

Divergent Quiescent Transition Systems^{*}

Willem G. J. Stokkink, Mark Timmer, and Mariëlle I. A. Stoelinga

Formal Methods and Tools, Faculty of EEMCS
University of Twente, The Netherlands
{w.g.j.stokkink, m.timmer, marielle}@utwente.nl

Abstract. Quiescence is a fundamental concept in modelling system behaviour, as it explicitly represents the fact that no output is produced in certain states. The notion of quiescence is also essential to model-based testing: if a particular implementation under test does not provide any output, then the test evaluation algorithm must decide whether or not to allow this behaviour. To explicitly model quiescence in all its glory, we introduce Divergent Quiescent Transition Systems (DQTSs).

DQTSs model quiescence using explicit δ -labelled transitions, analogous to Suspension Automata (SAs) in the well-known *ioco* framework. Whereas SAs have only been defined implicitly, DQTSs for the first time provide a fully-formalised framework for quiescence. Also, while SAs are restricted to convergent systems (i.e., without τ -cycles), we show how quiescence can be treated naturally using a notion of fairness, allowing systems exhibiting divergence to be modelled as well. We study compositionality under the familiar automata-theoretical operations of determinisation, parallel composition and action hiding. We provide a non-trivial algorithm for detecting divergent states, and discuss its complexity. Finally, we show how to use DQTSs in the context of model-based testing, for the first time presenting a full-fledged theory that allows *ioco* to be applied to divergent systems.

1 Introduction

Quiescence is a fundamental concept in modelling system behaviour. It explicitly represents the fact that in certain states no output is provided. The absence of outputs is often essential: an ATM, for instance, should deliver money only once per transaction. This means that its state just after payment should be quiescent: it should not produce any output until further input is given. On the other hand, the state before payment should clearly not be quiescent. Hence, quiescence may or may not be considered erroneous behaviour. Consequently, the notion of quiescence is essential in model-based testing, where it is detected by means of a timeout. If a particular implementation under test does not provide any output, then the test evaluation algorithm must decide whether to produce a pass verdict (allowing quiescence at this point) or a fail verdict (prohibiting quiescence at this point).

^{*} This research has been partially funded by NWO under grants 612.063.817 (SYRUP), Dn 63-257 (ROCKS) and 12238 (ArRangeer), and by the EU under grant 318490 (SENSATION).

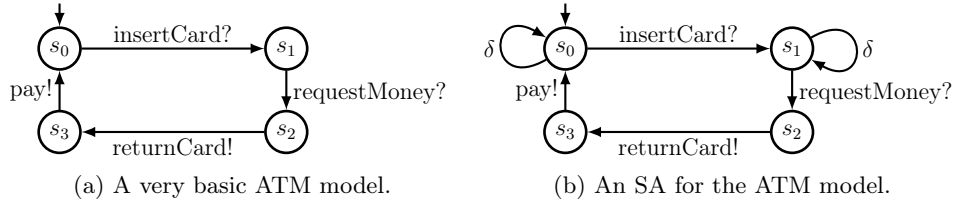


Fig. 1: Deriving a suspension automaton.¹

Origins. The notion of quiescence was first introduced by Vaandrager [1] to obtain a natural extension of blocking states: if a system is input-enabled (i.e., always ready to receive inputs), then no states are blocking, since each state has outgoing input transitions. Quiescence models the fact that a state would be blocking when considering only the internal and output actions. In the context of model-based testing, Tretmans introduced *repetitive quiescence* [2, 3]. This notion emerged from the need to continue testing, even in a quiescent state: in the ATM example above, we may need to test further behaviour arising from the (quiescent) state s_0 . To accommodate this, Tretmans introduced the *Suspension Automaton* (SA) as an auxiliary concept [4]. An SA is obtained from an Input-Output Transition System (IOTS) by first adding a self-loop labelled by the quiescence label δ to each quiescent state and subsequently determining the model. For instance, the ATM automaton in Fig. 1a has quiescent states s_0 and s_1 ; the corresponding SA is depicted in Fig. 1b.

Limitations of current treatments. While previous work [1–4] convincingly argued the need for quiescence, no comprehensive theory of quiescence existed thus far. A severe restriction is that SAs cannot cope with divergence (cycles consisting of internal actions only), since this may introduce newly quiescent states. The TGV framework [5] handles divergence by adding δ -labelled self-loops to such states. However, this treatment is in our opinion not satisfactory: quiescence due to divergence, expressing that no output will ever be produced, can in [5] be followed by an output action, which is counterintuitive. The current paper shows that an appropriate theory for quiescence that can cope with divergence is far from trivial.

Divergence does often occur in practice, e.g., due to action hiding. Therefore, current model-based testing approaches are not able to adequately handle such systems; in this paper, we fill this gap.

Example 1.1. Consider the simplified network protocol shown in Figure 2a. It is obtained as the parallel composition of a sending node (transmitting a message)

¹ Since we require systems to be input-enabled, these models are technically not correct. However, this could easily be fixed by adding self-loops to all states for each missing input. We chose to omit these for clarity of presentation.

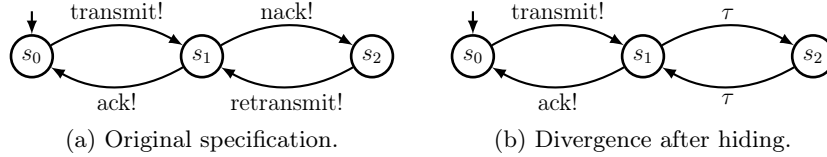


Fig. 2: A simple network protocol.

and a receiving node (sending positive and negative acknowledgements). If only the initial transmission and success of this transmission are considered observable behaviour, the other actions (needed for parallel composition, but irrelevant in the final system) can be hidden, and the system shown in Figure 2b appears. Here, divergence may occur in states s_1 and s_2 (for instance, when retransmission was implemented erroneously and never succeeds). So, observation of quiescence is possible from these states, but simply adding δ -loops does not work anymore. After all, quiescence indicates the *indefinite* absence of outputs, and adding δ -loops to these states would allow outputs to occur after the δ -transitions. Hence, more sophisticated constructs are needed.

In addition to the divergence issue, quiescence was never treated as a first-class citizen: SAs cannot be built from scratch, and, even though important conformance relations such as `ioco` are defined in terms of them, SAs have been defined as an auxiliary construct and have never been studied extensively in isolation. In particular, their closure properties under standard operations like parallel composition and action hiding have not been investigated much.

Our approach. This paper remediates the shortcomings of previous work by introducing *Divergent Quiescent Transition Systems* (DQTSs). DQTSs represent quiescence explicitly using special δ -transitions. We stipulate four well-formedness rules that formalise when δ -transitions may occur. For instance, no δ -transition may be followed by an output transition, since this would contradict the meaning of quiescence. Key in our work is the treatment of divergence: a divergent path leads to the observation of quiescence if and only if it is fair, i.e., models a reasonable execution. We use the notion of fairness from Input-Output Automata (IOAs) [6], based on task partitions.

We show that well-formed DQTSs are closed under parallel composition, determinisation and action hiding. In this way, they constitute a compositional theory for quiescence. Additionally, we formally explain how to obtain a DQTS from an existing IOA by a process called *deltafication*, and show that deltafication is commutative with parallel composition and action hiding. The addition of divergence (and correspondingly fairness) brought about a more involved process of deltafication and action hiding (which may introduce divergence), requiring a novel algorithm for detecting divergent states. We provide this algorithm, which allows us to check well-formedness on a given DQTS as well. Finally, we redefine the `ioco` conformance relation based on DQTSs, allowing it to be applied in the

presence of divergence and hence demonstrating the most important practical benefit of our model for testing: a more general class of systems can be handled.

A preliminary version of this work, already providing a fully formalised framework for dealing with quiescence as a first-class citizen, but not yet supporting divergence, appeared as [7].

Overview of the paper. Sec. 2 introduces the DQTS model, and Sec. 3 presents our well-formedness rules. Sec. 4 then provides operations and properties for DQTSs. In Sec. 5 we describe an algorithm to determine divergent states, and Sec. 6 discusses how to apply DQTSs in the `ioco` framework. Finally, conclusions and future work are presented in Sec. 7. Due to space limitations, we refer to [8] for proofs of all our results.

2 Divergent Quiescent Transition Systems

Preliminaries Given a set L , we use L^* to denote the set of all *finite sequences* $\sigma = a_1 a_2 \dots a_n$ over L . We write $|\sigma| = n$ for the *length* of σ , and ϵ for the *empty sequence*. We let L^ω denote the set of all *infinite sequences* over L , and use $L^\infty = L^* \cup L^\omega$. Given two sequences $\rho \in L^*$ and $v \in L^\infty$, we denote the *concatenation* of ρ and v by ρv . The *projection of an element $a \in L$ on $L' \subseteq L$* , denoted $a \upharpoonright L'$, is a if $a \in L'$ and ϵ otherwise. The projection of a sequence $\sigma = a\sigma'$ is defined inductively by $(a\sigma') \upharpoonright L' = (a \upharpoonright L') \cdot (\sigma' \upharpoonright L')$, and the projection of a set of sequences Z is defined as the sets of projections.

We use $\wp(L)$ to denote the *power set* of L . A set $P \subseteq \wp(L)$ such that $\emptyset \notin P$ is a *partition* of L if $\bigcup P = L$ and $p \neq q$ implies $p \cap q = \emptyset$ for all $p, q \in P$. Finally, we use the notation \exists^∞ for ‘there exist infinitely many’.

2.1 Basic Model and Definitions

Divergent Quiescent Transition Systems (DQTSs) are labelled transition systems that model quiescence, i.e., the absence of outputs or internal transitions, via a special δ -action. They are based on the well-known Input-Output Automata [9, 6]; in particular, their task partitions allow one to define fair paths.

Definition 2.1 (Divergent Quiescent Transition System). A Divergent Quiescent Transition System (DQTS) is a tuple $\mathcal{A} = \langle S, S^0, L^I, L^O, L^H, P, \rightarrow \rangle$, where S is a set of states; $S^0 \subseteq S$ is a non-empty set of initial states; L^I, L^O and L^H are disjoint sets of input, output and internal labels, respectively; P is a partition of $L^O \cup L^H$; and $\rightarrow \subseteq S \times L \cup \{\delta\} \times S$ is the transition relation, where $L = L^I \cup L^O \cup L^H$. We assume $\delta \notin L$.

Given a DQTS \mathcal{A} , we denote its components by $S_{\mathcal{A}}, S_{\mathcal{A}}^0, L_{\mathcal{A}}^I, L_{\mathcal{A}}^O, L_{\mathcal{A}}^H, P_{\mathcal{A}}, \rightarrow_{\mathcal{A}}$. We omit the subscript when it is clear from the context.

Example 2.1. The SA in Fig. 1b is a DQTS. □

Restrictions. We impose two important restrictions on DQTSs. (1) We require each DQTS \mathcal{A} to be *input-enabled*, i.e., always ready to accept any input. Thus, we require that for each $s \in S$ and $a \in L^I$, there exists an $s' \in S$ such that $(s, a, s') \in \rightarrow$. (2) We require each DQTS to be well-formed. Well-formedness requires technical preparation and is defined in Sec. 3.

Semantically, DQTSs assume progress. That is, DQTSs are not allowed to remain idle forever when output or internal actions are enabled. Without this assumption, each state would be potentially quiescent.

Actions. We use the terms label and action interchangeably. We often suffix a question mark (?) to input labels and an exclamation mark (!) to output labels. These are, however, not part of the label. A label without a suffix denotes an internal label. Output and internal actions are called *locally controlled*, because their occurrence is under the control of the DQTS. Thus, $L^{LC} = L^O \cup L^H$ denotes the set of all locally controlled actions. The special label δ is used to denote the occurrence of quiescence (see Def. 2.10). The task partition P partitions the locally controlled actions into blocks, allowing one to reason about fairness: an execution is fair if every task partition that is enabled infinitely often, is also given control infinitely often (see Sec. 2.2).

We use the standard notations for transitions.

Definition 2.2 (Transitional notations). *Let \mathcal{A} be a DQTS with $s, s' \in S$, $a, a_i \in L$, $b, b_i \in L^I \cup L^O$, and $\sigma \in (L^I \cup L^O)^+$, then:*

$$\begin{aligned}
s \xrightarrow{a} s' &=_{\text{def}} (s, a, s') \in \rightarrow \\
s \xrightarrow{a} &=_{\text{def}} \exists s'' \in S . s \xrightarrow{a} s'' \\
s \not\xrightarrow{a} &=_{\text{def}} \nexists s'' \in S . s \xrightarrow{a} s'' \\
s \xrightarrow{a_1 \dots a_n} s' &=_{\text{def}} \exists s_0, \dots, s_n \in S . s = s_0 \xrightarrow{a_1} \dots \xrightarrow{a_n} s_n = s' \\
s \xrightarrow{\epsilon} s' &=_{\text{def}} s = s' \text{ or } \exists a_1, \dots, a_n \in L^H . s \xrightarrow{a_1 \dots a_n} s' \\
s \xrightarrow{b} s' &=_{\text{def}} \exists s_0, s_1 \in S . s \xrightarrow{\epsilon} s_0 \xrightarrow{b} s_1 \xrightarrow{\epsilon} s' \\
s \xrightarrow{b_1 \dots b_n} s' &=_{\text{def}} \exists s_0, \dots, s_n \in S . s = s_0 \xrightarrow{b_1} \dots \xrightarrow{b_n} s_n = s' \\
s \xrightarrow{\sigma} &=_{\text{def}} \exists s'' \in S . s \xrightarrow{\sigma} s''
\end{aligned}$$

If $s \xrightarrow{a}$, we say that a is enabled in s . We use $L(s)$ to denote the set of all actions $a \in L$ that are enabled in state $s \in S$, i.e., $L(s) = \{a \in L \mid s \xrightarrow{a}\}$. The notions are lifted to infinite traces in the obvious way.

We use the following language notations for DQTSs and their behaviour.

Definition 2.3 (Language notations). *Let \mathcal{A} be a DQTS, then:*

- A finite path in \mathcal{A} is a sequence $\pi = s_0 a_1 s_1 a_2 s_2 \dots s_n$ such that $s_{i-1} \xrightarrow{a_i} s_i$ for all $1 \leq i \leq n$. Infinite paths are defined analogously. The set of all paths in \mathcal{A} is denoted $\text{paths}(\mathcal{A})$.
- Given any path, we write $\text{first}(\pi) = s_0$. Also, we denote by $\text{states}(\pi)$ the set of states that occur on π , and by $\omega\text{-states}(\pi)$ the set of states that occur infinitely often. That is, $\omega\text{-states}(\pi) = \{s \in \text{states}(\pi) \mid \exists^\infty j . s_j = s\}$.

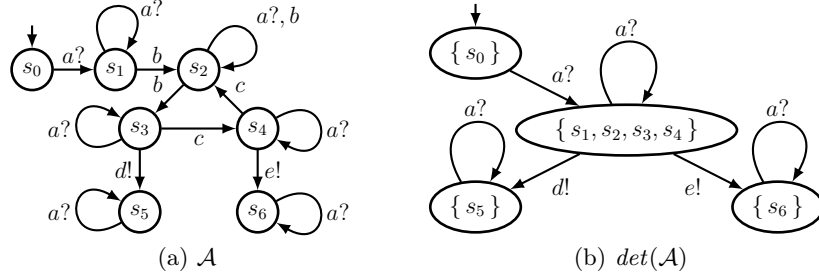


Fig. 3: Visual representations of the DQTSs \mathcal{A} and $\det(\mathcal{A})$.

- We define $\text{trace}(\pi) = \pi \upharpoonright (L^I \cup L^O)$, and say that $\text{trace}(\pi)$ is the trace of π . For every $s \in S$, $\text{traces}(s)$ is the set of all traces corresponding to paths that start in s , i.e., $\text{traces}(s) = \{ \text{trace}(\pi) \mid \pi \in \text{paths}(\mathcal{A}) \wedge \text{first}(\pi) = s \}$. We define $\text{traces}(\mathcal{A}) = \bigcup_{s \in S^0} \text{traces}(s)$, and say that two DQTSs \mathcal{B} and \mathcal{C} are trace-equivalent, denoted $\mathcal{B} \approx_{\text{tr}} \mathcal{C}$, if $\text{traces}(\mathcal{B}) = \text{traces}(\mathcal{C})$.
- For a finite trace σ and state $s \in S$, $\text{reach}(s, \sigma)$ denotes the set of states in \mathcal{A} that can be reached from s via σ , i.e., $\text{reach}(s, \sigma) = \{ s' \in S \mid s \xrightarrow{\sigma} s' \}$. For a set of states $S' \subseteq S$, we define $\text{reach}(S', \sigma) = \bigcup_{s \in S'} \text{reach}(s, \sigma)$.

When needed, we add subscripts to indicate the DQTS these notions refer to.

Definition 2.4 (Determinism). A DQTS \mathcal{A} is deterministic if $s \xrightarrow{a} s'$ and $s \xrightarrow{a} s''$ imply $a \notin L^H$ and $s' = s''$, for all $s, s', s'' \in S$ and $a \in L$. Otherwise, \mathcal{A} is nondeterministic.

Each DQTS has a trace-equivalent deterministic DQTS [10, 11]. Determinisation is carried out using the well-known subset construction procedure. This construction yields a system in which every state has a unique target per action, and internal transitions are not present anymore.

Definition 2.5 (Determinisation). The determinisation of a DQTS $\mathcal{A} = \langle S, S^0, L^I, L^O, L^H, P, \rightarrow \rangle$ is the DQTS $\det(\mathcal{A}) = \langle T, \{ S^0 \}, L^I, L^O, L^H, P, \rightarrow_D \rangle$, with $T = \wp(S) \setminus \emptyset$ and $\rightarrow_D = \{ (U, a, V) \in T \times L \times T \mid V = \text{reach}_{\mathcal{A}}(U, a) \wedge V \neq \emptyset \}$.

Example 2.2. The DQTS \mathcal{A} in Fig. 3a is nondeterministic; its determinisation $\det(\mathcal{A})$ is shown in Fig. 3b. \square

2.2 Fairness and Divergence

The notion of fairness also plays a crucial role in DQTSs. The reason for this is that parallel composition may yield unreasonable divergences. For instance, if the DQTS in Fig. 4 is the composition of a system consisting solely of an internal a -loop and a system outputting a b precisely once, the progress assumption on the

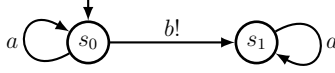


Fig. 4: Visual representation of a DQTS \mathcal{B} .

second component tells us that at some point we should observe this b -output. Therefore, we want to prohibit the divergent path $\pi = s_0 a s_0 a s_0 \dots$.

The following definition stems from [9, 6, 12], and states that if a subcomponent of the system infinitely often wants to execute some of its actions, it will indeed infinitely often execute some. Note that finite paths are fair by default.

Definition 2.6 (Fair path). *Let \mathcal{A} be a DQTS and $\pi = s_0 a_1 s_1 a_2 s_2 \dots$ a path of \mathcal{A} . Then, π is fair if, for every $A \in P$ such that $\exists^\infty j . L(s_j) \cap A \neq \emptyset$, we have $\exists^\infty j . a_j \in A$. The set of all fair paths of a DQTS \mathcal{A} is denoted $fpaths(\mathcal{A})$, and the set of corresponding traces is denoted $ftraces(\mathcal{A})$.*

Unfair paths are considered not to occur, so from now on we only consider $fpaths(\mathcal{A})$ and $ftraces(\mathcal{A})$ for the behaviour of \mathcal{A} .

Example 2.3. Consider again the DQTS \mathcal{B} in Fig. 4. The infinite path $\pi = s_0 a s_0 a s_0 \dots$ would not be fair if $P_{\mathcal{B}} = \{\{a\}, \{b\}\}$, as the b -output is ignored forever. It would however be fair if $P_{\mathcal{B}} = \{\{a, b\}\}$. \square

We can now formally define divergence: fair infinite internal behaviour.

Definition 2.7 (Divergent path). *Let \mathcal{A} be a DQTS and $\pi \in fpaths(\mathcal{A})$ a fair infinite path. The path π is divergent if it contains only transitions labelled with internal actions, i.e., $a_i \in L_{\mathcal{A}}^H$ for every action a_i on π . The set of all divergent paths of \mathcal{A} is denoted $dpaths(\mathcal{A})$.*

Example 2.4. Consider the DQTS \mathcal{A} in Fig. 3a with $L_{\mathcal{A}}^H = \{b, c\}$. The infinite paths $s_2 b s_2 b s_2 \dots$ and $s_2 b s_3 c s_4 c s_2 b s_3 \dots$ are both divergent. Note that divergent traces are not preserved by determinisation. \square

In contrast to SAs, we do allow divergent paths to occur in DQTSs. However, we assume that each divergent path in a DQTS only contains a finite number of states. This restriction serves to ensure that the deltafication of a DQTS, discussed in Sec. 4.1, always results in a correct DQTS. Since DQTSs typically contain a finite number of states, and even in infinite systems divergence often results from internal loops, this restriction is not a severe one.

Definition 2.8 (State-finite path). *Let \mathcal{A} be a DQTS and let $\pi \in fpaths(\mathcal{A})$ be an infinite path. If $|states(\pi)| < \infty$, then π is state-finite.*

When the system is on a state-finite divergent path, it continuously loops through a finite number of states on this path. We call these states divergent.

Definition 2.9 (Divergent state). *Let \mathcal{A} be a DQTS. A state $s \in S$ is divergent, denoted $d(s)$, if there is a (state-finite and fair) divergent path on which s occurs infinitely often, i.e., if there is a path $\pi \in dpaths(\mathcal{A})$ such that $s \in \omega\text{-states}(\pi)$. The set of all divergent states of \mathcal{A} is denoted $d(\mathcal{A})$.*

Example 2.5. Consider the DQTS \mathcal{A} in Fig. 3a. The path $\pi_1 = s_1 b s_2 b s_2 \dots$ is state-finite, fair and divergent. Since s_2 occurs infinitely often on π_1 , it is divergent; s_1 , on the other hand, is not. Whether s_3 is divergent depends on the task partition P . If P contains an element A such that $\{c, d, e\} \subseteq A$, then $\pi_2 = s_3 c s_4 c s_2 b s_3 \dots$ is fair and s_3 is divergent; otherwise, it is not. \square

2.3 Quiescence

Definition 2.10 (Quiescent state). *Let \mathcal{A} be a DQTS. A state $s \in S$ is quiescent, denoted $q(s)$, if it has no locally-controlled actions enabled. That is, $q(s)$ if $s \not\rightarrow$ for all $a \in L^{LC}$. The set of all quiescent states of \mathcal{A} is denoted $q(\mathcal{A})$.*

Example 2.6. States s_0 , s_5 and s_6 of the DQTS \mathcal{A} in Fig. 3a are quiescent. \square

Divergent paths in DQTSs may yield observations of quiescence in states that are not necessarily quiescent. Consider the DQTS \mathcal{B} in Fig. 4. State s_0 is not quiescent, since it enables output b . Nevertheless, this output is never observed on the divergent path $\pi = s_0 a s_0 a \dots$. Hence, quiescence might be observed in a non-quiescent state (here, if π is fair). After observing quiescence due to a divergent path, the system will reside in one of the divergent states on that path.

3 Well-formed DQTSs

To be meaningful, DQTSs have to adhere to four well-formedness rules that formalize the semantics of quiescence. As indicated before, we assume all DQTSs to do so.

Definition 3.1 (Well-formedness). *A DQTS \mathcal{A} is well-formed if it satisfies the following rules for all $s, s', s'' \in S$ and $a \in L^I$:*

Rule R1 (Quiescence should be observable): if $q(s)$ or $d(s)$, then $s \xrightarrow{\delta}$.

This rule requires that each quiescent or divergent state has an outgoing δ -transition, since in these states quiescence may be observed.

Rule R2 (Quiescent state after quiescence observation): if $s \xrightarrow{\delta} s'$, then $q(s')$.

Since there is no notion of timing in DQTSs, there is no particular observation duration associated with quiescence. Hence, the execution of a δ -transition represents that the system has not produced any outputs indefinitely; therefore, enabling any outputs after a δ -transition would clearly be erroneous.

Note that, even though the δ -transition may be due to divergence, it would not suffice to require $q(s') \vee d(s')$. After all, $d(s')$ does not exclude output actions from s' , and these should not be enabled directly after a δ -transitions.

Rule R3 (No new behaviour after quiescence observation): if $s \xrightarrow{\delta} s'$, then $\text{traces}(s') \subseteq \text{traces}(s)$.

There is no notion of timing in DQTSs. Hence, behaviour that is possible after an observation of quiescence, must also be possible beforehand. Still, the observation of quiescence may indicate the outcome of an earlier nondeterministic choice, thereby reducing possible behaviour. Hence, the potential inequality.

Rule R4 (Continued quiescence preserves behaviour): if $s \xrightarrow{\delta} s'$ and $s' \xrightarrow{\delta} s''$, then $\text{traces}(s'') = \text{traces}(s')$.

Since quiescence represents the fact that no outputs are observed, and there is no notion of timing in the DQTS model, there can be no difference between observing quiescence once or multiple times in succession.

In [13], four similar, but more complex, rules for *valid* SAs are discussed. However, these did not account for divergence.

Note that, by definition of divergent states, rule R1 does not require δ -transitions from states that have outgoing divergent paths on which they occur only finitely often. This simplifies the deltafication procedure, as will be made clear in Example 4.1. Also note that a path of a DQTS may contain multiple successive δ -transitions. This corresponds to the practical testing scenario of observing a time-out rather than an output more than once in a row [2, 3].

Since SAs are derived from IOTSs, and we assume that these IOTSs correctly capture system behaviour, we find that SAs are ‘well-formed’ in the sense that their observable behaviour (including quiescence) corresponds to that of realistic specifications. Since we also desire this property to hold for well-formed DQTSs, the above rules have been carefully crafted in such a way that well-formed DQTSs and SAs are equivalent in terms of expressible observable behaviour. The following theorem characterises this core motivation behind our design decisions: it shows that every trace in a DQTS can be obtained by starting with a traditional IOTS and adding δ -loops as for SAs, and vice versa. Hence, except for divergences, their expressivity coincides.

Theorem 3.1. *For every SA \mathcal{S} there exists a well-formed DQTS \mathcal{D} such that $\mathcal{S} \approx_{\text{tr}} \mathcal{D}$, and vice versa.*

Verifying rule R1 requires identifying divergent states; in Sec. 5 we provide an algorithm to do so. Rule R2 can be checked trivially, while R3 and R4 in practice could be checked heuristically. For R3, verify whether $s \xrightarrow{\delta} s'$ and $s' \xrightarrow{a?} s''$ imply $s \xrightarrow{a?} s''$, and for R4, verify whether $s \xrightarrow{\delta} s'$ and $s' \xrightarrow{\delta} s''$ imply that $s' = s''$.

4 Operations and Properties

4.1 Deltafication: from IOA to DQTS

Usually, specifications are modelled as IOAs (or IOTSs, which can easily be converted to IOAs by taking $L^H = \{\tau\}$ and $P = \{L^{LC}\}$). During testing,

however, we typically observe the outputs of the system generated in response to inputs from the tester; thus, it is useful to be able to refer to the absence of outputs explicitly. Hence, we need a way to convert an IOA to a well-formed DQTS that captures all possible observations of it, including quiescence. This conversion is called *deltafication*. It uses the notions of quiescence, divergence and state-finiteness, which were defined for DQTSs, but can just as well be used for IOAs (interpreting them as non-well-formed DQTSs without any δ -transitions). As for DQTSs, we require all IOAs to be input-enabled.

To satisfy rule R1, every state in which quiescence may be observed must have an outgoing δ -transition. When constructing SAs, δ -labelled self-loops are added to all quiescent states. This would not work for divergent states, however, since divergent states have outgoing internal transitions and possibly even output transitions (as in Fig. 4). So, a δ -labelled self-loop would contradict rule R2.

Our solution is to introduce a new state qos_s for every divergent state s , which acts as its *quiescence observation state*. When quiescence is observed in s , a δ -transition will lead to qos_s . To preserve the original behaviour, all inputs that are enabled in s must still be enabled in qos_s , and must lead to the same states that the original input transitions led to. All these considerations together lead to the following definition for the deltaxification procedure for IOAs.

Definition 4.1 (Deltaxification). *Let $\mathcal{A} = \langle S_{\mathcal{A}}, S^0, L^I, L^O, L^H, P, \rightarrow_{\mathcal{A}} \rangle$ be an IOA with $\delta \notin L$. The deltaxification of \mathcal{A} is $\delta(\mathcal{A}) = \langle S_{\delta}, S^0, L^I, L^O, L^H, P, \rightarrow_{\delta} \rangle$. We define $S_{\delta} = S_{\mathcal{A}} \cup \{ qos_s \mid s \in d(\mathcal{A}) \}$, i.e., S_{δ} contains a new state $qos_s \notin S_{\mathcal{A}}$ for every divergent state $s \in S_{\mathcal{A}}$ of \mathcal{A} . The transition relation \rightarrow_{δ} is as follows:*

$$\begin{aligned} \rightarrow_{\delta} = \rightarrow_{\mathcal{A}} \cup & \{ (s, \delta, s) \mid s \in q(\mathcal{A}) \} \\ & \cup \{ (s, \delta, qos_s) \mid s \in d(\mathcal{A}) \} \cup \{ (qos_s, \delta, qos_s) \mid s \in d(\mathcal{A}) \} \\ & \cup \{ (qos_s, a?, s') \mid s \in d(\mathcal{A}) \wedge a? \in L^I \wedge s \xrightarrow{a?}_{\mathcal{A}} s' \} \end{aligned}$$

Thus, the deltaxification of an IOA adds δ -labelled self-loops to all quiescent states. Furthermore, a new quiescence observation state qos_s is introduced for every divergent state $s \in S$, alongside the required inputs and δ -transitions.

Note that computing $q(\mathcal{A})$ is trivial: simply identify all states without outgoing output or internal transition. Determining $d(\mathcal{A})$ is more complex; an algorithm to do so is provided in Sec. 5.

Example 4.1. See Fig. 5 for IOA \mathcal{A} and its deltaxification, given $P_{\mathcal{A}} = \{ \{ b, c \} \}$. Hence, s_1 and s_2 are divergent, and q_0 and q_1 quiescence observation states. Note that s_0 has an outgoing divergent path, while in accordance to rule R1 it is not given an outgoing δ -transition. The reason is that, when observing quiescence, our progress assumption prescribes that the system can only reside in s_1 or s_2 . Hence, quiescence cannot be observed from s_0 , and therefore also the a -transition to s_3 should not be possible anymore after observation of quiescence. This is now taken care of by not having a direct δ -transition from s_0 . Because of this, no trace first having δ and then having the $b!$ output is present. \square

As expected, deltaxification indeed yields a well-formed DQTS.

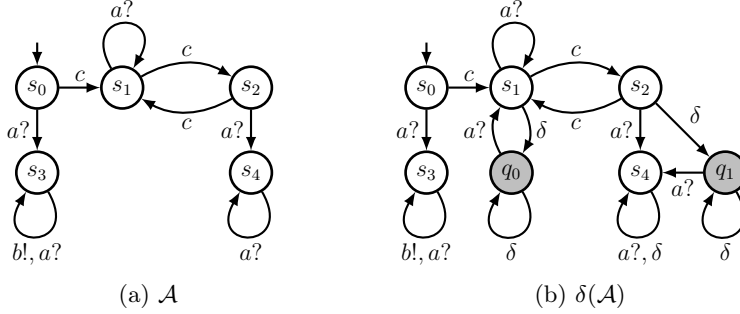


Fig. 5: An IOA \mathcal{A} and its deltafication $\delta(\mathcal{A})$. Newly introduced states are grey.

Theorem 4.1. *Given an IOA \mathcal{A} with $\delta \notin L$ such that all divergent paths in \mathcal{A} are state-finite, $\delta(\mathcal{A})$ is a well-formed DQTS.*

4.2 Operations on DQTSs

We introduce several standard operations on well-formed DQTSs. First, we define the well-known parallel composition operator. As usual, it requires every locally controlled action to be under the control of at most one component [6].

Definition 4.2 (Compatibility). *Two DQTSs \mathcal{A} and \mathcal{B} are compatible if $L_{\mathcal{A}}^O \cap L_{\mathcal{B}}^O = \emptyset$, $L_{\mathcal{A}}^H \cap L_{\mathcal{B}} = \emptyset$, and $L_{\mathcal{B}}^H \cap L_{\mathcal{A}} = \emptyset$.*

Definition 4.3 (Parallel composition). *Given two well-formed compatible DQTSs \mathcal{A} and \mathcal{B} , the parallel composition of \mathcal{A} and \mathcal{B} is the DQTS $\mathcal{A} \parallel \mathcal{B}$, with $S_{\mathcal{A} \parallel \mathcal{B}} = S_{\mathcal{A}} \times S_{\mathcal{B}}$, $S_{\mathcal{A} \parallel \mathcal{B}}^0 = S_{\mathcal{A}}^0 \times S_{\mathcal{B}}^0$, $L_{\mathcal{A} \parallel \mathcal{B}}^I = (L_{\mathcal{A}}^I \cup L_{\mathcal{B}}^I) \setminus (L_{\mathcal{A}}^O \cup L_{\mathcal{B}}^O)$, $L_{\mathcal{A} \parallel \mathcal{B}}^O = L_{\mathcal{A}}^O \cup L_{\mathcal{B}}^O$, $L_{\mathcal{A} \parallel \mathcal{B}}^H = L_{\mathcal{A}}^H \cup L_{\mathcal{B}}^H$, $P_{\mathcal{A} \parallel \mathcal{B}} = P_{\mathcal{A}} \cup P_{\mathcal{B}}$, and*

$$\begin{aligned} \rightarrow_{\mathcal{A} \parallel \mathcal{B}} = & \{ ((s, t), a, (s', t')) \in S_{\mathcal{A} \parallel \mathcal{B}} \times ((L_{\mathcal{A}} \cap L_{\mathcal{B}}) \cup \{ \delta \}) \times S_{\mathcal{A} \parallel \mathcal{B}} \mid \\ & s \xrightarrow{\mathcal{A}} s' \wedge t \xrightarrow{\mathcal{B}} t' \} \\ & \cup \{ ((s, t), a, (s', t)) \in S_{\mathcal{A} \parallel \mathcal{B}} \times (L_{\mathcal{A}} \setminus L_{\mathcal{B}}) \times S_{\mathcal{A} \parallel \mathcal{B}} \mid s \xrightarrow{\mathcal{A}} s' \} \\ & \cup \{ ((s, t), a, (s, t')) \in S_{\mathcal{A} \parallel \mathcal{B}} \times (L_{\mathcal{B}} \setminus L_{\mathcal{A}}) \times S_{\mathcal{A} \parallel \mathcal{B}} \mid t \xrightarrow{\mathcal{B}} t' \} \end{aligned}$$

We have $L_{\mathcal{A} \parallel \mathcal{B}} = L_{\mathcal{A} \parallel \mathcal{B}}^I \cup L_{\mathcal{A} \parallel \mathcal{B}}^O \cup L_{\mathcal{A} \parallel \mathcal{B}}^H = L_{\mathcal{A}} \cup L_{\mathcal{B}}$.

Note that we require DQTSs to synchronise on δ -transitions, as a parallel composition of two DQTSs can only be quiescent when both components are.

It is often useful to hide certain output actions of a given well-formed DQTS, treating them as internal actions. For example, actions used for synchronisation are often not needed anymore in the parallel composition. Action hiding is slightly more complicated for DQTSs than for IOAs, as transforming output actions to internal actions can lead to newly divergent states. Still, whereas in SAs this was forbidden, in DQTSs it is allowed. Consequently, after hiding, new quiescence observation states may have to be added for newly divergent states.

Definition 4.4 (Action hiding). Let $\mathcal{A} = \langle S_{\mathcal{A}}, S^0, L^I, L^O, L^H, P, \rightarrow_{\mathcal{A}} \rangle$ be a well-formed DQTS and $H \subseteq L^O$ a set of outputs, then hiding H in \mathcal{A} yields the DQTS $\mathcal{A} \setminus H = \langle S_H, S^0, L^I, L_H^O, L_H^H, P, \rightarrow_H \rangle$, with $L_H^O = L^O \setminus H$, $L_H^H = L^H \cup H$, and $S_H = S_{\mathcal{A}} \cup \{ qos_s \mid s \in d(\mathcal{A} \setminus H) \setminus d(\mathcal{A}) \}$. Finally, \rightarrow_H is defined by

$$\begin{aligned} \rightarrow_H = \rightarrow_{\mathcal{A}} \cup & \{ (s, \delta, qos_s) \mid s \in d(\mathcal{A} \setminus H) \setminus d(\mathcal{A}) \} \\ & \cup \{ (qos_s, \delta, qos_s) \mid s \in d(\mathcal{A} \setminus H) \setminus d(\mathcal{A}) \} \\ & \cup \{ (qos_s, a?, s') \mid s \in d(\mathcal{A} \setminus H) \setminus d(\mathcal{A}) \wedge a? \in L^I \wedge s \xrightarrow{\mathcal{A}} s' \} \end{aligned}$$

So, similar to deltafication, quiescence observation states are added for all newly divergent states, along with the required input transitions to preserve behaviour.

4.3 Properties of DQTSs

We present several important results regarding DQTSs. First, it turns out that well-formed DQTSs are closed under all operations defined thus far.

Theorem 4.2. *Well-formed DQTSs are closed under the operations of determinisation, parallel composition, and action hiding, i.e., given two well-formed and compatible DQTSs \mathcal{A} and \mathcal{B} , and a set of labels $H \subseteq L_{\mathcal{A}}^O$, we find that $det(\mathcal{A})$, $\mathcal{A} \setminus H$, and $\mathcal{A} \parallel \mathcal{B}$ are also well-formed DQTSs.*

Next, we investigate the commutativity of function composition of deltafication with the operations. We consider the function compositions of two operations to be commutative if the end results of applying both operations in either order are trace equivalent. After all, trace-equivalent DQTSs behave in the same way. (Note that this is not the case for IOAs or IOTSS, as trace-equivalent variants of such systems might have different quiescence behaviour.) We show that parallel composition and action hiding can safely be swapped with deltafication, but note that deltafication has to precede determinisation to get sensible results. This is immediate, since determinisation does not preserve quiescence.

Proposition 4.1. *Deltafication and determinisation do not commute, i.e., given an IOA \mathcal{A} such that $\delta \notin L$, not necessarily $det(\delta(\mathcal{A})) \approx_{tr} \delta(det(\mathcal{A}))$.*

Consequently, when transforming a nondeterministic IOA \mathcal{A} to a deterministic, well-formed DQTS, one should first derive $\delta(\mathcal{A})$ and afterwards determinise.

Deltafication does commute with action hiding and parallel composition. In the following theorem we use \setminus_I to denote basic action hiding for IOAs, and \setminus_D to denote action hiding for DQTSs (conform Def. 4.4).

Theorem 4.3. *Deltafication and action hiding commute: given an IOA \mathcal{A} such that $\delta \notin L$ and a set of labels $H \subseteq L_{\mathcal{A}}^O$, we have $\delta(\mathcal{A} \setminus_I H) \approx_{tr} \delta(\mathcal{A}) \setminus_D H$.*

Theorem 4.4. *Deltafication and parallel composition commute: given two compatible IOAs \mathcal{A} , \mathcal{B} , such that $\delta \notin L_{\mathcal{A}} \cup L_{\mathcal{B}}$, we have $\delta(\mathcal{A} \parallel \mathcal{B}) \approx_{tr} \delta(\mathcal{A}) \parallel \delta(\mathcal{B})$.*

The above results allow great modelling flexibility. After all, hiding and parallel composition are often already applied to the IOAs that describe a specification. We now showed that after deltafication these then yield the same well-formed DQTSs as in the case these operations are applied after deltafication.

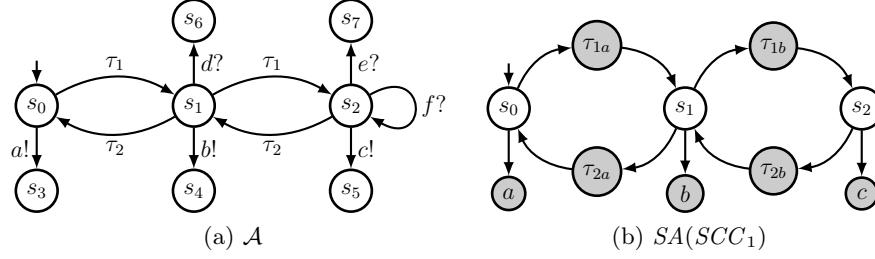


Fig. 6: An IOA \mathcal{A} and the Streett automaton $SA(SCC_1)$.

5 Algorithm for detecting divergent states

We present an algorithm to detect divergent states in an IOA or DQTS. This is vital for verifying conformance to well-formedness rule R1, and for deltafication, since additional states have to be added for all divergent states in the original IOA. Recall from Def. 2.9 that a state s is divergent if there exists a fair divergent path on which s occurs infinitely often. Consequently, we need to find a fair cycle that starts at s and consists of only internal transitions. The presence of ‘internal’ cycles can be determined using Tarjan’s well-known and efficient strongly connected components (SCCs) algorithm [14].

One way to efficiently verify fairness is to utilise Streett automata [15], which form a variation on Büchi automata [16]. The acceptance condition for a Streett automaton depends on pairs of sets of states (E_i, F_i) (called Streett pairs), that together form the acceptance component Ω . An ω -word is accepted with $\Omega = \{(E_1, F_1), \dots, (E_k, F_k)\}$, if there exists a corresponding run that, for each j , only visits a state from F_j infinitely often if it visits a state from E_j infinitely often. This acceptance condition corresponds nicely with our notion of fairness.

Given an internal cycle $\pi = s_0 a_1 s_1 a_2 \dots a_n s_0$ with $a_i \in L^H$, let $L(\pi) = \{a_1, a_2, \dots, a_n\}$ be the set of actions executed on the path π , and $L^{LC}(s_i)$ be the set of locally controlled actions enabled at a state $s_i \in \text{states}(\pi)$. Because we require every divergent path to be state-finite (see Def. 2.8), these sets can always be calculated. If the cycle π is to be fair, then for every component $A_i \in P$ such that $A_i \cap L^{LC}(s_i) \neq \emptyset$ for some $s_i \in \text{states}(\pi)$, there must be an action $a_i \in A_i$ such that $a_i \in L(\pi)$. By introducing additional states that, when visited, represent the fact that a particular locally controlled action is executed, we translate this fairness condition to a nonemptiness check on a Streett automaton.

To clarify this construction, assume we wish to obtain the deltafication of the IOA \mathcal{A} shown in Fig. 6a given $P = \{A_1, A_2, A_3\}$, where $A_1 = \{a, \tau_1\}$, $A_2 = \{b, \tau_2\}$, and $A_3 = \{c\}$. First, we calculate the SCCs of \mathcal{A} , while only considering the internal transitions; in this case, there is only one: $SCC_1 = \{s_0, s_1, s_2\}$. To illustrate the conditions for an internal cycle to be fair, consider $\pi = s_0 \tau_1 s_1 \tau_2 s_0$. Since $L^{LC}(s_0) \cap A_1 = \{a, \tau_1\}$ and $L^{LC}(s_1) \cap A_2 = \{b, \tau_2\}$, it follows that for π

to be fair, there must be actions $a_i \in A_1$ and $a_j \in A_2$ such that $a_i \in L(\pi)$ and $a_j \in L(\pi)$. This indeed is the case for π , i.e., it is fair.

However, we do not know a priori that the fair path π exists. To find it, consider Fig. 6b. There, we introduced intermediate ‘transition’ states (marked grey) for every locally controlled transition in and leading out of SCC_1 . For state s_0 to be visited infinitely often, it follows from $L^{LC}(s_0) \cap A_1 \neq \emptyset$ and $L^{LC}(s_0) \cap A_2 = L^{LC}(s_0) \cap A_3 = \emptyset$ that there must be actions $a_i \in A_1$ that are executed infinitely often as well. Hence, one of the states a, τ_{1a}, τ_{1b} of $SA(SCC_1)$ must be visited infinitely often if s_0 is. For s_1 , in addition, actions from A_2 must occur infinitely often. Finally, for s_2 similar reasoning applies. All this yields $\Omega = \{(E_1, F_1), (E_2, F_2), (E_3, F_3), (E_4, F_4), (E_5, F_5)\}$, where $(E_1, F_1) = (\{a, \tau_{1a}, \tau_{1b}\}, \{s_0\})$, $(E_2, F_2) = (\{a, \tau_{1a}, \tau_{1b}\}, \{s_1\})$, $(E_3, F_3) = (\{b, \tau_{2a}, \tau_{2a}\}, \{s_1\})$, $(E_4, F_4) = (\{b, \tau_{2a}, \tau_{2a}\}, \{s_2\})$ and $(E_5, F_5) = (\{c\}, \{s_2\})$. As mentioned earlier, an accepting run in $SA(SCC_1)$ must satisfy all Streett pairs in Ω . Consequently, if such an accepting run exists, then it immediately follows that a fair internal cycle exists in \mathcal{A} . Such a nonemptiness check can be carried out efficiently using an optimised algorithm by Henzinger and Telle [17].

However, a fair internal cycle only gives us a subset of all divergent states. To find all of them, we need to verify for every state if a fair internal cycle exists that contains that particular state. Therefore, if we wish to check if, e.g., state s_0 is divergent, we need to extend acceptance component Ω with an additional Streett pair to obtain $\Omega_{s_0} = \Omega \cup \{(\{s_0\}, SCC_1)\}$. This way, we ensure that internal cycles in SCC_1 are only considered fair if they also contain state s_0 . Hence, $SA(SCC_1)$ has an accepting run with acceptance component Ω_{s_0} if and only if s_0 is divergent. In a similar way, we can construct $\Omega_{s_1} = \Omega \cup \{(\{s_1\}, SCC_1)\}$ and $\Omega_{s_2} = \Omega \cup \{(\{s_2\}, SCC_1)\}$ to check whether s_1 and s_2 are divergent, respectively.

Based on the above, we give an algorithm (Fig. 7) to determine divergent states. For clarity, we range over all states s and check nonemptiness using their acceptance condition Ω_s . A trivial improvement would be to, when a fair cycle is found, mark all its states as divergent and refrain from checking Ω_{s_i} for them.

Complexity. We discuss the worst-case time complexity of this algorithm given a DQTS with n states, m transitions and k partitions.

First note that the size of the acceptance condition of the Streett automaton for an SCC of n' states and m' transitions is worst-case in $O(n'k + n'm')$. After all, each of the n' states yields at most k Streett pairs (yielding the term $n'k$). Moreover, all Streett pairs corresponding to a state, together contain at most all states that represent transitions, of which there are m' (yielding the term $n'm'$).

The time complexity of `construct_streett_automaton(C)` is bounded by the size of the acceptance condition, and hence is in $O(n'(k + m'))$ (with n' and m' taken from C). As the function is called once for each SCC of the system, the total contribution of this function to the full algorithm is in $O(n(k + m))$. Additionally, Tarjan is called once, adding $O(n + m)$. Finally, in the worst-case scenario, the Henzinger/Telle algorithm, which is in

$$O(m \min\{\sqrt{m \log n}, k, n\} + n(k + m) \min\{\log n, k\})$$

```

algorithm determine_divergent_states is
input: IOA  $\mathcal{A} = \langle S, S^0, L^I, L^O, L^H, \rightarrow \rangle$ 
output:  $d(\mathcal{A})$ : a set containing all divergent states of  $\mathcal{A}$ 

 $d(\mathcal{A}) := \emptyset$ 

// Use a modified version of Tarjan's algorithm to determine SCCs( $\mathcal{A}$ )
 $SCCs(\mathcal{A}) :=$  the set of all SCCs of  $\mathcal{A}$  that are connected with internal transitions

for each  $C \in SCCs(\mathcal{A})$ 
  // Build the Streett automaton  $SA(C)$  corresponding to SCC  $C$ 
   $\langle S_{SA}, \rightarrow_{SA}, \Omega \rangle :=$  construct_streett_automaton( $C$ )

  for each state  $s$  in  $C$ 
    // Add an additional Streett pair to ensure  $s$  is on any accepting cycle
     $\Omega_s := \Omega \cup \{s\}, S_C$ 

    // Use the algorithm by Henziger and Telle to check the emptiness of  $SA(C)$ 
    if  $SA(C)$  has an accepting run with acceptance component  $\Omega_s$ 
       $d(\mathcal{A}) := d(\mathcal{A}) \cup \{s\}$ 
    end for
  end for

// Auxiliary function to construct the Streett automaton  $SA(C)$ , alongside acceptance
// component  $\Omega$ , for the given SCC  $C$ 
function construct_streett_automaton( $C$ )
input: SCC  $C = \langle S_{SCC}, \bar{L}^I, L^O, L^H, P, \rightarrow_{SCC} \rangle$ 
output: a Streett automaton  $SA(C) = \langle S_{SA}, \rightarrow_{SA}, \Omega \rangle$ 

 $S_{SA} := S_{SCC}$ 
 $\rightarrow_{SA} := \Omega := ts\_map := \emptyset$ 

// First construct the Streett automaton
for each  $(s, a, t) \in \rightarrow_{SCC}$  such that  $s \in S_{SCC}$  and  $a \in L^{LC}$ 
  // We need to insert a transition state for the transition  $(s, a, t)$ 
  let  $ts_{(s,a,t)} \notin S_{SA}$  be a new state
   $S_{SA} := S_{SA} \cup \{ts_{(s,a,t)}\}$ 

  if  $t \in S_{SCC}$  then  $\rightarrow_{SA} := \rightarrow_{SA} \cup \{(s, a, ts_{(s,a,t)}), (ts_{(s,a,t)}, a, t)\}$ 
  else  $\rightarrow_{SA} := \rightarrow_{SA} \cup \{(s, a, ts_{(s,a,t)})\}$ 

  let  $A \in P$  be the component such that  $a \in A$ 
   $ts\_map(A) := ts\_map(A) \cup \{ts_{(s,a,t)}\}$ 
end for

// Now construct the acceptance component  $\Omega$ 
for each  $s \in S_{SCC}$ 
  // Add a new Streett pair for every component whose actions are enabled in  $s$ 
  for each  $A \in P$  such that  $s \xrightarrow{a}_{SCC}$  for some  $a \in A$ 
     $\Omega := \Omega \cup \{(ts\_map(A), \{s\})\}$ 
  end for
end for

return  $\langle S_{SA}, \rightarrow_{SA}, \Omega \rangle$ 
end function

```

Fig. 7: Algorithm for detecting divergent states.

as shown in [17], is called once for each state. Together, this yields

$$O(n(k+m) + (n+m) + n(m \min\{\sqrt{m \log n}, k, n\} + n(k+m) \min\{\log n, k\}))$$

Under the reasonable assumption that k is bounded, and after simplification, we find that the worst-case time complexity is in $O(n^2m)$.

6 DQTSs in a testing context

Our main motivation for introducing and studying the DQTS model was to enable a clean theoretical framework for model-based testing. Earlier, the TGV framework [5] already defined `ioco` also in the presence of divergence. Although this was an important first step, it is not completely satisfactory in the sense that quiescence observations may be followed by output actions; this is counterintuitive to our notion of quiescence. Now, we illustrate how DQTSs can be incorporated in the `ioco` testing theory without having this problem.

The core of the `ioco` framework is its *conformance relation*, relating specifications to implementations if and only if the latter is ‘correct’ with respect to the former. For `ioco`, this means that the implementation never provides an unexpected output (including quiescence) when it is only fed inputs that are allowed by the specification. Traditionally, this was formalised based on the SAs corresponding to the implementation and the specification. Now, we can apply well-formed DQTSs, as they already model the expected absence of outputs by explicit δ -transitions. In addition, since DQTSs support divergence, using them as opposed to SAs also allows `ioco` to be applied in the presence of divergence.

Definition 6.1 (`ioco`). *Let $\mathcal{A}_{impl}, \mathcal{A}_{spec}$ be well-formed DQTSs over the same alphabet. Then, $\mathcal{A}_{impl} \sqsubseteq_{ioco} \mathcal{A}_{spec}$ if and only if*

$$\forall \sigma \in traces(\mathcal{A}_{spec}) . out_{\mathcal{A}_{impl}}(\sigma) \subseteq out_{\mathcal{A}_{spec}}(\sigma),$$

where $out_{\mathcal{A}}(\sigma) = \{a \in L^O \cup \{\delta\} \mid \sigma a \in traces(\mathcal{A})\}$.

Since all DQTSs are required to be input-enabled, it is easy to see that `ioco`-conformance precisely corresponds to traditional trace inclusion over well-formed DQTSs.

This improved notion of `ioco`-correspondence can be used as before [4, 18], at each point in time during testing choosing to either try to provide an input, observe the behaviour of the system or stop testing. As long as the trace obtained this way (including the δ actions, which can now be the result of either quiescence or divergence) is also a trace of the specification, the implementation is correct.

Note that the implementation and specification do not necessarily need to have the same task partition. After all, these are only needed to establish fair paths and hence divergences. This is used during deltafication, to determine which states are divergent. Although this influences `ioco` conformance (since it induces δ transitions), the conformance relation itself is not concerned with the task partitions anymore.

7 Conclusions and Future Work

In this paper, we introduced Divergent Quiescent Transition Systems (DQTSs) and investigated their properties. Also, we showed how to detect divergent states in order to construct the deltafication of an IOA, and discussed its complexity. Like SAs, DQTSs can be used to describe all possible observations of a system, including the observation of quiescence, i.e., the absence of outputs. Hence, DQTSs are especially useful to model specifications of reactive systems in the context of model-based testing. DQTSs for the first time allow the modelling of systems that exhibit divergence and explicit quiescence.

There are two advantages of using DQTSs rather than SAs for model-based testing. First, DQTSs allow more systems to be modelled naturally, as convergence is not required. Second, DQTSs are stand-alone entities whose properties have been investigated thoroughly. Hence, DQTSs are a formal and comprehensive theory to model and analyse quiescence, even in the presence of divergence.

We have shown that DQTSs are equally potent as SAs in terms of expressible observable behaviour, and that DQTSs can be used as a drop-in replacement for SAs in the `ioco` framework. Furthermore, we have proven that well-formed DQTSs exhibit desirable compositional properties. Consequently, composite systems can be represented as the parallel composition of smaller subcomponents.

Future Work. The action hiding operation for the DQTS model is quite complex, as outlined in Def. 4.4. To improve this, it might be useful to investigate a different strategy to mark quiescent and divergent states, e.g., using state labels. Also, `ioco`-based model-based testing tools like TORX internally still use the SA model to represent the specification of the system under test, and an SA-like model to represent the actual test cases. Hence, such tools should be adapted to utilise the improved `ioco` framework based on DQTSs. Work is currently already underway to adapt the TORX tool. Finally, it would be interesting to see if our notions could be phrased in a coalgebraic setting.

References

1. Vaandrager, F.W.: On the relationship between process algebra and input/output automata (extended abstract). In: Proceedings of the 6th Annual Symposium on Logic in Computer Science (LICS), IEEE Computer Society (1991) 387–398
2. Tretmans, J.: Test generation with inputs, outputs, and quiescence. In: Proceedings of the 2nd International Workshop on Tools and Algorithms for Construction and Analysis of Systems (TACAS). Volume 1055 of Lecture Notes in Computer Science., Springer (1996) 127–146
3. Tretmans, J.: Test generation with inputs, outputs and repetitive quiescence. *Software - Concepts and Tools* **17**(3) (1996) 103–120
4. Tretmans, J.: Model based testing with labelled transition systems. In: Formal Methods and Testing. Volume 4949 of Lecture Notes in Computer Science., Springer (2008) 1–38
5. Jard, C., Jéron, T.: TGV: theory, principles and algorithms. *International Journal on Software Tools for Technology Transfer* **7**(4) (2005) 297–315

6. Lynch, N.A., Tuttle, M.R.: An introduction to input/output automata. *CWI Quarterly* **2** (1989) 219–246
7. Stokkink, W.G.J., Timmer, M., Stoelinga, M.I.A.: Talking quiescence: a rigorous theory that supports parallel composition, action hiding and determinisation. In: *Proceedings of the 7th Workshop on Model-Based Testing (MBT)*. Volume 80 of *EPTCS*. (2012) 73–87
8. Stokkink, W.G.J., Timmer, M., Stoelinga, M.I.A.: Divergent quiescent transition systems (extended version). Technical Report TR-CTIT-13-08, University of Twente (2013)
9. Lynch, N.A., Tuttle, M.R.: Hierarchical correctness proofs for distributed algorithms. In: *Proceedings of the 6th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, ACM (1987) 137–151
10. Sudkamp, T.A.: *Languages and machines*. Pearson Addison Wesley (2006)
11. Baier, C., Katoen, J.P.: *Principles of Model Checking*. The MIT Press (2008)
12. De Nicola, R., Segala, R.: A process algebraic view of input/output automata. *Theoretical Computer Science* **138** (1995) 391–423
13. Willemse, T.: Heuristics for ioco-based test-based modelling. In: *Proceedings of the 11th International Workshop on Formal Methods: Applications and Technology (FMICS)*. Volume 4346 of *Lecture Notes in Computer Science*. Springer (2007) 132–147
14. Tarjan, R.E.: Depth-first search and linear graph algorithms (working paper). In: *Proceedings of the 12th Annual Symposium on Switching and Automata Theory (SWAT)*, IEEE Computer Society (1971) 114–121
15. Latvala, T., Heljanko, K.: Coping with strong fairness. *Fundamenta Informaticae* **43**(1-4) (2000) 175–193
16. Farwer, B.: ω -automata. In: *Proceedings of Automata, Logics, and Infinite Games*. Volume 2500 of *Lecture Notes in Computer Science*. Springer (2002) 3–21
17. Henzinger, M.R., Telle, J.A.: Faster algorithms for the nonemptiness of Streett automata and for communication protocol pruning. In: *Proceedings of the 5th Scandinavian Workshop on Algorithm Theory (SWAT)*. Volume 1097 of *Lecture Notes in Computer Science*, Springer (1996) 16–27
18. Timmer, M., Brinksma, E., Stoelinga, M.I.A.: Model-based testing. In: *Software and Systems Safety: Specification and Verification*. Volume 30 of *NATO Science for Peace and Security Series D*. IOS Press, Amsterdam (2011) 1–32