')$, then $\mathscr{L}(t) = \mathscr{L}(t')$; and
(iii) if $s \xrightarrow{a} t$, $s' \xrightarrow{b} t'$, $\mathscr{L}(s) = \mathscr{L}(s')$ and $\mathscr{L}(t) = \mathscr{L}(t')$, then $a = b$.

Consistent $\mathrm{L}^2\mathrm{TS}$s were introduced in De Nicola & Vaandrager [3] for studying relationships between notions defined for Kripke structures and notions defined for LTSs. Condition (iii) says that the label of a transition is fully determined by the labels of its source and target state, whereas condition (ii) entails that

the label of a state t reachable from a state s is completely determined by the label of s and the sequence of labels of the transitions leading from s to t.

Example 3. The three L^{2}TSs from Figure 3a are not consistent because they violate condition (i), (ii), and (iii) respectively; the L^{2}TS in Figure 3b is consistent.
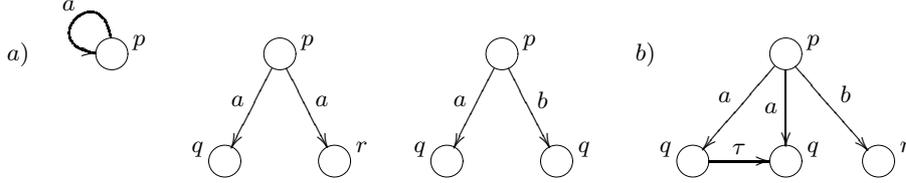

Fig. 3. a) Three inconsistent L^{2}TSs and b) a consistent L^{2}TS.

Many semantic equivalences on LTSs, such as $\underline{\leftrightarrow}_b$, $\underline{\leftrightarrow}_b^\lambda$ and $\underline{\leftrightarrow}_b^\Delta$, are considered in the literature; for an overview see [6].

Definition 5. Any semantic equivalence $\sim$ on LTSs extends to L^{2}TSs by declaring, for all states s and t in an L^{2}TS, that $s \sim t$ iff $\mathscr{L}(s) = \mathscr{L}(t)$ and $s \sim t$ in the associated LTS.

Any semantic equivalence $\sim$ on Kripke structures extends to L^{2}TSs by declaring, for all states s and t in an L^{2}TS, that $s \sim t$ iff $s \sim t$ in the associated Kripke structure.

The following theorem was proved in [3] for finite consistent L^{2}TSs. Here we drop the finiteness restriction.

Theorem 4. *On a consistent L^2TS, $\approx_{dbs}$ equals $\underline{\leftrightarrow}_b$, and $\approx_s$ equals $\underline{\leftrightarrow}_b^\lambda$.*

Proof. Suppose $s \approx_{dbs} t$ [or $s \approx_s t$]. Then there is a colouring $\mathcal{C}$ on the states of the L^{2}TS that is [fully] consistent on the associated Kripke structure K and satisfies $\mathcal{C}(s) = \mathcal{C}(t)$. By definition, this entails $\mathscr{L}(s) = \mathscr{L}(t)$. It remains to show that $\mathcal{C}$ is [fully] consistent on the associated LTS L. So let $\mathcal{C}(p) = \mathcal{C}(q)$, and let σ be a [complete] coloured trace of p in L. Using symmetry, it suffices to show that σ is also a [complete] coloured trace of q in L. Let ρ be obtained by omitting all actions from the alternating sequence of states and actions σ. Using direction "only if" of clause (i) in the definition of a consistent L^{2}TS, ρ must be a [complete] coloured trace of p in K. As $\mathcal{C}$ is [fully] consistent on K, ρ must also be a [complete] coloured trace of q in K. Finally, using clauses (i) "only if" and (iii), σ must be a [complete] coloured trace of q in L.

Now suppose $s \underline{\leftrightarrow}_b t$ [or $s \underline{\leftrightarrow}_b^\lambda t$]. Then $\mathscr{L}(s) = \mathscr{L}(t)$ and there is a colouring $\mathcal{C}$ on the states of the L^{2}TS, with $\mathcal{C}(s) = \mathcal{C}(t)$, that is [fully] consistent on L. Let $\mathcal{D}$ be the colouring given by $\mathcal{D}(p) := (\mathcal{C}(p), \mathscr{L}(p))$ for all $p \in S$, so that we have $\mathcal{D}(p) = \mathcal{D}(q) \Leftrightarrow [\mathcal{C}(p) = \mathcal{C}(q) \wedge \mathscr{L}(p) = \mathscr{L}(q)]$. It suffices to show that $\mathcal{D}$ is [fully]

consistent on K. The requirement $\mathcal{D}(p) = \mathcal{D}(q) \Rightarrow \mathscr{L}(p) = \mathscr{L}(q)$ is built into the definition of $\mathcal{D}$. So let $\mathcal{D}(p) = \mathcal{D}(q)$, and let ν be a [complete] $\mathcal{D}$-coloured trace of p in K. Using symmetry, it suffices to show that ν is also a [complete] $\mathcal{D}$-coloured trace of q in K. Using clause (i) "only if", there must be a [complete] $\mathcal{D}$-coloured trace ρ of p in L such that ν is obtained from ρ by omitting its actions. Let σ be the [complete] $\mathcal{C}$-coloured trace of s in K obtained from ρ by omitting the second component of each $\mathcal{D}$-colour of a state. As $\mathcal{C}(p) = \mathcal{C}(q)$ and $\mathcal{C}$ is [fully] consistent on L, σ must also be a [complete] $\mathcal{C}$-coloured trace of q in L. By applying clauses (i) "if" and (ii) one derives that ρ is a [complete] $\mathcal{D}$-coloured trace of q in L. Therefore, again using clause (i) "only if", ν must be a [complete] $\mathcal{D}$-coloured trace of q in K.   □

Observation 1. *For every Kripke structure* K *there exists a consistent L^2TS* D *such that* K *is the Kripke structure associated to* D.

One way to obtain D is to label any transition $s \longrightarrow t$ by the pair $(\mathscr{L}(s), \mathscr{L}(t))$ (or simply by $\mathscr{L}(t)$) when $\mathscr{L}(s) \neq \mathscr{L}(t)$, or τ when $\mathscr{L}(s) = \mathscr{L}(t)$. An alternative is the label $(\mathscr{L}(s) - \mathscr{L}(t), \mathscr{L}(t) - \mathscr{L}(s))$, where $(\emptyset, \emptyset)$ is identified with τ.

Unlike the situation for Kripke structures (Observation 1) it is not the case that every LTS can be obtained as the LTS associated to a consistent L^2TS. A simple counterexample is presented in [3]. Thus, in encoding LTSs as L^2TSs, it is in general not possible to keep the set of states the same.

Definition 6. An *LTS-to-L^2TS transformation* η consist of a function taking any LTS L to a consistent L^2TS $\eta(L)$, and in addition taking any state s in L to a state $\eta(s)$ in $\eta(L)$. Such a transformation should at least satisfy $s \underline{\leftrightarrow}_b^\lambda t \Leftrightarrow \eta(s) \underline{\leftrightarrow}_b^\lambda \eta(t)$, that is, it *preserves* ("$\Rightarrow$") and *reflects* ("$\Leftarrow$") divergence sensitive branching bisimulation equivalence, and likewise for $\underline{\leftrightarrow}_b$, and $\underline{\leftrightarrow}_b^\Delta$.

A common LTS-to-L^2TS transformation is presented in [3]. It takes an LTS L $= (S, \rightarrow)$ to an L^2TS $\eta(L)$ by inserting a new state halfway any transition $s \xrightarrow{a} t$ with $a \neq \tau$. This new state is labelled $\{a\}$, and each of the two transitions replacing $s \xrightarrow{a} t$ (from s to the new state and from there to t) is labelled a. Transitions $s \xrightarrow{\tau} t$ are untouched. One takes $\eta(s) = s$ for $s \in S$ and all such states from L are labelled $\{\bot\}$ in $\eta(L)$. (Consult [3] for the formal definition and examples.) In [3] it is shown that this transformation preserves and reflects $\underline{\leftrightarrow}_b^\lambda$; the same proof applies to $\underline{\leftrightarrow}_b$ and $\underline{\leftrightarrow}_b^\Delta$.

An LTS-to-L^2TS transformation η yields an LTS-to-Kripke-structure transformation that we also call η, namely the one transforming an LTS L into the Kripke structure associated to $\eta(L)$. In fact, using Theorem 4 and Observation 1, any LTS-to-Kripke-structure transformation η that preserves and reflects the required equivalences can be obtained in this way.

An LTS-to-L^2TS transformation η makes it possible to define when a state s in an LTS satisfies a $\mathsf{CTL}^*_{-\mathsf{X}}$ formula φ. Namely, one defines $s \models^\eta \varphi$ iff $\eta(s) \models \varphi$. This way, $\mathsf{CTL}^*_{-\mathsf{X}}$ can be used as temporal logic on LTSs.

Theorem 5. *Let s and t be states in an LTS, and let η be an LTS-to-L^2TS transformation. Then*

$$s \underline{\leftrightarrow}_b t \ \textit{iff} \ s \models^\eta_{db} \varphi \Leftrightarrow t \models^\eta_{db} \varphi \ \textit{for all} \ \mathsf{CTL}^*_{-\mathsf{X}} \ \textit{state formulas} \ \varphi$$
$$s \underline{\leftrightarrow}^\lambda_b t \ \textit{iff} \ s \models^\eta \varphi \Leftrightarrow t \models^\eta \varphi \ \textit{for all} \ \mathsf{CTL}^*_{-\mathsf{X}} \ \textit{state formulas} \ \varphi.$$

Proof. This is an immediate consequence of the requirement that η preserves and reflects $\underline{\leftrightarrow}_b$ and $\underline{\leftrightarrow}^\lambda_b$, in combination with Theorems 1, 2 and 4.   □

5   Parallel composition

For a behavioural equivalence to be useful in a process algebraic setting, it is essential that it is a congruence for the operations under consideration. In this section we prove that $\underline{\leftrightarrow}^\Delta_b$ and $\underline{\leftrightarrow}_b$ are congruences for the *merge* or *interleaving operator* $\|$. This operator is often used to represent (asynchronous) parallel composition. However, $\underline{\leftrightarrow}^\lambda_b$ fails to be a congruence for $\|$. We characterise the least discriminating congruence that makes all the distinctions of $\underline{\leftrightarrow}^\lambda_b$ as $\underline{\leftrightarrow}^\Delta_b$.

Definition 7. A binary operation $\|$ on the states of an LTS is a *merge* if for all $s, t, u \in S$ and for all $a \in A$ it holds that $s \| t \xrightarrow{a} u$ iff

  – there exists $s' \in S$ such that $s \xrightarrow{a} s'$ and $u = s' \| t$; or
  – there exists $t' \in S$ such that $t \xrightarrow{a} t'$ and $u = s \| t'$.

Using well known techniques from structural operational semantics [1] any LTS can be augmented with a merge. Moreover, an LTS with merge is generated by the structural operational semantics of any process calculus that includes an operation for pure interleaving. Henceforth we deal with LTSs with a merge $\|$.

Theorem 6. *The relation* $\underline{\leftrightarrow}^\Delta_b$ *is a congruence for* $\|$, *i.e., if* $s \underline{\leftrightarrow}^\Delta_b t$ *and* $u \underline{\leftrightarrow}^\Delta_b v$, *then* $s \| u \underline{\leftrightarrow}^\Delta_b t \| v$.

Proof. Let $\mathcal{R}$ be the reflexive and transitive closure of the relation

$$\{(p\|q, \, p'\|q') \mid p \underline{\leftrightarrow}^\Delta_b p' \ \& \ q \underline{\leftrightarrow}^\Delta_b q'\} \ .$$

Let $\mathcal{C}$ be the function that assigns to each state its equivalence class with respect to $\mathcal{R}$. It suffices to prove that $\mathcal{C}$ is a consistent divergence preserving colouring. So suppose $\mathcal{C}(r) = \mathcal{C}(r')$. Using Lemmas 3 and 5 it suffices to show that r and r' have the same $\mathcal{C}$-coloured traces of length three and the same divergent $\mathcal{C}$-coloured traces of length one. It is straightforward, but notationally cumbersome, to establish this in the special case that $r = p \| q$ and $r' = p' \| q'$ with $p \underline{\leftrightarrow}^\Delta_b p'$ and $q \underline{\leftrightarrow}^\Delta_b q'$. The general case then follows by induction on the length of a chain of pairs from the relation displayed above showing that the pair (r, r') is in its reflexive and transitive closure.   □

A similar proof shows that also $\underline{\leftrightarrow}_b$ is a congruence for $\|$. However, $\underline{\leftrightarrow}^\lambda_b$ is not.

Example 4. Consider an LTS with merge that contains a state 0 without outgoing transitions, a state $\Delta0$ with a τ-loop (an outgoing τ-labelled transition to itself) and no other outgoing transitions, and a state a with $a \xrightarrow{a} 0$ and no other outgoing transitions. (Note that such an LTS is, e.g., generated by the

structural operational semantics of CCS with recursion.) Then $0 \mathrel{\underline{\leftrightarrow}}_b^\lambda \Delta 0$. Figure 4a shows the fragment consisting of the states 0, $\Delta 0$ and a of the LTS under consideration. Figure 4b shows a fragment where the merge is applied. Observe that $0 \parallel a \mathrel{\underline{\not\leftrightarrow}}_b^\lambda \Delta 0 \parallel a$. The reason is that $\Delta 0 \parallel a$ has a maximal path that stays in its initial state, whereas $0 \parallel a$ has not. This problem does not apply to $\mathrel{\underline{\leftrightarrow}}_b$ because $0 \parallel a \mathrel{\underline{\leftrightarrow}}_b \Delta 0 \parallel a$. It does not apply to $\mathrel{\underline{\leftrightarrow}}_b^\Delta$ because $0 \mathrel{\underline{\not\leftrightarrow}}_b^\Delta \Delta 0$.
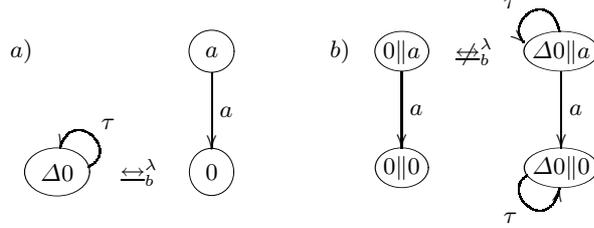
Fig. 4. $\mathrel{\underline{\leftrightarrow}}_b^\lambda$ is not a congruence for parallel composition

The example above involves a deadlock state, namely 0. This is unavoidable, as on LTSs without deadlock $\mathrel{\underline{\leftrightarrow}}_b^\lambda$ coincides with $\mathrel{\underline{\leftrightarrow}}_b^\Delta$ (cf. Section 3) and hence is a congruence for $\parallel$.

The standard solution to the problem of an equivalence $\sim$ failing to be a congruence for a desirable operator Op is to replace it by the coarsest congruence for Op that is included in $\sim$ [10]. Applying this technique to the current situation, the coarsest congruence for $\parallel$ included in $\mathrel{\underline{\leftrightarrow}}_b^\lambda$ turns out to be $\mathrel{\underline{\leftrightarrow}}_b^\Delta$.

Theorem 7. $\mathrel{\underline{\leftrightarrow}}_b^\Delta$ *is the coarsest congruence for* $\parallel$ *that is included in* $\mathrel{\underline{\leftrightarrow}}_b^\lambda$.[2]

Proof. We have already seen that $\mathrel{\underline{\leftrightarrow}}_b^\Delta$ is a congruence for $\parallel$, and that it is included in $\mathrel{\underline{\leftrightarrow}}_b^\lambda$. To show that it is the coarsest, it suffices to show that whenever $s \mathrel{\underline{\not\leftrightarrow}}_b^\Delta t$, we can find a state u such that $s \parallel u \mathrel{\underline{\not\leftrightarrow}}_b^\lambda t \parallel u$. In fact, we can always take u to be the state a from Example 4, for a an action that does not occur in any path from s or t.

Namely, suppose that $s \parallel a \mathrel{\underline{\leftrightarrow}}_b^\lambda t \parallel a$. Let $\mathcal{C}$ be a fully consistent colouring with $\mathcal{C}(s\|a) = \mathcal{C}(t\|a)$. Define the colouring $\mathcal{D}$ by $\mathcal{D}(p) = \mathcal{C}(p\|a)$ for p a state

---

[2] Strictly speaking, we merely show that $\mathrel{\underline{\leftrightarrow}}_b^\Delta$ is the coarsest congruence for $\parallel$ that is included in $\mathrel{\underline{\leftrightarrow}}_b^\lambda$ and satisfies the Fresh Atom Principle (FAP). This principle, described in [7], is satisfied by a semantic equivalence $\sim$ on LTSs when $\sim$ on an LTS L can always be obtained as the restriction of $\sim$ on any given larger LTS of which L is a subLTS, and whose transition labels may be drawn from a larger set of actions than those of L. FAP allows us to use the state a that figures in the proof of Theorem 7, regardless of whether such a state, or the fresh action a, occurs in the given LTS or not. FAP is satisfied by virtually all semantic equivalences documented in the literature, and can be used as a sanity check for meaningful equivalences [7].

reachable from s or t, and $\mathcal{D}(p) = p$ otherwise. Then $\mathcal{D}(s) = \mathcal{D}(t)$. It suffices to show that $\mathcal{D}$ is consistent and preserves divergence, implying $s \underline{\leftrightarrow}_b^\Delta t$.

So suppose $\mathcal{D}(p) = \mathcal{D}(q)$ with $p \neq q$. Then $\mathcal{C}(p\|a) = \mathcal{C}(q\|a)$.

First we show that p and q have the same $\mathcal{D}$-coloured traces. Let σ be a $\mathcal{D}$-coloured trace of p. Then σ is also a $\mathcal{C}$-coloured trace of $p\|a$. As $p\|a$ and $q\|a$ have the same complete $\mathcal{C}$-coloured traces, they surely have the same $\mathcal{C}$-coloured traces (for the coloured traces of a state are the prefixes of its complete coloured traces). Hence σ is a $\mathcal{C}$-coloured trace of $q\|a$. As p is reachable from s or t, the action a cannot occur in σ. Therefore, σ must also be a $\mathcal{D}$-coloured trace of q. By symmetry, any $\mathcal{D}$-coloured trace of q is also a $\mathcal{D}$-coloured trace of p, and hence p and q have the same $\mathcal{D}$-coloured traces.

Next, we show that p and q have the same divergent $\mathcal{D}$-coloured traces. So let σ be a divergent $\mathcal{D}$-coloured trace of p. Then σ is also a divergent $\mathcal{C}$-coloured trace of $p\|a$. Hence σ is a complete $\mathcal{C}$-coloured trace of $p\|a$ and thus also of $q\|a$. As the action a cannot occur in σ, it is not possible that σ stems from a finite maximal path from $q\|a$. Therefore, σ must be a divergent $\mathcal{C}$-coloured trace of $q\|a$, and hence a divergent $\mathcal{D}$-coloured trace of q. Again invoking symmetry, p and q have the same divergent $\mathcal{D}$-coloured traces.

It follows that $\mathcal{D}$ is consistent and preserves divergence; thus $s \underline{\leftrightarrow}_b^\Delta t$. $\qquad\square$

So if one is in search of a semantics such that, for s and t states in an LTS,
- if there is a $\mathsf{CTL}^*_{-\mathsf{X}}$ state formula φ such that $s \models^\eta \varphi$ but $t \not\models^\eta \varphi$, then s and t should be distinguished,
- if s and t can be distinguished after placing them in a context $\_\|u$ for some u, then they should be distinguished to start with
- and no two states should be distinguished unless this is required by the previous two conditions,

then branching bisimulation semantics with explicit divergence is the answer, for $s \underline{\leftrightarrow}_b^\Delta t$ iff for all u and all $\varphi \in \Phi$ we have $s\|u \models^\eta \varphi \Leftrightarrow t\|u \models^\eta \varphi$.

6 Adding deadlock detection to $\mathsf{CTL}^*_{-\mathsf{X}}$

We saw above that there are important properties of states s in an LTS that can be expressed in terms of a context $\_\|u$ and a $\mathsf{CTL}^*_{-\mathsf{X}}$ formula φ, namely as $s\|u \models^\eta \varphi$, but that cannot be directly expressed in terms of $\mathsf{CTL}^*_{-\mathsf{X}}$. This is somewhat unsatisfactory, and therefore we propose an extension of $\mathsf{CTL}^*_{-\mathsf{X}}$ in which this type of property can be expressed directly. We add a path modality ∞ that is valid on a path π iff π is infinite. This path modality, or actually an equally expressive one, was studied prior by Kaivola & Valmari [9] in the context of Linear Temporal Logic without the next state operator—see Section 9.

Definition 8. The syntax of CTL^*_∞ is given by

$$\varphi ::= p \mid \neg\varphi \mid \bigwedge \Phi' \mid \exists\psi \qquad \psi ::= \varphi \mid \neg\psi \mid \bigwedge \Psi' \mid \psi \,\mathsf{U}\, \psi \mid \infty$$

with $p \in \mathbf{AP}$, $\varphi \in \Phi$, $\Phi' \subseteq \Phi$, $\psi \in \Psi$ and $\Psi' \subseteq \Psi$.
Validity is defined as in Definition 2, but adding the clause

$-\ \pi \models \infty$ iff the path π is infinite.

We write $\exists^\infty \psi$ for $\exists(\infty \wedge \psi)$; this formula holds in a state s if there exists an infinite path π from s such that $\pi \models \psi$. Likewise $\forall^\infty \psi = \forall(\infty \rightarrow \psi)$ holds in s if for all infinite paths π from s we have that $s \models \psi$. These constructs are *dual*, in the sense that $s \models \neg\exists^\infty \psi$ iff $s \models \forall^\infty \neg\psi$.

The negation of ∞ holds for a maximal path π iff π is finite, and hence ends in a deadlock. It is tempting to simply extend $\mathsf{CTL}^*_{-\mathsf{X}}$ with a state formula δ such that $s \models \delta$ iff $\neg\exists s'.\ s \longrightarrow s'$. This would make it possible to express ∞ as $\neg\mathsf{F}\delta$. However, this would make the resulting logic too expressive: the two states in the Kripke structure $\circ \longrightarrow \circ$ (with the empty labelling) are branching bisimulation equivalent with explicit divergence, yet they would be distinguished by this extension of $\mathsf{CTL}^*_{-\mathsf{X}}$, as only the last state satisfies δ.

CTL^*_∞ is an extension of $\mathsf{CTL}^*_{-\mathsf{X}}$. There is no need for a similar extension of CTL^*, for δ can be expressed as $\neg\exists\mathsf{X}\top$. In particular, CTL^*_∞ is not more expressive than CTL^*.

The definition of branching bisimulation equivalence with explicit divergence lifts easily to Kripke structures: $s \underset{b}{\leftrightarrow}^\Delta t$, for s and t states in a Kripke structure, iff there exists a consistent and divergence preserving colouring $\mathcal{C}$ such that $\mathcal{C}(s) = \mathcal{C}(t)$. Here *divergence preserving* is defined as in Section 3; by Lemma 5, this time applied to Kripke structures, a consistent colouring preserves divergence iff, for any states s and t, $\mathcal{C}(s) = \mathcal{C}(t)$ implies

$$\text{for any infinite path } \pi \text{ from } s \text{ with } \mathcal{C}(\pi) = \mathcal{C}(s)$$
$$\text{there is an infinite path } \rho \text{ from } t \text{ with } \mathcal{C}(\rho) = \mathcal{C}(t).$$

Theorem 8. $s \underset{b}{\leftrightarrow}^\Delta t$ iff $s \models \varphi \Leftrightarrow t \models \varphi$ for all CTL^*_∞ state formulas φ.

Proof. "Only if" goes as in the proof of Theorem 1, reading $\models$ for $\models_{db}$, requiring $\mathcal{C}$ to be *consistent and divergence preserving*, and, in the second paragraph, requiring the paths π and ρ to be maximal and $\mathcal{C}(\pi)$ to be a *complete* coloured trace of s and t. Here we use that if a colouring is consistent and divergence preserving, then two states with the same colour must also have the same complete coloured traces. This follows from Lemma 6, this time applied to Kripke structures.

There is one extra case to check. Suppose $\mathcal{C}(\pi) = \mathcal{C}(\rho)$ and $\pi \models \infty$, but $\rho \not\models \infty$. Then the last state t of ρ has the same colour $\mathcal{C}(t)$ as one of the states s of π. Let π' be the (infinite) suffix of π starting at s. Then $\mathcal{C}(\pi') = \mathcal{C}(s) = \mathcal{C}(t)$, yet there is no infinite path from t, contradicting that $\mathcal{C}$ is divergence preserving.

"If" goes as in the proof of Theorem 1, but this time we have also to show that $\mathcal{C}$ preserves divergence. So let s and t be states and π an infinite path from s with $\mathcal{C}(\pi) = \mathcal{C}(s) = \mathcal{C}(t) = C$. Let

$$\mathcal{U} = \{u \mid \text{there is a path from } t \text{ to } u \text{ and } \mathcal{C}(u) \neq C\}.$$

For every $u \in \mathcal{U}$ pick a CTL^*_∞ formula $\varphi_u \in C - \mathcal{C}(u)$. Now $s \models \exists^\infty\mathsf{G}(\bigwedge_{u \in \mathcal{U}} \varphi_u)$ and, as $\mathcal{C}(s) = \mathcal{C}(t)$, also $t \models \exists^\infty\mathsf{G}(\bigwedge_{u \in \mathcal{U}} \varphi_u)$. Thus, there is an infinite path ρ from t such that $t' \models \bigwedge_{u \in \mathcal{U}} \varphi_u$ for all states t' in ρ. It follows that $t' \notin \mathcal{U}$. Hence $\mathcal{C}(t') = C$ and thus $\mathcal{C}(\rho) = C$. $\square$

7   Adding deadlock detection to $\mathsf{CTL_{-X}}$

$\mathsf{CTL_{-X}}$ is the sublogic of $\mathsf{CTL^*_{-X}}$ that only allows path formulas of the form $\varphi \,\mathsf{U}\, \varphi'$ and $\neg(\varphi \,\mathsf{U}\, \varphi')$, where φ and φ' are state formulas. Equivalently, it can be defined as only allowing path formulas of the form $\varphi \,\mathsf{U}\, \varphi'$ and $\mathsf{G}\varphi$, for we have

$$s \models \exists\mathsf{G}\varphi \ \text{ iff } \ s \models \exists\neg(\top \,\mathsf{U}\, \neg\varphi)$$

$$s \models \exists\neg(\varphi \,\mathsf{U}\, \varphi') \ \text{ iff } \ s \models \exists[(\neg\varphi') \,\mathsf{U}\, \neg(\varphi \vee \varphi')] \vee \exists\mathsf{G}\neg\varphi' \ .$$

Theorems 1 and 2 are also valid when using $\mathsf{CTL_{-X}}$ instead of $\mathsf{CTL^*_{-X}}$, for their proofs use no other temporal constructs than $\exists(\varphi \,\mathsf{U}\, \varphi')$ and $\exists\mathsf{G}\varphi$.

A natural proposal for $\mathsf{CTL_{\infty}}$ would be to add the path quantifier $\exists^{\infty}$ to $\mathsf{CTL_{-X}}$, thus yielding the syntax

$$\varphi ::= \ p \ \mid\ \neg\varphi \ \mid\ \bigwedge\Phi' \ \mid\ \exists(\varphi \,\mathsf{U}\, \varphi) \mid \exists^{\infty}(\varphi \,\mathsf{U}\, \varphi) \mid \exists\mathsf{G}\varphi \mid \exists^{\infty}\mathsf{G}\varphi \ .$$

However, we can economise on that, for

$$s \models \exists^{\infty}(\varphi \,\mathsf{U}\, \varphi') \ \text{ iff } \ s \models \exists(\varphi \,\mathsf{U}\, (\varphi' \wedge \exists^{\infty}\mathsf{G}\top))$$

$$s \models \exists\mathsf{G}\varphi \ \text{ iff } \ s \models \exists^{\infty}\mathsf{G}\varphi \vee \exists(\varphi \,\mathsf{U}\, (\forall\mathsf{G}\varphi))$$

where $\forall\mathsf{G}\varphi$ is an abbreviation for $\neg\exists(\top \,\mathsf{U}\, \neg\varphi)$. Hence $\mathsf{CTL_{\infty}}$ can be given by the syntax

$$\varphi ::= \ p \ \mid\ \neg\varphi \ \mid\ \bigwedge\Phi' \ \mid\ \exists(\varphi \,\mathsf{U}\, \varphi) \mid \exists^{\infty}\mathsf{G}\varphi \ .$$

It follows immediately from the proof of Theorem 8 that this language is sufficiently expressive to characterise branching bisimulation equivalence with explicit divergence:

Theorem 9. $s \underline{\leftrightarrow}^{\Delta}_b t \ \text{iff}\ s \models \varphi \Leftrightarrow t \models \varphi \ \text{for all } \mathsf{CTL_{\infty}} \text{ formulas } \varphi.$ $\square$

It is tempting to simply write $\exists^{\infty}\mathsf{G}$ as $\exists\mathsf{G}$; that is, to keep the same syntax as for $\mathsf{CTL_{-X}}$ but define its semantics in such a way that $\exists(\varphi \,\mathsf{U}\, \varphi')$ asks merely for a finite path, whereas $\exists\mathsf{G}\varphi$ asks for an infinite one. This *deadlock sensitive* interpretation of $\mathsf{CTL_{-X}}$ is an alternative for the interpretation of [3]. It is consistent with the classical interpretation of CTL [5, 2], as for total Kripke structures there is no difference between $\exists^{\infty}$ and $\exists$.

8   The deadlock extension of Kripke structures

Following De Nicola & Vaandrager [3] we have applied $\mathsf{CTL^*_{-X}}$ to non-total Kripke structures by using maximal instead of infinite paths in the definition of validity. As remarked in Section 2, the same effect can be obtained by transforming a non-total Kripke structure into a total one by adding a self loop $s \longrightarrow s$ to every deadlock state s, and applying the standard $\mathsf{CTL^*_{-X}}$ semantics to the resulting total Kripke structure. However, the latter does not apply to $\mathsf{CTL^*_{\infty}}$, because the self loop $s \longrightarrow s$ invalidates the formula $\exists\neg\infty$ that holds in any deadlock state s. Here we define another transformation on Kripke structures that makes any Kripke structure total, and allows encoding of $\mathsf{CTL^*_{\infty}}$ in terms of $\mathsf{CTL^*_{-X}}$.

Definition 9. The *deadlock extension* $D(\mathrm{K})$ of a Kripke structure K is obtained by the addition of a fresh state s_δ, labelled by the fresh atomic proposition δ, together with a transition from s_δ and from every deadlock state in K to s_δ.

An example of this transformation is depicted in Figure 5.
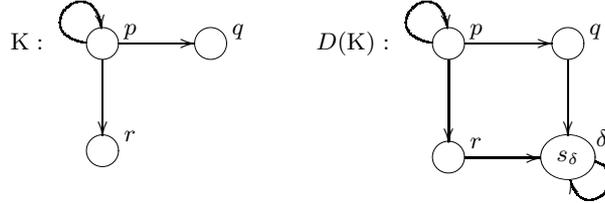
Fig. 5. Deadlock extension of a Kripke structure

Theorem 10. *Let* K *be a Kripke structure, with states* s *and* t. *Then* $s \underline{\leftrightarrow}_b^\Delta t$ *within the Kripke structure* K *iff* $s \underline{\leftrightarrow}_b^\Delta t$ *within the Kripke structure* $D(\mathrm{K})$.

Proof. "If": Let $\mathcal{D}$ be a consistent and divergence preserving colouring on $D(\mathrm{K})$. Note that $\mathcal{D}(s_\delta) \neq \mathcal{D}(s)$ for any state $s \neq s_\delta$ in $D(\mathrm{K})$. Let $\mathcal{C}$ be the restriction of $\mathcal{D}$ to the states of K. Then the $\mathcal{C}$-coloured traces of a state s in K equal the $\mathcal{D}$-coloured traces of s in $D(\mathrm{K})$, but with the colour $\mathcal{D}(s_\delta)$ omitted from the end of such traces. It follows that $\mathcal{C}$ is consistent. It preserves divergence by Lemma 5.

   "Only if": Let $\mathcal{C}$ be a consistent and divergence preserving colouring on K. Extend it to a colouring $\mathcal{D}$ on $D(\mathrm{K})$ by assigning a fresh colour δ to the extra state s_δ of $D(\mathrm{K})$. It suffices to check that $\mathcal{D}$ is consistent and divergence preserving.

   Claim. From any state s in K with the same colour as a deadlock state t in K there must be a path π to a deadlock state such that $\mathcal{C}(\pi) = \mathcal{C}(t)$.

   Proof of claim. As t has no $\mathcal{C}$-coloured traces of length two, neither does s, and as t has no divergent $\mathcal{C}$-coloured traces, neither does s. Thus, all paths from s are finite and only pass through states with colour $\mathcal{C}(t)$.

   Application of the claim. The $\mathcal{D}$-coloured traces of length two of a state $s \neq s_\delta$ in $D(\mathrm{K})$ are the $\mathcal{C}$-coloured traces of length two of the state s in K, together with the trace $\mathcal{C}(t)\delta$ in case s has the same colour as a deadlock state t in K. Thus $\mathcal{D}$ is consistent by Lemma 1, and preserves divergence by Lemma 5.   □

The "if"-direction of the theorem, with a similar proof, also applies to $\approx_s$ and $\approx_{dbs}$, but the "only if"-direction does not. As a counterexample, let K be a Kripke structure with a deadlock state d (having no outgoing transitions) and a livelock state l (with a self-loop as its only one outgoing transition); neither state satisfies any atomic propositions. In K we have $d \approx_s l$, and hence $d \approx_{dbs} l$, but in $D(\mathrm{K})$ we have $d \not\approx_{dbs} l$, and hence $d \not\approx_s l$.

   Considering that Kripke structures of the form $D(\mathrm{K})$ are total, and that on total Kripke structures $\approx_s$ and $\underline{\leftrightarrow}_b^\Delta$ coincide, it is in fact impossible to define a transformation like D for which Theorem 10 holds for both $\underline{\leftrightarrow}_b^\Delta$ and $\approx_s$.

Now let η be an arbitrary LTS-to-L^{2}TS-transformation, yielding an LTS-to-Kripke-structure transformation that is also called η (see Section 4). Then $D \circ \eta$ is not a valid LTS-to-Kripke-structure transformation as intended in [3], for it fails to preserve $\underline{\leftrightarrow}_b^\lambda / \approx_s$ and $\underline{\leftrightarrow}_b / \approx_{dbs}$ (cf. Definition 6). Yet, it satisfies

$$s \underline{\leftrightarrow}_b^\Delta t \quad \Leftrightarrow \quad D \circ \eta(s) \approx_s D \circ \eta(t)$$

(because $s \underline{\leftrightarrow}_b^\Delta t \Leftrightarrow \eta(s) \underline{\leftrightarrow}_b^\Delta \eta(t) \Leftrightarrow D \circ \eta(s) \underline{\leftrightarrow}_b^\Delta D \circ \eta(t)$ and on total Kripke structures $\underline{\leftrightarrow}_b^\Delta$ and $\approx_s$ coincide), and as such it is a suitable transformation for defining validity of $\mathsf{CTL}_{-\mathsf{X}}^*$ formula on states in LTSs. We obtain:

Corollary 2. *Let s and t be states in an LTS, and let η be an LTS-to-L^{2}TS transformation. Then*

$$s \underline{\leftrightarrow}_b^\Delta t \text{ iff } s \models^{D \circ \eta} \varphi \Leftrightarrow t \models^{D \circ \eta} \varphi \text{ for all } \mathsf{CTL}_{-\mathsf{X}}^* \text{ state formulas } \varphi. \qquad \square$$

Thus, one way to make $\mathsf{CTL}_{-\mathsf{X}}^*$ suitable for dealing with deadlock behaviour on LTSs is to stick to total Kripke structures and translate LTSs to Kripke structures by a translation $D \circ \eta$ instead of a transformation η as proposed in [3]. This way branching bisimulation equivalence with explicit divergence becomes the natural counterpart of stuttering equivalence on Kripke structures, and we have the modal characterisation of Corollary 2.

An alternative is to stick to more natural transformations η meeting the criteria on Definition 6, apply the definition of validity of $\mathsf{CTL}_{-\mathsf{X}}^*$ formulas to non-total Kripke structures as in [3], and extend $\mathsf{CTL}_{-\mathsf{X}}^*$ to CTL_∞^* as indicated in Section 6.

Below we show that these solutions lead to equally expressive logics on LTSs.

Definition 10. Given a set of atomic propositions, let CTL_δ^* be the logic $\mathsf{CTL}_{-\mathsf{X}}^*$ extended with an extra atomic proposition δ. The mappings $\mathscr{D}$ from CTL_∞^* to CTL_δ^* formula and $\mathscr{E}$ from CTL_∞^* to CTL_δ^* formula are defined inductively by

$$\begin{aligned}
\mathscr{D}(p) &= p & \mathscr{E}(p) &= p \\
\mathscr{D}(\neg\varphi) &= \neg\delta \wedge \neg\mathscr{D}(\varphi) & \mathscr{E}(\neg\varphi) &= \neg\mathscr{E}(\varphi) \\
\mathscr{D}(\textstyle\bigwedge_{i\in I}\varphi_i) &= \textstyle\bigwedge_{i\in I}\mathscr{D}(\varphi_i) & \mathscr{E}(\textstyle\bigwedge_{i\in I}\varphi_i) &= \textstyle\bigwedge_{i\in I}\mathscr{E}(\varphi_i) \\
\mathscr{D}(\exists\psi) &= \exists\mathscr{D}(\psi) & \mathscr{E}(\exists\psi) &= \exists\mathscr{E}(\psi) \\
\mathscr{D}(\neg\psi) &= \neg\delta \wedge \neg\mathscr{D}(\psi) & \mathscr{E}(\neg\psi) &= \neg\mathscr{E}(\psi) \\
\mathscr{D}(\textstyle\bigwedge_{i\in I}\psi_i) &= \textstyle\bigwedge_{i\in I}\mathscr{D}(\psi_i) & \mathscr{E}(\textstyle\bigwedge_{i\in I}\psi_i) &= \textstyle\bigwedge_{i\in I}\mathscr{E}(\psi_i) \\
\mathscr{D}(\psi \mathbin{\mathsf{U}} \psi') &= \mathscr{D}(\psi) \mathbin{\mathsf{U}} \mathscr{D}(\psi') & \mathscr{E}(\psi \mathbin{\mathsf{U}} \psi') &= (\mathscr{E}(\psi) \mathbin{\mathsf{U}} \delta_{\psi'}) \vee (\mathscr{E}(\psi) \mathbin{\mathsf{U}} \mathscr{E}(\psi')) \\
\mathscr{D}(\infty) &= \neg\mathsf{F}\delta & \mathscr{E}(\delta) &= \neg\top.
\end{aligned}$$

Here $\delta_{\psi'} = \begin{cases} \delta \text{ if } s_\delta \models \exists\psi' \\ \neg\top \text{ otherwise} \end{cases}$, and $\psi \mathbin{\mathsf{U}} \delta$ abbreviates $\neg\infty \wedge \mathsf{G}\psi$.

Theorem 11. *Let K be a Kripke structure and s a state in K.*
Then for any CTL_∞^* state formula φ we have $s \models \varphi$ in K iff $s \models \mathscr{D}(\varphi)$ in $D(\mathrm{K})$, and for any CTL_δ^* state formula φ we have $s \models \varphi$ in $D(\mathrm{K})$ iff $s \models \mathscr{E}(\varphi)$ in K.

Proof. For a state formula φ, let $[\![\varphi]\!]_K$ denote the set of states s in K with $s \models \varphi$. Likewise, for a path formula ψ, $[\![\psi]\!]_K$ denotes the set of maximal paths π in K with $\pi \models \varphi$. Note that there is a bijective correspondence between the maximal paths in K and those in $D(K)$ not starting in s_δ. A straightforward structural induction shows that $[\![\varphi]\!]_K = [\![\mathscr{D}(\varphi)]\!]_{D(K)}$ for any CTL^*_∞ state formula φ and, up to the aforementioned bijective correspondence, $[\![\psi]\!]_K = [\![\mathscr{D}(\psi)]\!]_{D(K)}$ for any CTL^*_∞ path formula ψ.

For the second statement, let π_δ be the unique path in $D(K)$ starting in s_δ. A straightforward structural induction shows that $[\![\varphi]\!]_{D(K)} - \{s_\delta\} = [\![\mathscr{E}(\varphi)]\!]_K$ for any CTL^*_δ state formula φ and, up to the above bijective correspondence, $[\![\psi]\!]_{D(K)} - \{\pi_\delta\} = [\![\mathscr{E}(\psi)]\!]_K$ for any CTL^*_δ path formula ψ. $\qquad\square$

In CTL^*_∞ the path modality ∞ is equally expressive as the path modality $\psi\,\mathsf{U}\,\delta$ of Definition 10, saying of a path that it is finite and all its suffixes satisfy ψ. This is because $\pi \models \psi\,\mathsf{U}\,\delta \Leftrightarrow \pi \models \neg\infty \wedge \mathsf{G}\psi$ and $\pi \models \infty \Leftrightarrow \pi \models \neg\mathsf{F}\delta \Leftrightarrow \pi \models \neg\top\,\mathsf{U}\,\delta$. In this light, the encoding $\mathscr{D}$ of CTL^*_∞ into CTL^*_δ merely adds a conjunct $\neg\delta$ here and there. These conjuncts are not optional; they enable, for instance, the correct translation of the CTL^*_∞ path formula $\mathsf{G}p$ by the CTL^*_δ formula $\neg\delta \wedge \mathsf{G}(\delta \vee p)$.

Theorem 12. *Also the logics CTL_δ and CTL_∞ are equally expressive.*

Proof. This follows because $\mathscr{D}$ can be restricted to a mapping from CTL_∞ to CTL_δ formula and $\mathscr{E}$ to a mapping from CTL_∞ to CTL_δ formula. In particular,

$$\mathscr{D}(\exists(\varphi\,\mathsf{U}\,\varphi')) \;=\; \exists(\mathscr{D}(\varphi)\,\mathsf{U}\,\mathscr{D}(\varphi')) \qquad \mathscr{D}(\exists\mathsf{G}^\infty\varphi) \;=\; \exists\mathsf{G}(\neg\delta \wedge \mathscr{D}(\varphi))$$

$$\mathscr{E}(\exists(\varphi\,\mathsf{U}\,\varphi')) = \begin{cases} \exists(\mathscr{E}(\varphi)\,\mathsf{U}\,\mathscr{E}(\varphi')) \vee \exists(\mathscr{E}(\varphi)\,\mathsf{U}\,(\neg\exists^\infty\mathsf{G}\top \wedge \exists\mathsf{G}\mathscr{E}(\varphi))) & \text{if } s_\delta \models \varphi' \\ \exists(\mathscr{E}(\varphi)\,\mathsf{U}\,\mathscr{E}(\varphi')) & \text{otherwise} \end{cases}$$

and
$$\mathscr{E}(\exists\mathsf{G}\varphi) = \begin{cases} \exists\mathsf{G}^\infty\mathscr{E}(\varphi) & \text{if } s_\delta \models \varphi' \\ \exists\mathsf{G}\,\mathscr{E}(\varphi) & \text{otherwise.} \end{cases} \qquad\square$$

9   Concluding remark: linear temporal logic

Linear Temporal Logic [11] (LTL) is the sublogic of CTL^* that allows propositional variables $p \in \mathbf{AP}$ but no other state formulas to be used as path formulas. Path formulas are applied to states by an implicit universal quantification: $s \models \pi$ iff $s \models \forall\pi$. Kaivola & Valmari [9] study equivalences on LTSs with the property that under all plausible transformations of LTSs into Kripke structures two equivalent states (transformed into states of Kripke structures) satisfy the same formulas in either $\mathsf{LTL}_{-\mathsf{X}}$ (LTL without the next state modality) or LTL^∞ ($\mathsf{LTL}_{-\mathsf{X}}$ augmented with an operator that is equally expressive as our modality ∞). They characterise the coarsest such congruences for a selection of standard process algebra operators—including the merge—as *NDFD*-equivalence (for $\mathsf{LTL}_{-\mathsf{X}}$) and *CFFD*-equivalence (for LTL^∞). There are two striking differences with the branching time case: Whereas $\underline{\leftrightarrow}_b^\Delta$, the equivalence characterising validity in

CTL_∞, is a congruence for the selected process algebra operators directly, in the case of LTL^∞ the resulting equivalence becomes much finer as a result of congruence closure. Furthermore, the addition of ∞ to $\mathsf{LTL}_{-\mathsf{X}}$ leads to distinctions that are not made by CFFD-equivalence. Kaivola & Valmari do not consider the question of raising the expressiveness of $\mathsf{LTL}_{-\mathsf{X}}$ to the level where it characterises NDFD- or CFFD-equivalence directly.

References

1. L. Aceto, W.J. Fokkink & C. Verhoef (2001): *Structural operational semantics.* In J.A. Bergstra, A. Ponse & S.A. Smolka, editors: *Handbook of Process Algebra*, Elsevier, pp. 197-292.
2. M.C. Browne, E.M. Clarke & O. Grumberg (1988): *Characterizing finite Kripke structures in propositional temporal logic.* Theoretical Computer Science 59, pp. 115–131.
3. R. De Nicola & F.W. Vaandrager (1995): *Three logics for branching bisimulation.* Journal of the ACM 42(2), pp. 458–487.
4. E.A. Emerson & E.M. Clarke (1982): *Using branching time temporal logic to synthesize synchronization skeletons.* Science of Computer Programming 2(3), pp. 241–266.
5. E.A. Emerson & J.Y. Halpern (1986): *'Sometimes' and 'Not Never' revisited: on branching time versus linear time temporal logic.* Journal of the ACM 33(1), pp. 151–178.
6. R.J. van Glabbeek (1993): *The linear time - branching time spectrum II.* In E. Best, editor: *Proceedings of CONCUR'93*, LNCS 715, Springer, pp. 66–81.
7. R.J. van Glabbeek (2005): *A characterisation of weak bisimulation congruence.* In A. Middeldorp, V. van Oostrom, F. van Raamsdonk & R. de Vrijer, editors: *Processes, Terms and Cycles: Steps on the Road to Infinity: Essays Dedicated to Jan Willem Klop on the Occasion of His 60th Birthday*, LNCS 3838, Springer, pp. 26–39.
8. R.J. van Glabbeek & W.P. Weijland (1996): *Branching time and abstraction in bisimulation semantics.* Journal of the ACM 43(3), pp. 555–600.
9. R. Kaivola & A. Valmari (1992): *The weakest compositional semantic equivalence preserving Nexttime-less Linear Temporal Logic.* In W.R. Cleaveland, editor: *Proceedings of CONCUR'02*, LNCS 630, pp. 207–221.
10. R. Milner (1989): *Communication and Concurrency.* Prentice Hall, Englewood Cliffs.
11. A. Pnueli (1977): *The Temporal Logic of Programs.* In: *Proceedings of FOCS'77*, IEEE Computer Society Press, pp. 46-57.