Fault-Tolerant Bisimulation and Process Transformations *

Tomasz Janowski **

Department of Computer Science
University of Warwick, Coventry CV4 7AL, UK

Abstract. We provide three methods of verifying concurrent systems which are tolerant of faults in their operating environment - algebraic, logical and transformational. The first is an extension of the bisimulation equivalence, the second is rooted in the Hennessy-Milner logic, and the third involves transformations of CCS processes. Based on the common semantic model of labelled transition systems, which is also used to model faults, all three methods are proved equivalent for certain classes of faults.

1 Introduction

Many models of concurrent systems have been proposed in the literature, based on either actions or states. Examples include sequences [MP91], trees [Mil89], machines [LT87], partial orders [Pra86] and event structures [Win89]. They offer different ways of representing executions of systems (linear or branching), their concurrent activity (interleaving or non-interleaving) and interaction (shared memory or message-passing). A concept which unifies various models is a *labelled transition system* [Kel76], a triple $(\mathcal{P}, \mathcal{A}, \rightarrow)$ where $\mathcal{P}$ is a set of processes, $\mathcal{A}$ a set of actions and $\rightarrow \subseteq \mathcal{P} \times \mathcal{A} \times \mathcal{P}$ a labelled transition relation. Labelled transition relations are often defined by induction on the structure of processes, providing the *structured operational semantics* [Plo81] of process description languages. An example of such a language is CCS [Mil89].

As models of processes, labelled transition systems describe their behaviour in detail, including particulars of their internal computation. However, in order to specify a process and then to prove its correctness, it is useful to decide which properties of the model are relevant and which can be ignored. Following [Mil89], it is most common to ignore these properties which cannot be observed in the finite interval of time. Two ways to do so are as follows:

- we can identify a process with its equivalence class, according to the (weak) *bisimulation equivalence* $\approx$ [Par81];
- we can identify a process with its properties, specified by the (weak) formulas of the Hennessy-Milner logic [HM85] and verified by satisfaction relation $\models$.

* To be presented at the Third International Symposium "Formal Techniques in Real-Time and Fault-Tolerant Systems", Lübeck, Germany, September 1994.

** Supported by the University of Warwick, under its Scholarship Scheme for East Europe, and by an Overseas Students Award from CVCP.

A vital test of the usefulness of any formal theory is that statements of this theory must be confirmed in practice (by experiment). Given such statements as $P \approx Q$ or $Q \models M$, it is expected that the low-level process Q, when placed in the real environment, behaves respectively as specified by the high-level process P or the formula M. In practice however, such Q depends on various hardware components which often malfunction because of the physical faults. Such faults affect the semantics of Q so that it may no longer behave as specified. Moreover, physical faults do not exist before Q is put into practice and so cannot be removed beforehand, they must be *tolerated*.

Clearly, it is not possible to tolerate arbitrary faults. We have to decide which faults are *anticipated* (and thus should be tolerated) and which are not (such faults are *catastrophical*). To represent the effect of the anticipated faults on the semantics of processes, we will use the set $\dashrightarrow$ of the *faulty transitions*. To verify fault-tolerance is then to prove that the low-level process behaves 'correctly' in the presence of the transitions $\dashrightarrow$. As such, fault-tolerance depends on the chosen notion of correctness. In this paper we provide three methods to verify fault-tolerance for bisimulation equivalence and the Hennessy-Milner logic:

1. A fault-tolerant bisimilarity $\boxplus$ where $P \boxplus Q$ if observing P in the fault-free environment (performing transitions $\longrightarrow$) and Q in the environment which contains anticipated faults (performing transitions $\longrightarrow$ and $\dashrightarrow$), we cannot tell them apart in the finite interval of time.
2. A relation $\Vdash$ to verify satisfaction of formulas of the Hennessy-Milner logic in the presence of the anticipated faults (when processes $\mathcal{P}$ undergo both normal and faulty transitions $\longrightarrow \cup \dashrightarrow$).
3. A language $\mathcal{D}$ for specifying faults and a process transformation $\mathcal{T}(Q, \Psi)$ where given the CCS process Q, the effect of transitions $\overset{\cdots}{\underset{\Psi}{\rightarrow}}$ (specified by $\Psi \in \mathcal{D}$) on Q are represented syntactically. Then, verifying that Q is fault-tolerant involves proving either:
   (a) $P \approx \mathcal{T}(Q, \Psi)$ for the high-level process P, or
   (b) $\mathcal{T}(Q, \Psi) \models M$ for the formula M.

We show that, for wide classes of faults, all these methods are equivalent:

$$\begin{array}{ccc} P \boxplus Q & \longleftrightarrow & \forall_{M \in \mathcal{M}}\, P \models M \ \text{ iff } \ Q \Vdash M \\ \updownarrow & & \updownarrow \\ P \approx \mathcal{T}(Q, \Psi) & \longleftrightarrow & \forall_{M \in \mathcal{M}}\, P \models M \ \text{ iff } \ \mathcal{T}(Q, \Psi) \models M \end{array} \qquad (1)$$

When the 'full' fault-tolerance is either impossible or too expensive to ensure, we may be still satisfied with its conditional version, given certain assumption about the quantity of faults. To this end we will use $n \in \mathbb{N} \cup \{\infty\}$ as the maximal number of times transitions $\dashrightarrow$ can occur successively (if $n = \infty$ then $\dashrightarrow$ can occur at any time; if $n = 0$ then not at all). As before, we provide and prove the equivalence of the three methods for verifying n-conditional fault-tolerance: relations $\boxplus_n$ and $\Vdash_n$, and transformation $\widetilde{\mathcal{T}}(\ .\ , \Psi, n)$.

The rest of this paper is as follows. In Section 2 we describe the semantic model. 'Intolerant' bisimulation equivalence, its fault-tolerant and conditional fault-tolerant versions are defined in Section 3. Their logical characterisation, in terms of the Hennessy-Milner logic, is given in Section 4. Both languages, of processes and faults, are defined in Sections 5 and 6, followed by transformations $\mathcal{T}$ and $\widetilde{\mathcal{T}}$ which are shown to provide the third, equivalent method of verifying fault-tolerance and conditional fault-tolerance in Section 7. Finally, in Section 8, we draw some conclusions and comment on the directions for future work.

2 Semantic Model

Consider the labelled transition relation $\longrightarrow$. If $(P, \alpha, P') \in \longrightarrow$ then we write $P \xrightarrow{\alpha} P$ and say that P performs α and evolves into P' (we also use $P \xrightarrow{s} P'$ for the action sequence $s \in \mathcal{A}^*$). One kind of transition we wish to largely ignore is $\xrightarrow{\tau}$ where action τ is unobservable and represents the outcome of a joint activity (interaction) between two processes. Interaction takes place on the pair of complementary actions $a, \overline{a} \in \mathcal{L}$ where $\mathcal{L} =_{def} \mathcal{A} - \{\tau\}$ is the set of observable actions and $\overline{\alpha}$ is a function over $\mathcal{A}$ which is bijective and such that $\overline{\overline{a}} = a, \overline{a} \neq \tau$ and $\overline{\tau} = \tau$. We use α to range over $\mathcal{A}$ and β to range over $\mathcal{L}_\varepsilon =_{def} \mathcal{L} \cup \{\varepsilon\}$ where ε denotes the empty sequence. We also let $\widehat{\cdot}$ be the function over $\mathcal{A}^*$ such that $\widehat{\varepsilon} = \varepsilon$, $\widehat{\tau : s} = \widehat{s}$ and $\widehat{a : s} = a : \widehat{s}$ (: denotes concatenation).

When placed in the 'real' environment, a process may not behave according to $\longrightarrow$: it may either perform transitions $\dashrightarrow$ which do not belong to $\longrightarrow$, $\dashrightarrow \cap \longrightarrow = \emptyset$, or it may refuse to perform some of the transitions $\longrightarrow$. The first case is demonstrated by transition $P \xrightarrow{c} P''$ in Figure 1, the second (in part) by $Q \xrightarrow{b} Q'''$. To represent the second case in full, we should physically remove $Q \xrightarrow{b} Q'''$ from the diagram. This would complicate our model so suppose only that this transition *may be* refused. This is achieved by two more transitions $Q \xrightarrow{\tau} Q'' \xrightarrow{a} Q'$ which may preempt (due to the occurrence of τ) transition $Q \xrightarrow{b} Q'''$. Given the set $\dashrightarrow$ of faulty transitions as the effect of faults, we let $\mapsto$ contain both kinds of transitions, $\mapsto =_{def} \longrightarrow \cup \dashrightarrow$ (we also use $\stackrel{s}{\mapsto}$ for the action sequence $s \in \mathcal{A}^*$).

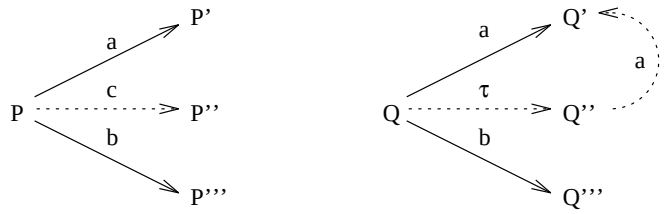

Fig. 1. Transition diagrams of P and Q with faulty transitions.

3 Bisimulation and Fault-Tolerance

There are many equivalences by which to abstract from the behavioural details of the transition relation $\longrightarrow$. They differ, among other things, in the adopted model of execution (linear- or branching-time) and concurrency (interleaving or non-interleaving). The best known of them and deemed to be the strongest among interleaving and branching-time equivalences is that of bisimulation equivalence [Par81], $\approx$. Two bisimilar processes, their semantics defined by transition relation $\longrightarrow$, cannot be distinguished by observing them in the finite interval of time. This property may no longer hold in the presence of faults which result in the additional transitions $\dashrightarrow$ (of the low-level process). As such, $\approx$ allows to verify correctness in the absence of faults only, it is *fault-intolerant*. How to verify *fault-tolerance*, i.e. correctness in the presence of transitions $\dashrightarrow$, and *conditional fault-tolerance*, where transitions $\dashrightarrow$ occur under assumption n about their quantity, is the topic of the current section.

3.1 Fault-Intolerance

Bisimulation equivalence is defined as the maximal fixed point of the functional $\mathcal{F}$ on the set of binary relations B on $\mathcal{P}$, $(P,Q) \in \mathcal{F}(B)$ iff

$$\begin{array}{l} \text{whenever } P \xrightarrow{\alpha} P' \text{ then } \exists_{Q',s}\ Q \xrightarrow{s} Q' \wedge \widehat{s} = \widehat{\alpha} \wedge (P',Q') \in B \\ \text{whenever } Q \xrightarrow{\alpha} Q' \text{ then } \exists_{P',s}\ P \xrightarrow{s} P' \wedge \widehat{s} = \widehat{\alpha} \wedge (P',Q') \in B \end{array} \qquad (2)$$

This maximal fixed point exists because $\mathcal{F}$ is monotonic: if $B_1 \subseteq B_2$ then $\mathcal{F}(B_1) \subseteq \mathcal{F}(B_2)$. Originally, it was reached 'from above', as the limit of the sequence $\mathcal{F}^n(\mathcal{P} \times \mathcal{P})$ for all $n \geq 0$. Unfortunately, unless infinite ordinals n are taken into account, this requires that transition relations $\xrightarrow{\alpha}$ are weak-image-finite i.e. that for all $P \in \mathcal{P}$, the set $\{P' \mid P \xrightarrow{s} P' \wedge \widehat{s} = \widehat{\alpha}\}$ is finite. No such assumption is needed to reach $\approx$ 'from below', as the union of all pre-fixed points B of $\mathcal{F}$, $\approx =_{def} \bigcup\{B \mid B \subseteq \mathcal{F}(B)\}$ [Par81]. An additional advantage is the useful technique for proving $P \approx Q$. It is enough to find a pre-fixed point B of $\mathcal{F}$ such that $(P,Q) \in B$. Such a B is called a *bisimulation*. We have:

$$\begin{array}{r} P \approx Q \text{ iff whenever } P \xrightarrow{\alpha} P' \text{ then } \exists_{Q',s}\ Q \xrightarrow{s} Q' \wedge \widehat{s} = \widehat{\alpha} \wedge P' \approx Q' \\ \text{whenever } Q \xrightarrow{\alpha} Q' \text{ then } \exists_{P',s}\ P \xrightarrow{s} P' \wedge \widehat{s} = \widehat{\alpha} \wedge P' \approx Q' \end{array} \qquad (3)$$

3.2 Fault-Tolerance

If observing two processes, the high-level in the fault-free environment (performing transitions $\longrightarrow$) and the lower-level in the environment which is affected by the anticipated faults (performing transitions $\longmapsto$), we cannot tell them apart in the finite interval of time, then we say that the lower-level process is fault-tolerant (with respect to the high-level one). To verify this property we provide two relations, *must-bisimilarity* $\boxdot$ and *may-bisimilarity* $\boxtimes$.

The first is the direct extension of $\approx$ to take account of transitions $\stackrel{\cdots\cdots}{\longrightarrow}$. We have $P \sqsubset\!\!\sim Q$ if P and Q are 'bisimilar', the first performing transitions $\longrightarrow$ and the second both $\longrightarrow$ and $\stackrel{\cdots\cdots}{\longrightarrow}$. We define $\sqsubset\!\!\sim$ using a *must-bisimulation* B which is a binary relation such that if $(P,Q) \in B$ then any $\longrightarrow$ transition of P is matched by some transition sequence $\longmapsto$ of Q and any $\longmapsto$ transition of Q is matched by some transition sequence $\longrightarrow$ of P, such that the matched transitions have the same observable actions and B is preserved:

$$\begin{array}{l} \text{whenever } P \stackrel{\alpha}{\longrightarrow} P' \text{ then } \exists_{Q',s}\ Q \stackrel{s}{\longmapsto} Q' \wedge \widehat{s} = \widehat{\alpha} \wedge (P',Q') \in B \\ \text{whenever } Q \stackrel{\alpha}{\longmapsto} Q' \text{ then } \exists_{P',s}\ P \stackrel{s}{\longrightarrow} P' \wedge \widehat{s} = \widehat{\alpha} \wedge (P',Q') \in B \end{array} \qquad (4)$$

Then $P \sqsubset\!\!\sim Q$ iff $(P,Q) \in B$ for some must-bisimulation B. Such Q satisfies the basic postulate: no external observer can distinguish between P which behaves according to transitions $\longrightarrow$ and Q which may additionally perform transitions $\stackrel{\cdots\cdots}{\longrightarrow}$. In one aspect however, such Q is unsatisfactory. Because $\sqsubset\!\!\sim$ allows to match transitions of the high-level process by faulty transitions of the low-level one, such Q may not behave properly in the environment where not all transitions $\stackrel{\cdots\cdots}{\longrightarrow}$ are provided. For Q to behave as specified, transitions $\stackrel{\cdots\cdots}{\longrightarrow}$ *must* occur. In practice however, it is more useful is to assume the mere possibility of faults (that faults *may* occur), not their necessity (that they *must* occur).

This assumption is met by may-bisimilarity $P \sqsubset\!\!\approx Q$ where only normal transitions of Q are allowed to match transitions of P. As before, may-bisimilarity $\sqsubset\!\!\approx$ is defined as the largest *may-bisimulation* which is a binary relation B such that if $(P,Q) \in B$ then:

$$\begin{array}{l} \text{whenever } P \stackrel{\alpha}{\longrightarrow} P' \text{ then } \exists_{Q',s}\ Q \stackrel{s}{\longrightarrow} Q' \wedge \widehat{s} = \widehat{\alpha} \wedge (P',Q') \in B \\ \text{whenever } Q \stackrel{\alpha}{\longmapsto} Q' \text{ then } \exists_{P',s}\ P \stackrel{s}{\longrightarrow} P' \wedge \widehat{s} = \widehat{\alpha} \wedge (P',Q') \in B \end{array} \qquad (5)$$

Example 1. Consider the high-level process P in Figure 2 and four low-level, fault-affected processes Q, R, S and T. We have:

1. $P \approx Q$ because $\{(P,Q),(P',Q'),(P',Q'')\}$ is a bisimulation but $P \not\sqsubset\!\!\sim Q$ and $P \not\sqsubset\!\!\approx Q$ because there is no must- or may-bisimulation which contains (P,Q): $Q \stackrel{\tau}{\longmapsto} Q'''$ and $P \stackrel{s}{\longrightarrow} P$ $(\widehat{s} = \varepsilon)$ only, however $P \stackrel{a}{\longrightarrow} P'$ but $Q''' \not\longmapsto$.
2. $P \sqsubset\!\!\sim R$ because $\{(P,R),(P,R'''),(P',R'),(P',R'')\}$ is a must-bisimulation but $P \not\approx R$ and $P \not\sqsubset\!\!\approx R$ because there is no bisimulation or may-bisimulation which contains (P,R): $P \stackrel{a}{\longrightarrow} P'$ and $R \stackrel{s}{\longrightarrow} R'$ $(\widehat{s} = a)$ only, however $P' \stackrel{b}{\longrightarrow} P$ but R' has no normal transitions, $R' \not\longrightarrow$.
3. $P \sqsubset\!\!\sim S$ because $\{(P,S),(P,S'''),(P',S'),(P',S'')\}$ is a must-bisimulation and $P \approx S$ because $\{(P,S),(P',S'),(P',S'')\}$ is a bisimulation. Also $P \not\sqsubset\!\!\approx S$ because there is no may-bisimulation which contains (P,S): $S \stackrel{\tau}{\longmapsto} S'''$ and $P \stackrel{s}{\longrightarrow} P$ $(\widehat{s} = \varepsilon)$ only, however $P \stackrel{a}{\longrightarrow} P'$ but $S''' \not\stackrel{a}{\longmapsto}$.
4. $P \sqsubset\!\!\sim T$, $P \sqsubset\!\!\approx T$ and $P \approx T$ because $\{(P,T),(P,T'''),(P',T'),(P',T'')\}$ is simultaneously a must-bisimulation, may-bisimulation and bisimulation. $\square$

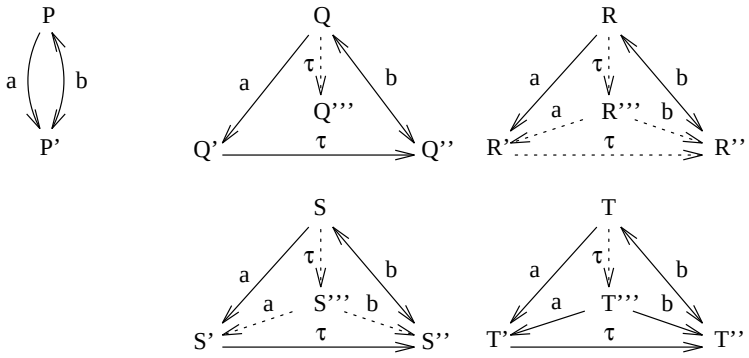

Fig. 2. Transition diagrams of P, Q, R, S and T with faulty transitions.

The example shows that $\approx$ and $\sqsubset$ are not comparable: $\sqsubset$ does not imply $\approx$, nor does $\approx$ imply $\sqsubset$. However, it is easy to show that any may-bisimulation is simultaneously a must-bisimulation and a bisimulation. As a result, because all relations are defined as the union of the corresponding bisimulations, we have:

$$\sqsubset \quad \subseteq \quad \sqsubset \cap \approx \tag{6}$$

The example ($P \sqsubset S$ and $P \approx S$ but $P \not\sqsubset S$) also shows that this inclusion is proper i.e. that $P \sqsubset Q$ and $P \approx Q$ together are not enough to establish $P \sqsubset Q$. That is a pity since $P \sqsubset Q$ which is more desirable than $P \sqsubset Q$, is also more difficult to establish (the equivalence diagram (1) for $\sqsubset$ is only partly valid for may-bisimilarity $\sqsubset$). However, for B to be a may-bisimulation, it is not only necessary but also sufficient that B is a bisimulation and a must-bisimulation:

B is a may-bisimulation iff it is a must-bisimulation and a bisimulation. (7)

Thus in order to prove $P \sqsubset Q$, it is enough to show that $(P, Q) \in B$ for B which is a bisimulation and a must-bisimulation at the same time. This justifies our efforts to establish the properties of $\sqsubset$ in the first instance. What both relations have in common is that neither of them is reflexive or symmetric (they are not preorders). For processes in Figure 2 we have:

- $Q \not\sqsubset Q$ because there is no must-bisimulation which contains the pair (Q, Q): $Q \stackrel{\tau}{\longmapsto} Q'''$ and $Q \stackrel{s}{\longrightarrow} Q$ ($\widehat{s} = \varepsilon$) only, however $Q \stackrel{a}{\longrightarrow} Q'$ but $Q''' \stackrel{a}{\not\longmapsto}$. Consequently $Q \not\sqsubset Q$ because of (6).
- $Q \sqsubset P$ because $\{(Q, P), (Q', P'), (Q'', P')\}$ is a may-bisimulation but $P \not\sqsubset Q$ because $Q \stackrel{\tau}{\longmapsto} Q'''$ and $P \stackrel{s}{\longrightarrow} P$ ($\widehat{s} = \varepsilon$) only, however $P \stackrel{a}{\longrightarrow} P'$ but $Q''' \stackrel{a}{\not\longmapsto}$. Consequently we have $Q \sqsubset P$ and $P \not\sqsubset Q$, as well as $Q \sqsubset P$ and $P \not\sqsubset Q$.

The lack of these properties is not unexpected when verifying correctness in the presence of faults. Because one and the same process has two different semantics, as the high-level (fault-free) process and as the low-level (fault-affected) one, we cannot ensure that the underlying relation is reflexive or symmetric.

Transitivity is most desirable to support the stepwise development of processes and to support the reasoning in the presence of faults where it may be helpful to deal with only some of transitions $\cdots\!\!\rightarrow$ (not all) at a time. To this end let us partition $\cdots\!\!\rightarrow$ among $m > 0$ nonempty disjoint sets $\cdots\!\!\rightarrow_j$, $\cdots\!\!\rightarrow = \bigcup_{j=1}^{m} \cdots\!\!\rightarrow_j$. Given $j = 0, \ldots, m$ we define $\stackrel{}{\underset{j}{\longmapsto}}$ as the union of the normal transitions $\longrightarrow$ and the first j partitions of $\cdots\!\!\rightarrow$: $\stackrel{}{\underset{j}{\longmapsto}} =_{def} \longrightarrow \cup \bigcup_{i=1}^{j} \cdots\!\!\rightarrow_i$. This gives an ascending sequence $\longrightarrow = \stackrel{}{\underset{0}{\longmapsto}} \subseteq \cdots \subseteq \stackrel{}{\underset{m}{\longmapsto}} = \longmapsto$ of the transition relations.

Suppose now that for $0 \leq j \leq l \leq m$, $\boxtimes_l^j$ and $\boxplus_l^j$ denote the corresponding bisimulation relations where transitions $\stackrel{}{\underset{j}{\longmapsto}}$ are regarded as normal and transitions $\bigcup_{i=j+1}^{l} \cdots\!\!\rightarrow_i$ as abnormal:

$$\begin{array}{r} P \boxtimes_l^j Q \text{ iff whenever } P \stackrel{\alpha}{\underset{j}{\longmapsto}} P' \text{ then } \exists_{Q',s}\ Q \stackrel{s}{\underset{j}{\longmapsto}} Q' \wedge \widehat{s} = \widehat{\alpha} \wedge P' \boxtimes_l^j Q' \\ \text{whenever } Q \stackrel{\alpha}{\underset{l}{\longmapsto}} Q' \text{ then } \exists_{P',s}\ P \stackrel{s}{\underset{j}{\longmapsto}} P' \wedge \widehat{s} = \widehat{\alpha} \wedge P' \boxtimes_l^j Q' \end{array} \tag{8}$$

Relation $\boxplus_l^j$ is defined alike. Then, given $j \leq k \leq l$, we can easily prove the following transitive properties of $\boxplus$ and $\boxtimes$ ($\circ$ is the relational composition):

$$\boxtimes_k^j \circ \boxtimes_l^j \subseteq \boxtimes_l^j \qquad\qquad \boxplus_k^j \circ \boxplus_l^k \subseteq \boxplus_l^j \tag{9}$$

According to the first inclusion, to tolerate transitions $\bigcup_{i=j+1}^{l} \cdots\!\!\rightarrow_i$ (given $\boxtimes$), at least once we must tolerate them altogether. According to the second, to tolerate transitions $\bigcup_{i=j+1}^{l} \cdots\!\!\rightarrow_i$ (with respect to $\boxplus$), it is enough to first tolerate transitions $\bigcup_{i=j+1}^{k} \cdots\!\!\rightarrow_i$ and then transitions $\bigcup_{i=k+1}^{l} \cdots\!\!\rightarrow_i$. Following the first inclusion, it is easy to see that $\boxtimes$ is transitive. This is not the case for $\boxplus$ and in general the first inclusion does not hold for $\boxplus$ and the second for $\boxtimes$, as shown by processes in Figure 3. We have: $P \boxplus Q \boxplus R$ but clearly $P \not\boxplus R$. Also, $P \boxtimes_1^0 S$ and $S \boxtimes_2^1 T$ but $P \not\boxtimes_2^0 T$ because $T \stackrel{\tau}{\cdots\!\!\rightarrow_2} T'''$ and $P \stackrel{s}{\underset{0}{\longmapsto}} P$ ($\widehat{s} = \varepsilon$) only, however $P \stackrel{a}{\underset{0}{\longmapsto}} P'$ but $T''' \not\stackrel{}{\underset{0}{\longmapsto}}$.

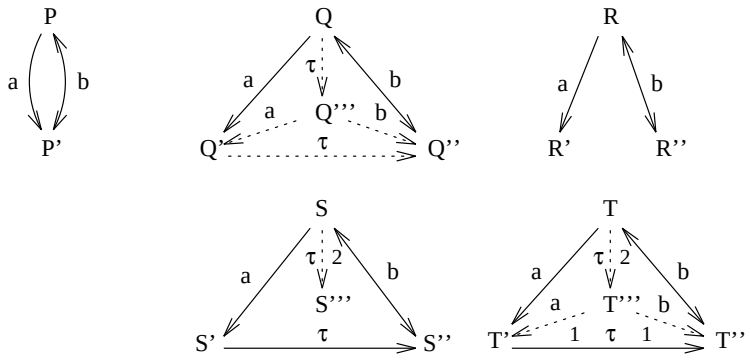

Fig. 3. Transition diagrams of P, Q, R, S and T with faulty transitions.

3.3 Conditional Fault-Tolerance

For any may- or must-bisimilar processes P and Q, $\cdots\!\!\rightarrow$ is the assumption about faults of the operating environment of Q, where Q is guaranteed to behave 'properly', as specified by P. We call $\cdots\!\!\rightarrow$ a *qualitative* assumption, in opposite to the *quantitative* assumptions $n \in \mathbb{N} \cup \{\infty\}$ which are introduced in this section and specify the maximal number of times transitions $\cdots\!\!\rightarrow$ can occur successively (if $n = 0$ then $\cdots\!\!\rightarrow$ are assumed not to occur at all; if $n = \infty$ then they can occur at any time). The reasons for introducing such assumptions are threefold:

- For certain sets $\cdots\!\!\rightarrow$, we cannot ensure fault-tolerance in full. In these circumstances, we must be satisfied with its degraded, conditional version, for certain assumptions about the quantity of $\cdots\!\!\rightarrow$.
- Even when the 'full' fault-tolerance is (in theory) possible, we may choose its conditional version because it is often easier to do so. This argument is true for applications which are not safety-critical.
- Conditional fault-tolerance may facilitate the stepwise procedure where Q is first designed for restricted assumptions about faults and then stepwise transformed for increasingly relaxed assumptions.

Recall that if $Q \stackrel{s}{\longmapsto} Q'$ then Q evolves into Q' performing the sequence s of transitions $\longrightarrow$ and $\cdots\!\!\rightarrow$. This may be no longer the case if transitions $\cdots\!\!\rightarrow$ can only occur under assumption n. It these circumstances we will use the family $\{\stackrel{i}{\longmapsto}_j\}_{i,j=0}^{n}$ of relations $\stackrel{i}{\longmapsto}_j \subseteq \mathcal{P} \times \mathcal{A}^* \times \mathcal{P}$. If $(Q, s, Q') \in \stackrel{i}{\longmapsto}_j$ then we write $Q \stackrel{s}{\stackrel{i}{\longmapsto}}_j Q'$ and say that Q evolves into Q' by the sequence $s \in \mathcal{A}^*$ of transitions $\longmapsto$, under assumption n, and given that i is the number of times transitions $\cdots\!\!\rightarrow$ have successively occurred before and j after $Q \stackrel{s}{\stackrel{i}{\longmapsto}}_j Q'$. Formally, relations $\stackrel{i}{\longmapsto}_j$ are defined by the following inductive rules:

$$\begin{array}{l} Q \stackrel{\varepsilon}{\stackrel{i}{\longmapsto}}_i Q \\ Q \stackrel{\alpha:s}{\stackrel{i}{\longmapsto}}_j Q'' \text{ whenever } \exists_{Q'}\ Q \stackrel{\alpha}{\longrightarrow} Q' \stackrel{s}{\stackrel{0}{\longmapsto}}_j Q'' \quad \vee \\ \qquad\qquad\qquad\qquad\qquad\quad Q \stackrel{\alpha}{\cdots\!\!\rightarrow} Q' \stackrel{s}{\stackrel{i+1}{\longmapsto}}_j Q'' \wedge i \neq n \end{array} \tag{10}$$

The induction above is well-defined: the first rule provides the base, for the empty sequence ε, and the second rule decreases the length of the action sequence by one. Given $n = 0$, we always have $i = n$, so transitions $\cdots\!\!\rightarrow$ cannot occur at all. Given $n = \infty$, it is never the case that $i = n$, so $\cdots\!\!\rightarrow$ can occur at any time.

Consider conditional version of $\boxtimes$, $\boxtimes_n$. We have $P \boxtimes_n Q$ if observing P in the fault-free environment (performing transitions $\longrightarrow$) and Q in any environment where it *may* also perform transitions $\cdots\!\!\rightarrow$ (provided no more than n times in a row), we cannot distinguish between them in the finite amount of time. In order to keep track of the number i of the successive transitions $\cdots\!\!\rightarrow$, $\boxtimes_n$ is defined using $[0, n]$-indexed families $\{B_i\}_{i=0}^{n}$ of binary relations $B_i \subseteq \mathcal{P} \times \mathcal{P}$. Such a family $\{B_i\}_{i=0}^{n}$ is called a conditional must-bisimulation iff for all $i, j \in [0, n]$ and $\alpha \in \mathcal{A}$, if $(P, Q) \in B_i$ then:

$$\begin{array}{l} \text{whenever } P \stackrel{\alpha}{\longrightarrow} P' \text{ then } \exists_{Q',s,j}\ Q \stackrel{s}{\stackrel{i}{\longmapsto}}_j Q' \wedge \widehat{s} = \widehat{\alpha} \wedge (P', Q') \in B_j \\ \text{whenever } Q \stackrel{\alpha}{\stackrel{i}{\longmapsto}}_j Q' \text{ then } \exists_{P',s}\ P \stackrel{s}{\longrightarrow} P' \wedge \widehat{s} = \widehat{\alpha} \wedge (P', Q') \in B_j \end{array} \tag{11}$$

Similarly, $\{B_i\}_{i=0}^n$ is called a conditional may-bisimulation iff for all $i, j \in [0,n]$ and $\alpha \in \mathcal{A}$, if $(P,Q) \in B_i$ then:

$$\begin{array}{l} \text{whenever } P \xrightarrow{\alpha} P' \text{ then } \exists_{Q',s}\ Q \xrightarrow{s} Q' \wedge \widehat{s} = \widehat{\alpha} \wedge (P',Q') \in B_i \\ \text{whenever } Q \overset{\alpha}{\longmapsto}{}_j^i Q' \text{ then } \exists_{P',s}\ P \xrightarrow{s} P' \wedge \widehat{s} = \widehat{\alpha} \wedge (P',Q') \in B_j \end{array} \tag{12}$$

Let $i \in [0,n]$. We define relations $\sqsubset\!\!\sim_n^i$ and $\sqsubset\!\!\approx_n^i$ as follows:

$P \sqsubset\!\!\sim_n^i Q$ iff $(P,Q) \in B_i$ for some $\{B_i\}_{i=0}^n$
which is a conditional must-bisimulation
$P \sqsubset\!\!\approx_n^i Q$ iff $(P,Q) \in B_i$ for some $\{B_i\}_{i=0}^n$
which is a conditional may-bisimulation

Then we have $P \sqsubset\!\!\sim_n Q$ iff $P \sqsubset\!\!\sim_n^0 Q$ and $P \sqsubset\!\!\approx_n Q$ iff $P \sqsubset\!\!\approx_n^0 Q$.

Example 2. Consider processes in Figure 4. We have:

- $P \not\sqsubset\!\!\sim Q$ and $P \not\sqsubset\!\!\sim_2 Q$ because of transition $Q''' \overset{a}{\dashrightarrow} Q$ which enables two subsequent actions a. However $P \sqsubset\!\!\sim_1 Q$ because then $Q''' \overset{a}{\dashrightarrow} Q$ (as the second one in a row) cannot be chosen.
- $P \not\sqsubset\!\!\approx R$ and $P \not\sqsubset\!\!\approx_3 R$ because of transition $R''' \overset{\tau}{\dashrightarrow} R''$ which, after performing a and b, leads to R'' where action a is not possible. However $P \sqsubset\!\!\approx_2 R$ because then transition $R''' \overset{\tau}{\dashrightarrow} R''$ (as the third one in a row) cannot be chosen. Finally $P \not\sqsubset\!\!\approx_1 R$ because then $R' \overset{b}{\dashrightarrow} R'''$ cannot be taken but is needed to match transitions of P. □

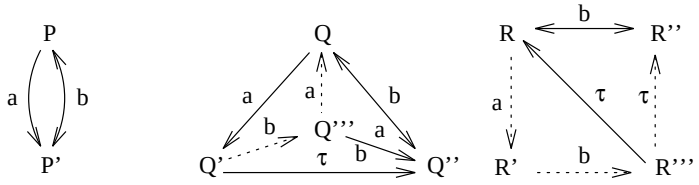

Fig. 4. Transition diagrams of P, Q and R with abnormal transitions.

Conditional may-bisimilarity $\sqsubset\!\!\approx_n$ is monotone decreasing with respect to n. This is not the case for $\sqsubset\!\!\sim_n$, as shown by $P \sqsubset\!\!\sim_2 R$ and $P \not\sqsubset\!\!\sim_1 R$ in Figure 4. For $n = \infty$, $\sqsubset\!\!\sim_\infty$ and $\sqsubset\!\!\approx_\infty$ coincide with their unconditional versions; for $n = 0$, they coincide with $\approx$. We have the following diagram of inclusions:

$$\begin{array}{ccccccccc} \sqsubset\!\!\approx & = & \sqsubset\!\!\approx_\infty & \subset \cdots \subset & \sqsubset\!\!\approx_{n+1} & \subset & \sqsubset\!\!\approx_n & \subset \cdots \subset & \sqsubset\!\!\approx_0 = \approx \\ & & \cap & & \cap & & \cap & & \| \\ \sqsubset\!\!\sim & = & \sqsubset\!\!\sim_\infty & & \sqsubset\!\!\sim_{n+1} & & \sqsubset\!\!\sim_n & & \sqsubset\!\!\sim_0 = \approx \end{array} \tag{13}$$

4 Logic of Processes

The Hennessy-Milner logic [HM85] is a simple modal logic for specifying properties of processes. It provides a language $\mathcal{M}$ of formulas M which extends propositional logic by the modal operators $\langle\beta\rangle M$. $\mathcal{M}$ is defined by the grammar:

$$M \ ::= \ true \mid M \wedge M \mid \neg M \mid \langle\beta\rangle M \tag{14}$$

The semantics of M (the set of all processes which satisfy M) is defined by relation $\models \subseteq \mathcal{P} \times \mathcal{M}$ where if $(P, M) \in \models$ then we write $P \models M$. Following [HM85], $\models$ is defined as the least set such that:

$$\begin{array}{l} P \models true \\ P \models M \wedge N \ \text{ iff } \ P \models M \ \wedge \ P \models N \\ P \models \neg M \quad \text{ iff not } P \models M \\ P \models \langle\beta\rangle M \quad \text{ iff } \exists_{P',s} \ \ P \xrightarrow{s} P' \ \wedge \ \widehat{s} = \beta \ \wedge \ P' \models M \end{array} \tag{15}$$

We abbreviate $\neg true$ as $false$, $\neg\langle\beta\rangle\neg M$ as $[\beta]M$ and for $m > 1$ define $\langle\beta\rangle^m M$ as $\langle\beta\rangle\langle\beta\rangle^{m-1}M$ and $\langle\beta\rangle^1 M$ as $\langle\beta\rangle M$ ($[\beta]^m M$ is defined alike).

Algebraically, we can identify a process with its equivalence class. However, given a logic where properties of processes can be stated and verified, we can identify a process with its properties. When both algebraic and logical views agree, that is when two processes are equivalent iff they have the same properties, then we say that the equivalence is characterised by the logic. Following [HM85], if $P \approx Q$ then P and Q satisfy the same formulas $M \in \mathcal{M}$ and the other way round but only for weak-image-finite $\longrightarrow$:

$$P \approx Q \ \text{ iff } \ \forall_{M \in \mathcal{M}} \ P \models M \Leftrightarrow Q \models M \tag{16}$$

The aim of this section is to provide similar statements for fault-tolerant and conditional fault-tolerant extensions of $\approx$. Consider the new relation $\models\!\!\models \subseteq \mathcal{P} \times \mathcal{M}$ which is defined like $\models$ except that the transitions $\longmapsto$ are now used to define the semantics of formulas $\langle\beta\rangle M$:

$$Q \models\!\!\models \langle\beta\rangle M \ \text{ iff } \ \exists_{Q',s} \ \ Q \stackrel{s}{\longmapsto} Q' \ \wedge \ \widehat{s} = \beta \ \wedge \ Q' \models\!\!\models M \tag{17}$$

Applying $\models$ for the high-level process and $\models\!\!\models$ for the low-level one, we can show that for weak-image-finite relations $\longmapsto$, must-bisimilarity $\sqsubseteq$ is characterised by the Hennessy-Milner logic:

Proposition 1. For weak-image-finite relation $\longmapsto$ we have:

$$P \sqsubseteq Q \ \text{ iff } \ \forall_{M \in \mathcal{M}} \ \ P \models M \ \Leftrightarrow \ Q \models\!\!\models M$$

Proof. ($\Rightarrow$) By induction on the structure of M. ($\Leftarrow$) We show that the relation $B =_{def} \{(P, Q) \mid \forall_{M \in \mathcal{M}} \ \ P \models M \ \Leftrightarrow \ Q \models\!\!\models M\}$ is a must-bisimulation. For details see Appendix A. □

Because $P \boxtimes Q$ implies $P \sqsubset Q$, we also have $P \models M$ iff $Q \models\!\!\!\models M$ for any $P \boxtimes Q$ and $M \in \mathcal{M}$. The inverse however does not hold, as demonstrated by P and S in Figure 2 which have the same properties (P with respect to $\models$ and S according to $\models\!\!\!\models$) but still $P \not\boxtimes S$.

Example 3. Consider processes in Figure 2 and two formulas for $m \geq 0$:

- $M = [b]^{2m}\langle a\rangle true$ which asserts that the first action a is always possible after an even number of b's;
- $N = [a]\langle b\rangle^{2m+1}\langle a\rangle true$ where the second occurrence of a may be possible after an odd number of b's.

We have $P \models M \wedge N$ and thus $Q \models M \wedge N$ and $S \models M \wedge N$ because of $P \approx Q \approx S$. We also have $P \sqsubset R$ what gives $R \models\!\!\!\models M \wedge N$ and $P \sqsubset S$ what results in $S \models\!\!\!\models M \wedge N$. Finally $Q \models\!\!\!\models \neg M \wedge N$ and $R \models M \wedge \neg N$. □

Consider $n \in \mathbb{N} \cup \{\infty\}$ which specifies the maximal number of times transitions $\dashrightarrow$ can occur successively. The logical characterisation of $\boxtimes_n$ involves conditional satisfaction relation $\models\!\!\!\models_n$ which is defined in terms of the family of relations $\models\!\!\!\models^i_n$, indexed by $[0,n]$. Consider $i \in [0,n]$. For all formulas except $\langle\beta\rangle M$, $\models\!\!\!\models^i_n$ is defined like $\models$. For $\langle\beta\rangle M$ we have:

$$Q \models\!\!\!\models^i_n \langle\beta\rangle M \;\text{ iff }\; \exists_{Q',s,j} \;\; Q \stackrel{s}{\longmapsto}{}^i_j Q' \;\wedge\; \widehat{s} = \beta \;\wedge\; Q' \models\!\!\!\models^j_n M \qquad (18)$$

Thus $Q \models\!\!\!\models^i_n M$ if Q satisfies M, performing transitions $\longrightarrow$ and $\dashrightarrow$ but the last no more than n times in a row and no more than $n-i$ of them initially. Finally, we define $P \models\!\!\!\models_n M$ iff $P \models\!\!\!\models^0_n M$. Given such $\models\!\!\!\models_n$ and provided that $\longmapsto$ is weak-image-finite, we can prove the following characterisation theorem:

Proposition 2. For weak-image-finite relation $\longmapsto$ we have:

$$P \boxtimes_n Q \;\text{ iff }\; \forall_{M\in\mathcal{M}} \;\; P \models M \;\Leftrightarrow\; Q \models\!\!\!\models_n M$$

Proof. ($\Rightarrow$) By induction on the structure of M. ($\Leftarrow$) We show that the family $\{B_i\}^n_{i=0}$ of relations $B_i =_{def} \{(P,Q) \mid \forall_{M\in\mathcal{M}} \;\; P \models M \;\Leftrightarrow\; Q \models\!\!\!\models^i_n M\}$ is a conditional must-bisimulation. For details see Appendix A. □

As a result, because $P \boxtimes_n Q$ implies $P \sqsubset_n Q$, we have $P \models M$ iff $Q \models\!\!\!\models_n M$ for all $P \boxtimes_n Q$ and $M \in \mathcal{M}$.

Example 4. Consider processes in Figure 4, $m \geq 0$ and the following formulas:

- $M = \langle a\rangle true \wedge [a][a]false$
  where action a is possible but then it cannot be followed by another a;
- $N = \langle b\rangle true \wedge [b][a]false$
  where action b is possible but not followed by a either.

We have $P \models [a][b]^{2m+1}(M \wedge N)$ and thus $Q \models [a][b]^{2m+1}(M \wedge N)$ because of $P \approx Q$. For Q we have $Q \models\!\!\!\models [a][b]^{2m+1}N$ and $Q \models\!\!\!\models \neg[a][b]^{2m+1}M$, however $Q \models\!\!\!\models_1 [a][b]^{2m+1}(M \wedge N)$ because then transition $Q''' \stackrel{a}{\dashrightarrow} Q$, which enables two subsequent actions a, cannot be chosen. For R we have $R \models\!\!\!\models \neg[a][b]^{2m+1}M$ and $R \models\!\!\!\models \neg[a][b]^{2m+1}N$, however $R \models\!\!\!\models_2 [a][b]^{2m+1}(M \wedge N)$ because then transition $R''' \stackrel{\tau}{\dashrightarrow} R''$ cannot be chosen. □

5 Language of Processes

The structure of a process P has been ignored so far, it was defined as an element of the abstract set $\mathcal{P}$. The more complex is the behaviour of P however, the greater is the need to treat P structurally. In this section, following [Mil89], we define $\mathcal{P}$ as the language of processes which is given the structured operational semantics [Plo81] in terms of the labelled transition system $(\mathcal{P}, \mathcal{A}, \longrightarrow)$.

The syntax of $\mathcal{P}$ is based on two sets of symbols, $\mathcal{A}$ of actions and $\mathcal{X}$ of process identifiers, and involves two syntactic categories, $\mathcal{E}$ of process expressions and $\mathcal{D}$ of declarations. Let $X, Y \in \mathcal{X}$, $L \subseteq \mathcal{L}$ and f be a function over $\mathcal{A}$ such that $f(\tau) = \tau$, $f(a) \neq \tau$ and $f(\overline{a}) = \overline{f(a)}$. $\mathcal{E}$ is defined by the grammar:

$$E \quad ::= \quad X \mid \mathbf{0} \mid \alpha.E \mid E + E \mid E|E \mid E \backslash L \mid E[f] \tag{19}$$

Informally, $\mathbf{0}$ is unable to take any action and $\alpha.E$ performs α and then behaves like E. The operator $+$ represents summation, $|$ parallel composition, $\backslash$ restriction and $[\,]$ renaming. One derived operator is $E \frown F$ where E and F proceed in parallel with actions *out* of E and *in* of F 'joined' and restricted, $E \frown F =_{def} (E[mid/out]|F[mid/in]) \backslash \{mid\}$ where *mid* is not used by E or F.

We use $\mathcal{X}(E)$ for the set of all identifiers in E and $E\{F/X\}$ for the process expression E where all identifiers X are replaced by F. In order to interpret $X \in \mathcal{X}(E)$, we use declarations of the form $X \mathrel{\widehat{=}} F$. If also $X \in \mathcal{X}(F)$ then such X is defined by recursion. Given $X \mathrel{\widehat{=}} F$ and $Y \mathrel{\widehat{=}} G$ where $X \in \mathcal{X}(G)$ and $Y \in \mathcal{X}(F)$, such X and Y are defined by the mutual recursion. In the sequel we will often need to manipulate declarations for mutually recursive identifiers. Then, it will be helpful to use a simple language $\mathcal{D}$ for specifying collections of such declarations. $\mathcal{D}$, ranged over by Δ and ∇, is defined by the grammar:

$$\Delta \quad ::= \quad [\,] \mid \Delta[X \mathrel{\widehat{=}} E] \mid \alpha \odot \Delta \mid \Delta \oplus \nabla \tag{20}$$

Informally, $[\,]$ is an empty declaration and $\Delta[X \mathrel{\widehat{=}} E]$ declares X as E and other identifiers as in Δ. Moreover, $\alpha \odot \Delta$ and $\Delta \oplus \nabla$ perform respective operations (α-prefix and summation) on the right sides of all the corresponding declarations in Δ and ∇. Formally, $\Delta \in \mathcal{D}$ is assigned a partial function $[\![\Delta]\!]$ from $\mathcal{X}$ to $\mathcal{E}$ which is defined in Figure 5 by induction on the structure of Δ. We use $dom(\Delta)$ as the domain of $[\![\Delta]\!]$ ($dom([\,]) =_{def} \emptyset$) and $ran(\Delta)$ as its range.

$$[\![\alpha \odot \Delta]\!](X) =_{def} \alpha.[\![\Delta]\!](X) \quad \text{if } X \in dom(\Delta)$$

$$[\![\Delta[Y \mathrel{\widehat{=}} E]]\!](X) =_{def} \begin{cases} E & \text{if } X = Y \\ [\![\Delta]\!](X) & \text{if } X \neq Y,\ X \in dom(\Delta) \end{cases}$$

$$[\![\Delta \oplus \nabla]\!](X) =_{def} \begin{cases} [\![\Delta]\!](X) & \text{if } X \in dom(\Delta) - dom(\nabla) \\ [\![\Delta]\!](X) + [\![\nabla]\!](X) & \text{if } X \in dom(\Delta) \cap dom(\nabla) \\ [\![\nabla]\!](X) & \text{if } X \in dom(\nabla) - dom(\Delta) \end{cases}$$

Fig. 5. Denotational semantics of declarations $\mathcal{D}$.

We abbreviate $[][X \mathrel{\widehat{=}} E][Y \mathrel{\widehat{=}} F]$ as $[X \mathrel{\widehat{=}} E, Y \mathrel{\widehat{=}} F]$ and write $[X \mathrel{\widehat{=}} E \mid p]$ for all declarations $X \mathrel{\widehat{=}} E$ such that the predicate p holds. Δ is said to be closed if all identifiers in the right side expressions of Δ are declared in Δ:

$$\bigcup_{F \in ran(\Delta)} \mathcal{X}(F) \subseteq dom(\Delta) \tag{21}$$

A process $P \in \mathcal{P}$ is finally the pair $\langle E, \Delta \rangle$ of the process expression E and the closed declaration Δ for all identifiers of E, $\mathcal{X}(E) \subseteq dom(\Delta)$. We write $\langle E, \Delta \rangle \equiv \langle F, \nabla \rangle$ if $\langle E, \Delta \rangle$ and $\langle F, \nabla \rangle$ are identical.

The semantics of $\langle E, \Delta \rangle \in \mathcal{P}$ is defined in terms of the labelled transition system $(\mathcal{P}, \mathcal{A}, \longrightarrow)$ by induction on the structure of E. If $E = X$ then transitions of $\langle X, \Delta \rangle$ involve the semantics of Δ, they are inferred from the transitions of $\langle [\![\Delta]\!](X), \Delta \rangle$. Following [Mil89], transition relation $\longrightarrow$ is the least set defined by inference rules in Figure 6.

$$\frac{}{\langle \alpha.E, \Delta \rangle \xrightarrow{\alpha} \langle E, \Delta \rangle} \qquad \frac{\langle E, \Delta \rangle \xrightarrow{\alpha} \langle E', \Delta \rangle}{\langle E+F, \Delta \rangle \xrightarrow{\alpha} \langle E', \Delta \rangle} \qquad \frac{\langle F, \Delta \rangle \xrightarrow{\alpha} \langle F', \Delta \rangle}{\langle E+F, \Delta \rangle \xrightarrow{\alpha} \langle F', \Delta \rangle}$$

$$\frac{\langle E, \Delta \rangle \xrightarrow{\alpha} \langle E', \Delta \rangle}{\langle E|F, \Delta \rangle \xrightarrow{\alpha} \langle E'|F, \Delta \rangle} \qquad \frac{\langle F, \Delta \rangle \xrightarrow{\alpha} \langle F', \Delta \rangle}{\langle E|F, \Delta \rangle \xrightarrow{\alpha} \langle E|F', \Delta \rangle}$$

$$\frac{\langle E, \Delta \rangle \xrightarrow{a} \langle E', \Delta \rangle \quad \langle F, \Delta \rangle \xrightarrow{\overline{a}} \langle F', \Delta \rangle}{\langle E|F, \Delta \rangle \xrightarrow{\tau} \langle E'|F', \Delta \rangle}$$

$$\frac{\langle E, \Delta \rangle \xrightarrow{\alpha} \langle E', \Delta \rangle}{\langle E \setminus L, \Delta \rangle \xrightarrow{\alpha} \langle E' \setminus L, \Delta \rangle}, \quad \alpha, \overline{\alpha} \notin L \qquad \frac{\langle E, \Delta \rangle \xrightarrow{\alpha} \langle E', \Delta \rangle}{\langle E[f], \Delta \rangle \xrightarrow{f(\alpha)} \langle E'[f], \Delta \rangle}$$

$$\frac{\langle [\![\Delta]\!](X), \Delta \rangle \xrightarrow{\alpha} \langle E, \Delta \rangle}{\langle X, \Delta \rangle \xrightarrow{\alpha} \langle E, \Delta \rangle}, \quad X \in dom(\Delta)$$

Fig. 6. Operational semantics of processes $\mathcal{P}$

Example 5. Consider $n \in \mathbb{N} \cup \{\infty\}$ and the process R_n which performs actions a and b; the first at any time; the second no more than n times in a row. We have $R_n =_{def} \langle Y_0, \Delta_n \oplus \nabla_n \rangle$ where:

$$\begin{aligned} \Delta_n &=_{def} [Y_i \mathrel{\widehat{=}} a.Y_0 \quad \mid 0 \le i \le n] \\ \nabla_n &=_{def} [Y_i \mathrel{\widehat{=}} b.Y_{i+1} \mid 0 \le i < n] \end{aligned}$$

□

We compose processes by composing their expressions, using the process combinators (19). For binary operators $+$ and $|$ we assume that the component processes have disjoint sets of identifiers. If $dom(\Delta) \cap dom(\nabla) = \emptyset$ then:

$$\begin{aligned} \alpha.\langle E, \Delta \rangle &=_{def} \langle \alpha.E, \Delta \rangle \\ \langle E, \Delta \rangle \setminus L &=_{def} \langle E \setminus L, \Delta \rangle \\ \langle E, \Delta \rangle[f] &=_{def} \langle E[f], \Delta \rangle \\ \langle E, \Delta \rangle + \langle F, \nabla \rangle &=_{def} \langle E + F, \Delta \oplus \nabla \rangle \\ \langle E, \Delta \rangle \mid \langle F, \nabla \rangle &=_{def} \langle E \mid F, \Delta \oplus \nabla \rangle \end{aligned} \tag{22}$$

In the language defined so far, processes interact by synchronising on complementary actions a and $\overline{a}$. There is no directionality or value which passes between them. For pragmatic reasons, we also need a value-passing language for the set V of values (we assume, for simplicity, that V is finite). To this end we introduce value constants (like ε), value variables (like x and s), value and boolean expressions (like e and p respectively), built using constants, variables and any function symbols we need. The last include $\sharp s$ as the length of the sequence s, s_0 its first element, s' all but the first element and $s : x$ as the sequence s with value x appended. We also introduce parameters into process identifiers: $X(e1, .., en)$ for X of arity n. Then we extend the basic language by input and output prefixes $a(x).E$, $\overline{a}(e).E$ and conditionals **if** p **then** E **else** F. For their translation into the basic language see [Mil89].

Example 6. Consider a buffer Buf_m of capacity $m > 0$, which receives (by action in) and subsequently transmits (by action $\overline{out}$) all values unchanged, in the same order and with at most m of them received but not sent. We have:

$$\begin{array}{lll} Buf_m & =_{def} & \langle X(\varepsilon), \Delta \oplus \nabla \rangle \\ \text{where} & \Delta =_{def} & [X(s) \mathrel{\widehat{=}} in(x).X(s:x) \mid 0 \leq \sharp s < m] \\ & \nabla =_{def} & [X(s) \mathrel{\widehat{=}} \overline{out}(s_0).X(s') \mid 0 < \sharp s \leq m] \end{array}$$

□

6 Language of Faults

Although a fault is modelled by a set of transitions, using this set directly is not the most convenient way of specifying faults in practice, especially when the abnormal behaviour we want to describe is complex. The purpose of this section is to define the language where faults can be specified and combined.

The idea is to use process identifiers as 'states' which can be affected by faults. Consider a process $\langle X, \Delta \rangle$. Transitions of $\langle X, \Delta \rangle$ can be only inferred from the transitions of $\langle [\![\Delta]\!](X), \Delta \rangle$ where $[\![\Delta]\!](X)$ is the process expression assigned to X by Δ. In order to specify faults, we will use an alternative, 'faulty' declaration $\Psi \in \mathcal{D}$. Suppose that $X \in dom(\Psi)$. Then X is assigned yet another expression $[\![\Psi]\!](X)$ which determines abnormal transitions of $\langle X, \Delta \rangle$, following transitions of $\langle [\![\Psi]\!](X), \Delta \rangle$ and denoted by $\stackrel{}{\underset{\Psi}{\longmapsto}}$:

$$\frac{\langle [\![\Psi]\!](X), \Delta \rangle \stackrel{\alpha}{\underset{\Psi}{\longmapsto}} \langle E, \Delta \rangle}{\langle X, \Delta \rangle \stackrel{\alpha}{\underset{\Psi}{\longmapsto}} \langle E, \Delta \rangle} \tag{23}$$

Transition relation $\stackrel{}{\underset{\Psi}{\longmapsto}}$ is defined as the least set which satisfies inference rules in Figure 7 and used to denote Ψ-affected semantics of $\mathcal{P}$. We also define:

$$Q \stackrel{\alpha}{\underset{\Psi}{\cdots\!\!\rightarrow}} Q' \text{ iff } Q \stackrel{\alpha}{\underset{\Psi}{\longmapsto}} Q' \text{ and } Q \not\stackrel{\alpha}{\longmapsto} Q' \tag{24}$$

and relations $\stackrel{i}{\underset{\Psi}{\longmapsto}}_j \subseteq \mathcal{P} \times \mathcal{A}^* \times \mathcal{P}$ (10), given $i, j \in [0, n]$ and faulty transitions $\stackrel{}{\underset{\Psi}{\cdots\!\!\rightarrow}}$, specified by Ψ.

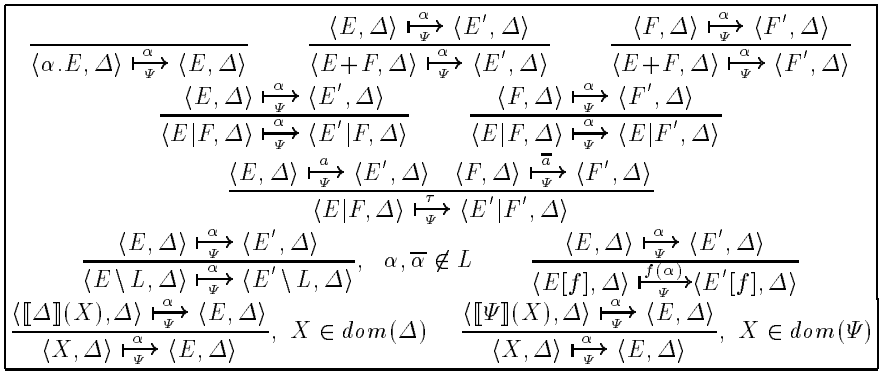

Fig. 7. Operational semantics of processes $\mathcal{P}$ affected by the fault Ψ.

Ψ is not assumed to be closed. However, in order to ensure that $\overset{}{\underset{\Psi}{\longmapsto}}$ does not lead from the well-defined process (where all identifiers are declared) to the ill-defined one, we assume that all process identifiers in the right-side expressions of Ψ are declared by Δ:

$$\bigcup_{F \in ran(\Psi)} \mathcal{X}(F) \subseteq dom(\Delta) \tag{25}$$

We use $\mathcal{P}_\Psi \subseteq \mathcal{P}$ for the set of such $\langle E, \Delta \rangle$ and assume that $\overset{}{\underset{\Psi}{\longmapsto}} \subseteq \mathcal{P}_\Psi \times \mathcal{A} \times \mathcal{P}_\Psi$ and $\overset{i}{\underset{\Psi}{\longmapsto}}_j \subseteq \mathcal{P}_\Psi \times \mathcal{A}^* \times \mathcal{P}_\Psi$.

Example 7. Consider the following declarations which specify various communication faults of the bounded buffer Buf_m, creation (Ψe), corruption (Ψc), omission (Ψo), replication (Ψr) and permutation (Ψp) of messages:

$$\begin{array}{ll}
\Psi e =_{def} [X(s) \mathrel{\widehat{=}} \tau.\overline{out}(\surd).X(s) & \mid 0 \leq \sharp s \leq m] \\
\Psi c =_{def} [X(s) \mathrel{\widehat{=}} \tau.\overline{out}(\surd).X(s') & \mid 0 < \sharp s \leq m] \\
\Psi o =_{def} [X(s) \mathrel{\widehat{=}} \tau.X(s') & \mid 0 < \sharp s \leq m] \\
\Psi r =_{def} [X(s) \mathrel{\widehat{=}} \tau.\overline{out}(s_0).X(s) & \mid 0 < \sharp s \leq m] \\
\Psi p =_{def} [X(s) \mathrel{\widehat{=}} \tau.\overline{out}((s')_0).X(s_0 : s'') & \mid 1 < \sharp s \leq m]
\end{array}$$

We use $\surd$ to denote messages which has been corrupted or created. This is a way to abstract from their particular value which is immaterial. Also, we assume that when permuted, only one message is delayed. In order to specify more complex faults, e.g. simultaneous creation, omission and permutation of messages, we can use the summation $\Psi e \oplus \Psi o \oplus \Psi p$ of declarations:

$$\begin{array}{l}
[X(s) \mathrel{\widehat{=}} \tau.\overline{out}(\surd).X(s) \mid \sharp s = 0] \\
[X(s) \mathrel{\widehat{=}} \tau.\overline{out}(\surd).X(s) + \tau.X(s') \mid \sharp s = 1] \\
[X(s) \mathrel{\widehat{=}} \tau.\overline{out}(\surd).X(s) + \tau.X(s') + \tau.\overline{out}((s')_0).X(s_0 : s'') \mid 1 < \sharp s \leq m] \quad \square
\end{array}$$

7 Fault Transformation of Processes

The primary effect of faults is that a process no longer behaves according to the normal transition relation $\longrightarrow$. In addition to $\longrightarrow$, it can also perform transitions $\stackrel{\cdots}{\Psi}\!\!\rightarrow$, specified by $\Psi \in \mathcal{D}$. This is a direct, semantic method to represent effects of faults on the behaviour of the process. In this section we present an alternative, syntactic method. The idea is to capture the effect of faults, specified by Ψ, by the process transformation $\mathcal{T}(\,\cdot\,, \Psi)$ where for any $Q \in \mathcal{P}_\Psi$, its behaviour in the Ψ-affected environment is 'the same' as the behaviour of $\mathcal{T}(Q, \Psi)$ in the environment which is free of faults [Liu91, LJ91]. We show that $\mathcal{T}(Q, \Psi)$ yields the binary relation on $\mathcal{P} \times \mathcal{P}_\Psi$ which coincides with $\sqsubset$ and the satisfaction relation which agrees with $\models$. In the conditional case we provide a transformation $\widetilde{\mathcal{T}}(\,\cdot\,, \Psi, n)$ which is shown to coincide with $\sqsubset_n$ and $\models_n$.

7.1 Fault-Tolerance

Consider $\Psi \in \mathcal{D}$ which specifies transitions $\stackrel{\cdots}{\Psi}\!\!\rightarrow$ and a process $\langle E, \Delta\rangle \in \mathcal{P}_\Psi$ with the well-defined, Ψ-affected semantics $\stackrel{}{\mapsto_\Psi}$. If $\langle E, \Delta\rangle$ has no identifiers in common with Ψ then $\langle E, \Delta\rangle$ is not affected by the transitions $\stackrel{\cdots}{\Psi}\!\!\rightarrow$. Suppose that $X \in dom(\Delta) \cap dom(\Psi)$ and that transitions of $\langle E, \Delta\rangle$ can be inferred from $\langle X, \Delta\rangle \stackrel{\alpha}{\mapsto_\Psi} \langle E', \Delta\rangle$. We have either

$$\frac{\langle [\![\Delta]\!](X), \Delta\rangle \stackrel{\alpha}{\mapsto_\Psi} \langle E', \Delta\rangle}{\langle X, \Delta\rangle \stackrel{\alpha}{\mapsto_\Psi} \langle E', \Delta\rangle} \quad \text{or} \quad \frac{\langle [\![\Psi]\!](X), \Delta\rangle \stackrel{\alpha}{\mapsto_\Psi} \langle E', \Delta\rangle}{\langle X, \Delta\rangle \stackrel{\alpha}{\mapsto_\Psi} \langle E', \Delta\rangle} \tag{26}$$

where the first transition is normal (it uses Δ to interpret X) and the second is faulty (it uses Ψ). The ability of $\langle E, \Delta\rangle$ to perform the second transition (with respect to $\longrightarrow$) can be syntactically represented by summation, by redefining its process identifier X as $[\![\Delta]\!](X) + [\![\Psi]\!](X)$. To represent the capacity for all transitions in $\stackrel{\cdots}{\Psi}\!\!\rightarrow$, such a summation must be performed for all identifiers $X \in dom(\Delta) \cap dom(\Psi)$. This leads to the following transformation:

$$\mathcal{T}(\langle E, \Delta\rangle, \Psi) =_{def} \langle E, \Delta \oplus \Psi\rangle \tag{27}$$

We would like to show that $\mathcal{T}(\langle E, \Delta\rangle, \Psi)$ captures the effect of faults, specified by Ψ on $\langle E, \Delta\rangle$. Observe first that declarations (and thus transformation $\mathcal{T}$) 'persist' through the transitions $\stackrel{}{\mapsto_\Psi}$ and $\longrightarrow$:

$$\begin{array}{l} \text{If } \langle E, \Delta\rangle \stackrel{\alpha}{\mapsto_\Psi} Q' \text{ then } Q' \equiv \langle E', \Delta\rangle \text{ for some } E'. \\ \text{If } \mathcal{T}(\langle E, \Delta\rangle, \Psi) \stackrel{\alpha}{\longrightarrow} R' \text{ then } R' \equiv \mathcal{T}(\langle E', \Delta\rangle, \Psi) \text{ for some } E'. \end{array} \tag{28}$$

Both statements can be shown by transitional induction. Transitional induction is also employed to prove the following lemma (for details see Appendix B):

Lemma 3.

If $\quad Q \in \mathcal{P}_\Psi$
then $Q \stackrel{\alpha}{\mapsto_\Psi} Q'$ iff $\mathcal{T}(Q, \Psi) \stackrel{\alpha}{\longrightarrow} \mathcal{T}(Q', \Psi)$.

Transformation $\mathcal{T}(\ \cdot\ , \Psi)$ induces the satisfaction relation on $\mathcal{P}_\Psi \times \mathcal{M}$ which holds between a process $Q \in \mathcal{P}_\Psi$ and a formula $M \in \mathcal{M}$ iff $\mathcal{T}(Q, \Psi) \models M$. Applying Lemma 3, we can show that this relation coincides with $\Vdash$:

Proposition 4.

If $M \in \mathcal{M}$, $Q \in \mathcal{P}_\Psi$ and
transitions $\cdots\!\!\rightarrow$ are specified by Ψ
then $Q \Vdash M$ iff $\mathcal{T}(Q, \Psi) \models M$

Proof. By induction on the structure of formulas M, applying Lemma 3. For details see Appendix B. □

Transformation $\mathcal{T}(\ \cdot\ , \Psi)$ also induces the binary relation on $\mathcal{P} \times \mathcal{P}_\Psi$ which holds between $P \in \mathcal{P}$ and $Q \in \mathcal{P}_\Psi$ iff $P \approx \mathcal{T}(Q, \Psi)$. The following proposition asserts that this relation coincides with must-bisimilarity $\sqsubset$:

Proposition 5.

If $P \in \mathcal{P}$, $Q \in \mathcal{P}_\Psi$ and
transitions $\cdots\!\!\rightarrow$ are specified by Ψ
then $P \sqsubset Q$ iff $P \approx \mathcal{T}(Q, \Psi)$

Proof. It is easy to see that for weak-image-finite $\stackrel{\Psi}{\longmapsto}$, this statement follows from the characterisation theorem (16) and Propositions 1 and 4. For any $\stackrel{\Psi}{\longmapsto}$, it follows from the fact that:

$$B \subseteq \mathcal{P} \times \mathcal{P}_\Psi \text{ is a must-bisimulation iff } \mathfrak{T}(B, \Psi) \text{ is a bisimulation} \qquad (29)$$

where $\mathfrak{T}(B, \Psi) =_{def} \{(P, \mathcal{T}(Q, \Psi)) \mid (P, Q) \in B\}$. For details see Appendix B. □

Thus for must-bisimilarity $\sqsubset$, we have the equivalence diagram (1) of all three approaches to verify fault-tolerance, given $\stackrel{\Psi}{\cdots\!\!\rightarrow}$ such that $\stackrel{\Psi}{\longmapsto}$ is weak-image-finite. For $\stackrel{\Psi}{\longmapsto}$ which is not weak-image-finite, we cannot guarantee that

$$P \models M \text{ iff } Q \Vdash M \text{ for all } M \in \mathcal{M} \text{ implies } P \sqsubset Q.$$

Moreover, for may-bisimilarity $\sqsupset$, we cannot guarantee that $P \approx \mathcal{T}(Q, \Psi)$ implies $P \sqsupset Q$. However, applying equivalences (7) and (29) it is easy to see that $B \subseteq \mathcal{P} \times \mathcal{P}_\Psi$ is a may-bisimulation iff it is a bisimulation together with $\mathfrak{T}(B, \Psi)$. As a result, given transitions $\stackrel{\Psi}{\cdots\!\!\rightarrow}$, we have the following statements:

$P \approx Q$ iff $(P, Q) \in B$ for B which is a bisimulation
$P \sqsubset Q$ iff $(P, Q) \in B$ for B such that $\mathfrak{T}(B, \Psi)$ is a bisimulation.
$P \sqsupset Q$ iff $(P, Q) \in B$ for B which is a bisimulation together with $\mathfrak{T}(B, \Psi)$.

Example 8. Consider $m, w > 0$ and the task to ensure a reliable communication, specified by the bounded buffer Buf_w, over a medium of capacity m which omits and replicates messages. To this end, we will use a version of the sliding window protocol with the window size w. The protocol consists of two processes, the sender So and the receiver Ro. The first transmits all messages with their sequence numbers i modulo $w+1$, such that at most w messages are sent without being acknowledged. Suppose, for simplicity, that acknowledgements take place by synchronising So and Ro on the action ack. We use s for the sequence of messages sent but not acknowledged ($\sharp s \leq w$) and repeatedly retransmit s_0. For all arithmetic operations taken modulo $w+1$ we have:

$$\begin{array}{lll}
So =_{def} \langle Zs(0,\varepsilon), & [Zs(i,s) & \widehat{=}\ in(x).Zs(i,s,x) \quad | \ 0 \leq \sharp s < w] \oplus \\
& [Zs(i,s) & \widehat{=}\ ack.Zs(i,s') \quad | \ 0 < \sharp s \leq w] \oplus \\
& [Zs(i,s) & \widehat{=}\ \overline{out}(i - \sharp s, s_0).Zs(i,s) \quad | \ 0 < \sharp s \leq w] \\
& [Zs(i,s,x) & \widehat{=}\ \overline{out}(i,x).Zs(i+1,s:x) \quad | \ 0 \leq \sharp s < w] \oplus \\
& [Zs(i,s,x) & \widehat{=}\ ack.Zs(i,s',x) \quad | \ 0 < \sharp s < w] \oplus \\
& [Zs(i,s,x) & \widehat{=}\ \overline{out}(i-1,s_0).Zs(i,s,x) \quad | \ 0 < \sharp s < w] \rangle \\
Ro =_{def} \langle Zr(0), & [Zr(i) & \widehat{=}\ in(j,x).\ \textbf{if}\ i = j \\
& & \quad \textbf{then}\ \overline{out}(x).\overline{ack}.Zr(i+1) \\
& & \quad \textbf{else}\ Zr(i)] \rangle
\end{array}$$

Given such So and Ro, we can prove that:

$$Buf_w \approx (So \frown \mathcal{T}(Buf_m, \Psi o \oplus \Psi r) \frown Ro) \setminus \{ack\}$$

and consequently $Buf_w \boxplus (So \frown Buf_m \frown Ro) \setminus \{ack\}$, following Proposition 5 and the fact that processes So and Ro are not affected by $\Psi o \oplus \Psi r$ (they have disjoint sets of identifiers). However, the above statement does not depend on the transitions specified by $\Psi o \oplus \Psi r$ and we can find a bisimulation B such that $(Buf_w, (So \frown Buf_m \frown Ro) \setminus \{ack\}) \in B$ and $\mathfrak{T}(B, \Psi o \oplus \Psi r)$ is a bisimulation. Thus we have $Buf_w \boxplus (So \frown Buf_m \frown Ro) \setminus \{ack\}$. □

7.2 Conditional Fault-Tolerance

Consider declaration Ψ and a process $\langle E, \Delta \rangle \in \mathcal{P}_\Psi$. Suppose that we want to verify $\langle E, \Delta \rangle$ in the presence of transitions $\overset{\cdots}{\Psi}\!\!\rightarrow$ and under assumption n about their quantity. To this end, like before, we will use process transformations. The idea is to use a family $\{X_i\}_{i=0}^n$ of identifiers for each identifier X of Δ or Ψ. Consider $i \in [0,n]$ and the following transformations $\widehat{\mathcal{T}}_i(\cdot, n)$ and $\widetilde{\mathcal{T}}_i(\cdot, \Psi, n)$:

$$\begin{array}{l}
\widehat{\mathcal{T}}_i(\langle E, \Delta \rangle, n) \quad =_{def} \langle E_i, \Delta_n \rangle \quad \text{and} \\
\widetilde{\mathcal{T}}_i(\langle E, \Delta \rangle, \Psi, n) =_{def} \mathcal{T}(\widehat{\mathcal{T}}_i(\langle E, \Delta \rangle, n), \Psi_n) \\
\text{where} \quad E_i =_{def} E\{X_i/X \mid X \in \mathcal{X}(E)\} \\
\qquad \Delta_n =_{def} [X_i \widehat{=} F\{Y_0/Y \quad \mid Y \in \mathcal{X}(F)\} \mid [\![\Delta]\!](X) = F \wedge 0 \leq i \leq n] \\
\qquad \Psi_n =_{def} [X_i \widehat{=} F\{Y_{i+1}/Y \mid Y \in \mathcal{X}(F)\} \mid [\![\Psi]\!](X) = F \ \wedge 0 \leq i < n]
\end{array}$$

Thus $\widetilde{\mathcal{T}}_i$ is defined in terms of transformations $\mathcal{T}$ and $\widehat{\mathcal{T}}_i$. Also, Δ_n is obtained from Δ by replacing each declaration $X \mathrel{\widehat{=}} F$ with the family of declarations $X_i \mathrel{\widehat{=}} F\{Y_0/Y \mid Y \in \mathcal{X}(F)\}$, for all $i \in [0,n]$, and similarly for Ψ_n but only given $i \in [0,n)$. Finally, we define $\widehat{\mathcal{T}}(\cdot,n)$ and $\widetilde{\mathcal{T}}(\cdot,\Psi,n)$ taking $i=0$:

$$\begin{array}{ll} \widehat{\mathcal{T}}(\langle E,\Delta\rangle,n) & =_{def}\ \widehat{\mathcal{T}}_0(\langle E,\Delta\rangle,n) \\ \widetilde{\mathcal{T}}(\langle E,\Delta\rangle,\Psi,n) & =_{def}\ \widetilde{\mathcal{T}}_0(\langle E,\Delta\rangle,\Psi,n) \end{array} \tag{30}$$

Recall the family $\{\overset{i}{\underset{\Psi}{\longmapsto}}_j\}_{i,j=0}^{n}$ of relations which denotes the effect of transitions $\overset{}{\underset{\Psi}{\cdots\!\!\rightarrow}}$ on the semantics of $\mathcal{P}_\Psi$, under assumption n about their quantity. There are two problems to obtain the same effect using transformations:

1. Consider $X \in dom(\Delta) \cap dom(\Psi)$ and $\langle X,\Delta\rangle \overset{\alpha}{\underset{\Psi}{\longmapsto}} \langle E,\Delta\rangle$ which can be inferred from $\langle [\![\Delta]\!](X),\Delta\rangle \overset{\alpha}{\underset{\Psi}{\longmapsto}} \langle E,\Delta\rangle$ or $\langle [\![\Psi]\!](X),\Delta\rangle \overset{\alpha}{\underset{\Psi}{\longmapsto}} \langle E,\Delta\rangle$. The problem appears when $\langle X,\Delta\rangle \overset{\alpha}{\underset{\Psi}{\longmapsto}} \langle E,\Delta\rangle$ can be inferred from both of them. Then it is always regarded as 'normal' by $\overset{i}{\underset{\Psi}{\longmapsto}}_j$ but not always by $\widetilde{\mathcal{T}}(\cdot,\Psi,n)$. If no such E and α exists then we say that Ψ *has the proper effect on* Δ.
2. The second problem is that in case of $\widetilde{\mathcal{T}}(\cdot,\Psi,n)$ (but not $\overset{i}{\underset{\Psi}{\longmapsto}}_j$), some transitions do not 'update' the index i. Suppose that $\Psi =_{def} [X \mathrel{\widehat{=}} \tau.a.X]$ and $\Delta =_{def} [X \mathrel{\widehat{=}} b.X]$. Then we have: $\langle X,\Delta\rangle \overset{\tau}{\underset{\Psi}{\cdots\!\!\rightarrow}} \langle a.X,\Delta\rangle \overset{a}{\longrightarrow} \langle X,\Delta\rangle$ and thus $\langle X,\Delta\rangle \overset{\tau\,a}{\underset{\Psi}{\longmapsto}}{}^{0}_{0} \langle X,\Delta\rangle$, however $\widetilde{\mathcal{T}}_0(\langle X,\Delta\rangle,\Psi,n) \overset{\tau\,a}{\longrightarrow} \widetilde{\mathcal{T}}_1(\langle X,\Delta\rangle,\Psi,n)$. We can solve this problem assuming that all expressions involved are *linear* i.e. they are of the form $\sum_{i=1}^{k} \alpha_i.X_i$ (Δ is linear if all $F \in ran(\Delta)$ are linear).

Under assumption about the linear form of Δ and Ψ, we can easily show that:

$$\begin{array}{l} \text{if } \langle X,\Delta\rangle \overset{\alpha}{\underset{\Psi}{\longmapsto}} Q' \text{ then } Q' \equiv \langle Y,\Delta\rangle \text{ and} \\ \text{if } \widetilde{\mathcal{T}}_i(\langle X,\Delta\rangle,\Psi,n) \overset{\alpha}{\longrightarrow} R' \text{ then } R' \equiv \widetilde{\mathcal{T}}_j(\langle Y,\Delta\rangle,\Psi,n) \end{array} \tag{31}$$

where $Y \in \mathcal{X}$ and $j \in [0,n]$. Then we have the following lemma:

Lemma 6.

If $\langle X,\Delta\rangle \in \mathcal{P}_\Psi$ where
Ψ has the proper effect on Δ and Ψ and Δ are linear
then $\langle X,\Delta\rangle \overset{s}{\underset{\Psi}{\longmapsto}}{}^{i}_{j} \langle Y,\Delta\rangle$ iff $\widetilde{\mathcal{T}}_i(\langle X,\Delta\rangle,\Psi,n) \overset{s}{\longrightarrow} \widetilde{\mathcal{T}}_j(\langle Y,\Delta\rangle,\Psi,n)$

Proof. By induction on the length of s. For details see Appendix C. □

Like before, transformation $\widetilde{\mathcal{T}}(\,\cdot\,,\Psi,n)$ induces two relations: the satisfaction relation which holds between $Q \in \mathcal{P}_\Psi$ and $M \in \mathcal{M}$ iff $\widetilde{\mathcal{T}}(Q,\Psi,n) \models M$ and the binary relation which holds between $P \in \mathcal{P}$ and $Q \in \mathcal{P}_\Psi$ iff $P \approx \widetilde{\mathcal{T}}(Q,\Psi,n)$. Under assumptions of Lemma 6, we can show that the first relation coincides with conditional satisfaction relation and the second with conditional must-bisimilarity:

Proposition 7.

If $M \in \mathcal{M}$ and $\langle X,\Delta\rangle \in \mathcal{P}_\Psi$ where
Ψ has the proper effect on Δ and Ψ and Δ are linear
then $\langle X,\Delta\rangle \models_n M$ iff $\widetilde{\mathcal{T}}(\langle X,\Delta\rangle,\Psi,n) \models M$

Proof. We show that for all $i \in [0,n]$, $\langle X,\Delta\rangle \models^i_n M$ iff $\widetilde{\mathcal{T}}_i(\langle X,\Delta\rangle,\Psi,n) \models M$. The proof proceeds by induction on the structure of M, applying Lemma 6. For details see Appendix C. □

Proposition 8.

If $P \in \mathcal{P}$ and $\langle X,\Delta\rangle \in \mathcal{P}_\Psi$ where
Ψ has the proper effect on Δ and Ψ and Δ are linear
then $P \mathrel{\mathbb{L}}_n \langle X,\Delta\rangle$ iff $P \approx \widetilde{\mathcal{T}}(\langle X,\Delta\rangle,\Psi,n)$

Proof. For weak-image-finite $\overset{}{\underset{\Psi}{\longmapsto}}$, this statement follows from (16) and Propositions 2 and 7. For any $\overset{}{\underset{\Psi}{\longmapsto}}$ and family $\{B_i\}_{i=0}^n$ such that if $(P,Q) \in B_i$ then $Q \equiv \langle X,\Delta\rangle$, it follows from the fact that:

$$\begin{array}{l}\{B_i\}_{i=0}^n \text{ is a conditional must-bisimulation iff}\\ \widetilde{\mathfrak{T}}(\{B_i\}_{i=0}^n,\Psi,n) \text{ is a bisimulation}\end{array} \tag{32}$$

where $\widetilde{\mathfrak{T}}(\{B_i\}_{i=0}^n,\Psi,n) =_{def} \bigcup_{i=0}^n\{(P,\widetilde{\mathcal{T}}_i(Q,\Psi,n)) \mid (P,Q) \in B_i\}$.
For details of the proof see Appendix C. □

Assuming additionally that all right-side expressions in Ψ are of the form $\sum_{i=1}^k \tau.X_i$, the same result can be obtained applying auxiliary actions, concurrent composition and processes R_n (Example 5):

$$\widetilde{\mathcal{T}}'(\langle X,\Delta\rangle,\Psi,n) =_{def} (\langle X,(\Delta \odot \overline{a}) \oplus (\overline{b} \odot \Psi)\rangle \mid R_n)\setminus\{a,b\}$$

where $[\![\Delta \odot \overline{a}]\!](X) = [\![\Delta]\!](X)\{\overline{a}.Y/Y \mid Y \in \mathcal{X}([\![\Delta]\!](X))\}$.

Thus for $\mathbb{L}_n$ and under assumptions of Lemma 6, we have the equivalence diagram similar to (1), given $\overset{}{\underset{\Psi}{\cdots\!\rightarrow}}$ such that $\overset{}{\underset{\Psi}{\longmapsto}}$ is weak-image-finite. Consider $\widehat{\mathfrak{T}}(\{B_i\}_{i=0}^n,n) =_{def} \bigcup_{i=0}^n\{(P,\widehat{\mathcal{T}}_i(Q,n)) \mid (P,Q) \in B_i\}$. For $\mathbb{L}_n$ we can show that $\{B_i\}_{i=0}^n$ is a conditional may-bisimulation iff $\widetilde{\mathfrak{T}}(\{B_i\}_{i=0}^n,\Psi,n)$ and $\widehat{\mathfrak{T}}(\{B_i\}_{i=0}^n,n)$ are bisimulations.

Although the linear form of Δ and Ψ is necessary to establish these results, the meaning of $\widetilde{\mathcal{T}}$ and $\widehat{\mathcal{T}}$ for non-linear Δ and Ψ is also well-understood. While in the first case all transitions are significant, they all 'update' the index i, in the second case only chosen ones are significant. The main reason for 'mismatch' between $\widetilde{\mathcal{T}}(\cdot,\Psi,n)$ and $\overset{i}{\underset{\Psi}{\longmapsto}}_j$ for non-linear expressions lies in the restrictive form of the latter. In the following example we will illustrate using transformations $\widetilde{\mathcal{T}}$ to verify conditional fault-tolerance for $\langle E,\Delta\rangle$ and Ψ which is not linear.

Example 9. Consider $n,m > 0$ and the task to ensure a reliable communication (specified by Buf_{m+n+2}) over a medium of capacity m which permutes messages. To this end we will use two processes: the sender Sp_n and the receiver Rp_n. In order to determine the proper transmission order, messages will be send by Sp_n with their sequence numbers modulo n. The value of n determines the number of parallel components St_i of Rp_n ($i = 0,\ldots,n-1$), each one used to store a message with the sequence value i, received out-of-order. The value of $\perp$ means that no message is stored. Suppose that the summation $i+1$ below is taken modulo n. Then we have:

$$
\begin{array}{lll}
Sp_n & =_{def} & \langle Zs(0), [Zs(i) \mathrel{\widehat{=}} in(x).\overline{out}(x,i).Zs(i+1) \mid 0 \leq i < n] \rangle \\
Rp_n & =_{def} & (Ctr \mid St_0 \mid \cdots \mid St_{n-1}) \setminus \{st_0, \ldots, st_{n-1}\} \\
Ctr & =_{def} & \langle Zr(0), \; [Zr(i) \mathrel{\widehat{=}} in(x,j).\ \textbf{if } i = j \\
 & & \qquad \textbf{then } \overline{out}(x).Zm(i+1) \\
 & & \qquad \textbf{else } \overline{st_j}(x).Zr(i) \qquad \mid 0 \leq i < n] \\
 & & [Zm(i) \mathrel{\widehat{=}} st_i(x).\ \textbf{if } x = \perp \\
 & & \qquad \textbf{then } Zr(i) \\
 & & \qquad \textbf{else } \overline{out}(x).Zm(i+1) \mid 0 \leq i < n] \rangle \\
St_i & =_{def} & \langle Zm_i, \quad [Zm_i \mathrel{\widehat{=}} st_i(x).\overline{st_i}(x).Zm_i + \overline{st_i}(\perp).Zm_i] \rangle, \; 0 \leq i < n
\end{array}
$$

We can shown that the process $Sp_n \frown Buf_m \frown Rp_n$ tolerates Ψp, provided the number of successive permutations is not greater than n:

$$Buf_{m+n+2} \approx \widetilde{\mathcal{T}}(Sp_n \frown Buf_m \frown Rp_n, \Psi p, n)$$ □

8 Conclusions

Currently, there is a number of methods for specifying and proving correctness of systems which are tolerant of faults in the operating environment [Cri85, JH87, LJ91, Nor92, Pel91, PJ93, Pra87]. Based on different formalisms and various semantic models, of systems and faults, using different ways to represent effects of faults on the behaviour of systems, they are difficult to compare and relate. In particular, it is not certain whether a system which is fault-tolerant with respect to one of these methods is also fault-tolerant according to the others.

This relationship is clear for three methods defined in this paper: algebraic, logical and transformational. Based on the common semantic model of labelled transition systems, which is also used to model faults, all three methods have been proved equivalent for certain classes of faults. The equivalence holds in two cases, unconditional, where no assumption is made about the quantity of faults, and conditional, given the maximal number of times they can occur successively.

There is a number of directions that we plan to develop this work. We plan to study the use of other bisimulation-like relations, like the partial [Wal90], the 'terminating' [AH92] and the context-dependent [Lar87] bisimilarities for fault-tolerance. In the presence of faulty transitions, a convergent process may diverge and the one which terminates (successfully) may deadlock. We plan to relate our theory with modal specifications [LT88] which constrain possible implementations by two kinds of transitions, necessary and admissible (any necessary transition is also admissible). Bisimulation gives rise to the refinement ordering between modal specifications which is different however from the relations defined in this paper. We plan to determine the class of contexts (built from the operators of the process language) where our relations, especially the stronger, may-bisimilarity, is substitutive. Last but not least, we plan to support the development of fault-tolerant processes, based on the verification theory of this paper and using the decomposition of faults specified in our language.

Acknowledgements

I am grateful to my supervisor, Mathai Joseph, for many valuable comments on draft versions of this paper, to Zhiming Liu for stimulating discussions on fault-tolerance, and to David Walker for his reading, helpful comments and for putting some literature to my attention.

References

[AH92] L. Aceto and M. Hennessy. Termination, deadlock and divergence. *Journal of ACM*, 39(1):147–187, 1992.

[Cri85] F. Cristian. A rigorous approach to fault-tolerant programming. *IEEE Transactions on Software Engineering*, 11(1):23–31, 1985.

[HM85] M. Hennessy and R. Milner. Algebraic laws for nondeterminism and concurrency. *Journal of the ACM*, 32(1):137–161, 1985.

[JH87] He Jifeng and C.A.R. Hoare. Algebraic specification and proof of a distributed recovery algorithm. *Distributed Computing*, 2:1–12, 1987.

[Kel76] R. Keller. Formal verification of parallel programs. *Communications of ACM*, 19(7):561–572, 1976.

[Lar87] K.G. Larsen. A context dependent equivalence between processes. *Theoretical Computer Science*, 49:185–215, 1987.

[Liu91] Z. Liu. *Fault-Tolerant Programming by Transformations.* PhD thesis, University of Warwick, 1991.

[LJ91] Z. Liu and M. Joseph. Transformations of programs for fault-tolerance. *Formal Aspects of Computing*, 4:442–469, 1991.

[LT87] N. Lynch and M. Tuttle. Hierarchical correctness proofs for distributed algorithms. Technical report, MIT Laboratory for Computer Science, 87.

[LT88] K.G. Larsen and B. Thomsen. A modal process logic. In *Proc. 3rd Annual Symposium on Logic in Computer Science*, pages 203–210, 88.

[Mil89] R. Milner. *Communication and Concurrency.* Prentice-Hall International, 1989.

[MP91] Z. Manna and A. Pnueli. *The Temporal Logic of Reactive and Concurrent Systems*, volume 1. Springer-Verlag, 1991.

[Nor92] J. Nordahl. *Specification and Design of Dependable Communicating Systems.* PhD thesis, Technical University of Denmark, 1992.

[Par81] D. Park. Concurrency and automata on infinite sequences. *LNCS*, 104, 81.

[Pel91] J. Peleska. Design and verification of fault tolerant systems with CSP. *Distributed Computing*, 5:95–106, 1991.

[PJ93] D. Peled and M. Joseph. A compositional approach for fault-tolerance using specification transformation. *LNCS*, 694, 1993.

[Plo81] G. Plotkin. A structural approach to operational semantics. Technical report, Computer Science Department, Aarhus University, 81.

[Pra86] V. Pratt. Modeling concurrency with partial orders. *International Journal of Parallel Programming*, 15(1):33–71, 1986.

[Pra87] K.V.S. Prasad. *Combinators and Bisimulation Proofs for Restartable Systems.* PhD thesis, Department of Computer Science, University of Edinburgh, 1987.

[Wal90] D.J. Walker. Bisimulation and divergence. *Information and Computation*, 85:202–241, 90.

[Win89] G. Winskel. An introduction to event structures. *LNCS*, 354:364–397, 1989.

A Proofs from Section 4

Lemma 9.
B is a must-bisimulation iff for all $(P,Q) \in B$ and $s \in \mathcal{A}^*$:

$$\text{whenever } P \xrightarrow{s} P' \text{ then } \exists_{Q',t}\ Q \overset{t}{\longmapsto} Q' \wedge \widehat{s} = \widehat{t} \wedge (P',Q') \in B$$
$$\text{whenever } Q \overset{s}{\longmapsto} Q' \text{ then } \exists_{P',t}\ P \xrightarrow{t} P' \wedge \widehat{s} = \widehat{t} \wedge (P',Q') \in B$$

Proof. ($\Rightarrow$) Consider $(P,Q) \in B$ and let $P \equiv P_0 \xrightarrow{\alpha_1} \cdots \xrightarrow{\alpha_n} P_n \equiv P'$ where $n \geq 0$ and $s = \alpha_1 \ldots \alpha_n$. If $n = 0$ then $P' \equiv P$ and it is enough to take $t = \varepsilon$ and $Q' \equiv Q$. If $n > 0$ then for all $k \in [1,n]$ there exists Q_k and t_k such that $Q_{k-1} \overset{t_k}{\longmapsto} Q_k$, $\widehat{t_k} = \widehat{\alpha_k}$ and $(P_k, Q_k) \in B$. Let $t =_{def} t_1 : \cdots : t_n$. Then $Q \overset{t}{\longmapsto} Q'$, $\widehat{t} = \widehat{s}$ and $(P',Q') \in B$. For $Q \overset{s}{\longmapsto} Q'$ the proof is the same.
($\Leftarrow$) It is enough to take s such that $\sharp s = 1$. □

Proof of Proposition 1.
($\Rightarrow$) Let $P \sqsubset\!\sqsupset Q$ and $M \in \mathcal{M}$. We will show that $P \models M$ iff $Q \models M$, by induction on M. Let $P \models \langle\beta\rangle M_1$ (for $true$, $M_1 \wedge M_2$ and $\neg M_1$ the proof is obvious). Then $P \xrightarrow{s} P'$ where $\widehat{s} = \beta$ and $P' \models M_1$ for some s and P'. Applying Lemma 9, there exists t and Q' such that $Q \overset{t}{\longmapsto} Q'$, $\widehat{t} = \widehat{s}$ and $P' \sqsubset\!\sqsupset Q'$. Then, by induction we have $Q' \models M_1$ and finally $Q \models \langle\beta\rangle M_1$. For $Q \models \langle\beta\rangle M_1$ the proof is similar.
($\Leftarrow$) Consider $B =_{def} \{(P,Q) \mid \forall_{M \in \mathcal{M}}\ P \models M \Leftrightarrow Q \models M\}$ and suppose that $\longmapsto$ is weak-image-finite. We will show that B is a must-bisimulation. Let $(P,Q) \in B$ and $P \xrightarrow{\alpha} P'$. Then $P \models \langle\widehat{\alpha}\rangle true$ and $Q \models \langle\widehat{\alpha}\rangle true$ what gives $Q \overset{s}{\longmapsto} Q'$ for some s ($\widehat{s} = \widehat{\alpha}$) and Q'. Let $\mathcal{Q}$ be the set of all such Q'. Because $\longmapsto$ is weak-image-finite, we have $\mathcal{Q} = \{Q_i\}_{i=1}^n$ for some $n > 0$. Then, it is enough to find $i \in [1,n]$ such that $(P', Q_i) \in B$. Suppose on the contrary, that for all $i \in [1,n]$ there exists $M_i \in \mathcal{M}$ such that $P' \models M_i$ and $Q_i \not\models M_i$. Let $M =_{def} M_1 \wedge \cdots \wedge M_n$. Then we have $P' \models M$ and so $P \models \langle\widehat{\alpha}\rangle M$, however $Q \not\models \langle\widehat{\alpha}\rangle M$ in spite of $(P,Q) \in B$. For $Q \overset{\alpha}{\longmapsto} Q'$ the proof is the same ($\longrightarrow$ is also weak-image-finite). □

Lemma 10.
If $Q \overset{s}{\longmapsto}{}^i_j Q' \overset{t}{\longmapsto}{}^j_k Q''$ then $Q \overset{s:t}{\longmapsto}{}^i_k Q''$.

Proof. By induction on the length of s. If $Q \overset{\varepsilon}{\longmapsto}{}^i_j Q' \overset{t}{\longmapsto}{}^j_k Q''$ then we have $Q' \equiv Q$ and $i = j$ what implies $Q \overset{\varepsilon:t}{\longmapsto}{}^i_k Q''$. If $Q \overset{\alpha:s}{\longmapsto}{}^i_j Q' \overset{t}{\longmapsto}{}^j_k$ then either $i < n$ and $Q \overset{\alpha}{\cdots\!\rightarrow} Q''' \overset{s}{\longmapsto}{}^{i+1}_j Q' \overset{t}{\longmapsto}{}^j_k Q''$ or $Q \xrightarrow{\alpha} Q''' \overset{s}{\longmapsto}{}^0_j Q' \overset{t}{\longmapsto}{}^j_k Q''$ for some Q'''. In the first case by induction $Q''' \overset{s:t}{\longmapsto}{}^{i+1}_k Q''$ so we get $Q \overset{\alpha:s:t}{\longmapsto}{}^i_k Q''$. In the second case by induction $Q''' \overset{s:t}{\longmapsto}{}^0_k Q''$ so we finally get $Q \overset{\alpha:s:t}{\longmapsto}{}^i_k Q''$. □

Lemma 11.
Consider $\{B_i\}_{i=0}^n$ where $B_i \subseteq \mathcal{P} \times \mathcal{P}$ for $i \in [0,n]$. $\{B_i\}_{i=0}^n$ is a conditional must-bisimulation iff for all $i,j \in [0,n]$, $(P,Q) \in B_i$ and $s \in \mathcal{A}^*$:

$$\text{whenever } P \xrightarrow{s} P' \text{ then } \exists_{Q',j,t}\ Q \overset{t}{\longmapsto}{}^i_j Q' \wedge \widehat{t} = \widehat{s} \wedge (P',Q') \in B_j$$
$$\text{whenever } Q \overset{s}{\longmapsto}{}^i_j Q' \text{ then } \exists_{P',t}\ \ P \xrightarrow{t} P' \wedge \widehat{t} = \widehat{s} \wedge (P',Q') \in B_j$$

Proof. ($\Rightarrow$) Consider $i \in [0,n]$, $(P,Q) \in B_i$ and $s = \alpha_1 \ldots \alpha_n$. Suppose that $P \equiv P_0 \stackrel{\alpha_1}{\longrightarrow} \cdots \stackrel{\alpha_n}{\longrightarrow} P_n \equiv P'$ for $n \geq 0$. If $n = 0$ then it is enough to take $t = \varepsilon$, $j = i$ and $Q' \equiv Q$. It $n > 0$ then for all $k \in [1,n]$ there exists Q_k, t_k and i_k such that $Q_{k-1} \stackrel{t_k}{\longmapsto}_{i_k}^{i_{k-1}} Q_k$, $\widehat{t_k} = \widehat{\alpha_k}$ and $(P_k, Q_k) \in B_{i_k}$ where $Q_0 \equiv Q$ and $i_0 = i$. Let $t = t_1 : \cdots : t_n$. Then $\widehat{t} = \widehat{s}$, $(P', Q_n) \in B_{i_n}$ and from Lemma 10, $Q \stackrel{t}{\longmapsto}_{i_n}^{i} Q_n$. For $Q \stackrel{s}{\longmapsto}_{j}^{i} Q'$ the proof is similar.
($\Leftarrow$) It is enough to take s such that $\sharp s = 1$. □

Proof of Proposition 2.
($\Rightarrow$) Let $i \in [0,n]$, $P \,\mathrm{I\!I}_n^i Q$ and $M \in \mathcal{M}$. We will show that $P \models M$ iff $Q \models_n^i M$ by induction on M. Let $P \models \langle\beta\rangle M_1$ (for $true$, $M_1 \wedge M_2$ and $\neg M_1$ the proof is obvious). Then $P \stackrel{s}{\longrightarrow} P'$ and $\widehat{s} = \beta$ for some P' such that $P' \models M_1$. Applying Lemma 11, there exists t ($\widehat{t} = \beta$) and j such that $Q \stackrel{t}{\longmapsto}_{j}^{i} Q'$ and $P' \,\mathrm{I\!I}_n^j Q'$ for some Q'. Then by induction we have $Q' \models_n^j M_1$ and thus $Q \models_n^i \langle\beta\rangle M_1$. For $Q \models_n^i \langle\beta\rangle M_1$ the proof is similar.
($\Leftarrow$) Let $B_i =_{def} \{(P,Q) \mid \forall_{M \in \mathcal{M}}\ P \models M \Leftrightarrow Q \models_n^i M\}$ and suppose that $\longmapsto$ is weak-image-finite. We will show that $\{B_i\}_{i=0}^n$ is a conditional must-bisimulation. Suppose that $i \in [0,n]$, $(P,Q) \in B_i$ and $P \stackrel{\alpha}{\longrightarrow} P'$. Then we have $P \models \langle\widehat{\alpha}\rangle true$ and $Q \models_n^i \langle\widehat{\alpha}\rangle true$ so $Q \stackrel{s}{\longmapsto}_{j}^{i} Q'$ for some j and s such that $\widehat{s} = \widehat{\alpha}$. Let $\mathcal{Q}_j$ be the set of all such Q'. Because $\longmapsto$ is weak-image-finite, we have $\mathcal{Q}_j = \{Q_k^j\}_{k=1}^{k_j}$ where $k_j \geq 0$ and there exists j such that $k_j > 0$. Then it is enough to show that for some j there exists $k \in [1,k_j]$ such that $(P', Q_k^j) \in B_j$. Suppose on the contrary: for all $j \in [0,n]$ and for all $k \in [1,k_j]$ there exists $M_k^j \in \mathcal{M}$ such that $P' \models M_k^j$ and $Q_k^j \not\models M_k^j$. Let $M =_{def} \bigwedge_{j=0}^{n} \bigwedge_{k=1}^{k_j} M_k^j$. Then we have $P' \models M$ so $P \models \langle\widehat{\alpha}\rangle M$, however $Q \not\models_n^i \langle\widehat{\alpha}\rangle M$ in spite of $(P,Q) \in B_i$. For $Q \stackrel{s}{\longmapsto}_{j}^{i} Q'$ the proof is the same, because $\longrightarrow$ is also weak-image-finite. □

B Proofs from Section 7.1

Proof of Lemma 3.
We proceed by transitional induction. Let $\Gamma =_{def} \Delta \oplus \Psi$.
($\Rightarrow$) If $Q \equiv \langle E, \Delta\rangle \stackrel{\alpha}{\longmapsto}_{\Psi} Q'$ then $Q' \equiv \langle E', \Delta\rangle$ (28) and we will show that $\langle E, \Gamma\rangle \stackrel{\alpha}{\longrightarrow} \langle E', \Gamma\rangle$ by induction on the inference of transition $\langle E, \Delta\rangle \stackrel{\alpha}{\longmapsto}_{\Psi} \langle E', \Delta\rangle$. There are six cases:

1. $E \equiv \alpha E'$. Then $\langle \alpha.E', \Gamma\rangle \stackrel{\alpha}{\longrightarrow} \langle E', \Gamma\rangle$.
2. $E \equiv E_1 + E_2$. Then either $\langle E_1, \Delta\rangle \stackrel{\alpha}{\longmapsto}_{\Psi} \langle E', \Delta\rangle$ or $\langle E_2, \Delta\rangle \stackrel{\alpha}{\longmapsto}_{\Psi} \langle E', \Delta\rangle$. For the first (the second is symmetrical), by induction $\langle E_1, \Gamma\rangle \stackrel{\alpha}{\longrightarrow} \langle E', \Gamma\rangle$ and so we have $\langle E_1 + E_2, \Gamma\rangle \stackrel{\alpha}{\longrightarrow} \langle E', \Gamma\rangle$.
3. $E \equiv E_1 | E_2$. Then there are three cases:
   (a) $\langle E_1, \Delta\rangle \stackrel{\alpha}{\longmapsto}_{\Psi} \langle E_1', \Delta\rangle$ where $E' \equiv E_1' | E_2$. Then by induction we have $\langle E_1, \Gamma\rangle \stackrel{\alpha}{\longrightarrow} \langle E_1', \Gamma\rangle$ and so $\langle E_1 | E_2, \Gamma\rangle \stackrel{\alpha}{\longrightarrow} \langle E_1' | E_2, \Gamma\rangle$.
   (b) $\langle E_2, \Delta\rangle \stackrel{\alpha}{\longmapsto}_{\Psi} \langle E_2', \Delta\rangle$ where $E' \equiv E_1 | E_2'$ is similar.
   (c) $\langle E_1, \Delta\rangle \stackrel{a}{\longmapsto}_{\Psi} \langle E_1', \Delta\rangle$ and $\langle E_2, \Delta\rangle \stackrel{\overline{a}}{\longmapsto}_{\Psi} \langle E_2', \Delta\rangle$ for $\alpha = \tau$ and $E' \equiv E_1' | E_2'$. By induction $\langle E_1, \Gamma\rangle \stackrel{a}{\longrightarrow} \langle E_1', \Gamma\rangle$ and $\langle E_2, \Gamma\rangle \stackrel{\overline{a}}{\longrightarrow} \langle E_2', \Gamma\rangle$ and so we have $\langle E_1 | E_2, \Gamma\rangle \stackrel{\tau}{\longrightarrow} \langle E_1' | E_2', \Gamma\rangle$.

4. $E \equiv E_1 \setminus L$. Then $\langle E_1, \Delta\rangle \stackrel{\alpha}{\mapsto\!\!\!\!\!\rightarrow}_{\Psi} \langle E_1', \Delta\rangle$ for $\alpha, \overline{\alpha} \notin L$ and $E' \equiv E_1' \setminus L$. Thus by induction $\langle E_1, \Gamma\rangle \stackrel{\alpha}{\longrightarrow} \langle E_1', \Gamma\rangle$ and so $\langle E_1 \setminus L, \Gamma\rangle \stackrel{\alpha}{\longrightarrow} \langle E_1' \setminus L, \Gamma\rangle$.
5. $E \equiv E_1[f]$ is similar.
6. $E \equiv X$. Then either $\langle [\![\Delta]\!](X), \Delta\rangle \stackrel{\alpha}{\mapsto\!\!\!\!\!\rightarrow}_{\Psi} \langle E', \Delta\rangle$ or $\langle [\![\Psi]\!](X), \Delta\rangle \stackrel{\alpha}{\mapsto\!\!\!\!\!\rightarrow}_{\Psi} \langle E', \Delta\rangle$. Consider the first (the second is symmetrical), then $X \in dom(\Delta)$ and by induction $\langle [\![\Delta]\!](X), \Gamma\rangle \stackrel{\alpha}{\longrightarrow} \langle E', \Gamma\rangle$. There are two cases:
   (a) $X \notin dom(\Psi)$. Then $\langle X, \Gamma\rangle \stackrel{\alpha}{\longrightarrow} \langle E', \Gamma\rangle$ because $[\![\Gamma]\!](X) = [\![\Delta]\!](X)$.
   (b) $X \in dom(\Psi)$. Then $[\![\Gamma]\!](X) = [\![\Delta]\!](X) + [\![\Psi]\!](X)$ from (25) and so because $\langle [\![\Delta]\!](X) + [\![\Psi]\!](X), \Gamma\rangle \stackrel{\alpha}{\longrightarrow} \langle E', \Gamma\rangle$, we have $\langle X, \Gamma\rangle \stackrel{\alpha}{\longrightarrow} \langle E', \Gamma\rangle$.

($\Leftarrow$) Let $\mathcal{T}(Q, \Psi) \stackrel{\alpha}{\longrightarrow} R'$ where $Q \equiv \langle E, \Delta\rangle$. Then $R' \equiv \mathcal{T}(\langle E', \Delta\rangle, \Psi)$ applying (28) and we will show that $\langle E, \Delta\rangle \stackrel{\alpha}{\mapsto\!\!\!\!\!\rightarrow}_{\Psi} \langle E', \Delta\rangle$, using the similar induction on the inference of the transition $\langle E, \Gamma\rangle \stackrel{\alpha}{\longrightarrow} \langle E', \Gamma\rangle$. Consider $E \equiv X$ only. Then $\langle [\![\Gamma]\!](X), \Gamma\rangle \stackrel{\alpha}{\longrightarrow} \langle E', \Gamma\rangle$ and by induction we have $\langle [\![\Gamma]\!](X), \Delta\rangle \stackrel{\alpha}{\mapsto\!\!\!\!\!\rightarrow}_{\Psi} \langle E', \Delta\rangle$. Thus there are two cases:

1. $X \notin dom(\Psi)$. Then $[\![\Gamma]\!](X) = [\![\Delta]\!](X)$ so $\langle X, \Delta\rangle \stackrel{\alpha}{\mapsto\!\!\!\!\!\rightarrow}_{\Psi} \langle E', \Delta\rangle$.
2. $X \in dom(\Psi)$. Then $[\![\Gamma]\!](X) = [\![\Delta]\!](X) + [\![\Psi]\!](X)$ applying (25) so we either have $\langle [\![\Delta]\!](X), \Delta\rangle \stackrel{\alpha}{\mapsto\!\!\!\!\!\rightarrow}_{\Psi} \langle E', \Delta\rangle$ or $\langle [\![\Psi]\!](X), \Delta\rangle \stackrel{\alpha}{\mapsto\!\!\!\!\!\rightarrow}_{\Psi} \langle E', \Delta\rangle$. In both cases we finally get $\langle X, \Delta\rangle \stackrel{\alpha}{\mapsto\!\!\!\!\!\rightarrow}_{\Psi} \langle E', \Delta\rangle$. □

Proof of Proposition 4.
We will show that $Q \Vdash M$ iff $\mathcal{T}(Q, \Psi) \models M$ by induction on the structure of M. Let $\mathcal{T}(Q, \Psi) \models \langle\beta\rangle M_1$ (for $true$, $M_1 \wedge M_2$ and $\neg M_1$ the proof is obvious). Then $\mathcal{T}(Q, \Psi) \stackrel{s}{\longrightarrow} R'$ where $\widehat{s} = \beta$, $R' \models M_1$ and $R' \equiv \mathcal{T}(Q', \Psi)$ applying (28). Thus $Q \stackrel{s}{\mapsto\!\!\!\!\!\rightarrow}_{\Psi} Q'$ from Lemma 3 and by induction we have $Q' \Vdash M_1$. As a result $Q \Vdash \langle\beta\rangle M_1$. For $Q \Vdash \langle\beta\rangle M_1$ the proof is similar. □

Proof of Proposition 5.
($\Rightarrow$) Let $P \mathrel{\text{⫤}} Q$. Then $(P, Q) \in B$ which a must-bisimulation for $\stackrel{\cdots}{\longrightarrow}_{\Psi}$. We will show that $\mathfrak{T}(B, \Psi) =_{def} \{(P, \mathcal{T}(Q, \Psi)) \mid (P, Q) \in B\}$ is a bisimulation. If $P \stackrel{\alpha}{\longrightarrow} P'$ then $Q \stackrel{s}{\mapsto\!\!\!\!\!\rightarrow}_{\Psi} Q'$ where $\widehat{s} = \widehat{\alpha}$ and $(P', Q') \in B$ for some Q'. But then $(P', \mathcal{T}(Q', \Psi)) \in \mathfrak{T}(B, \Psi)$ and $\mathcal{T}(Q, \Psi) \stackrel{s}{\longrightarrow} \mathcal{T}(Q', \Psi)$ what follows from Lemma 3. If $\mathcal{T}(Q, \Psi) \stackrel{\alpha}{\longrightarrow} R'$ then we get $R' \equiv \mathcal{T}(Q', \Psi)$ (28) and applying Lemma 3 we have $Q \stackrel{\alpha}{\mapsto\!\!\!\!\!\rightarrow}_{\Psi} Q'$. Thus there exists P' such that $P \stackrel{s}{\longrightarrow} P'$, $\widehat{s} = \widehat{\alpha}$ and $(P', Q') \in B$. Because then $(P', \mathcal{T}(Q', \Psi)) \in \mathfrak{T}(B, \Psi)$, $\mathfrak{T}(B, \Psi)$ is a bisimulation.
($\Leftarrow$) If $P \approx \mathcal{T}(Q, \Psi)$ then $(P, \mathcal{T}(Q, \Psi)) \in B$ which is the least bisimulation with this property. Applying (28), it is easy to see that there exists $C \subseteq \mathcal{P} \times \mathcal{P}_{\Psi}$ such that $B = \mathfrak{T}(C, \Psi)$ and it is enough to prove that C is a must-bisimulation. Let $(P, Q) \in C$ and $P \stackrel{\alpha}{\longrightarrow} P'$. Then $(P, \mathcal{T}(Q, \Psi)) \in B$ and applying (28) there exists Q' and s such that $\mathcal{T}(Q, \Psi) \stackrel{s}{\longrightarrow} \mathcal{T}(Q', \Psi)$, $\widehat{s} = \widehat{\alpha}$ and $(P', \mathcal{T}(Q', \Psi)) \in B$. Then we have $(P', Q') \in C$ and $Q \stackrel{s}{\mapsto\!\!\!\!\!\rightarrow}_{\Psi} Q'$ from Lemma 3. Let $Q \stackrel{\alpha}{\mapsto\!\!\!\!\!\rightarrow}_{\Psi} Q'$. Then $\mathcal{T}(Q, \Psi) \stackrel{\alpha}{\longrightarrow} \mathcal{T}(Q', \Psi)$ from Lemma 3 and there exists P' and s such that $P \stackrel{s}{\longrightarrow} P'$, $\widehat{s} = \widehat{\alpha}$ and $(P', \mathcal{T}(Q', \Psi)) \in B$. Thus $(P', Q') \in C$ what completes the proof that C is a must-bisimulation. □

C Proofs from Section 7.2

Proof of Lemma 6:
We proceed by induction on the length of s. Let $\Gamma_n =_{def} \Delta_n \oplus \Psi_n$.
($\Rightarrow$) For $s = \varepsilon$ it is immediate. Let $\langle X, \Delta\rangle \overset{\alpha:s}{\underset{\Psi}{\longmapsto}}{}^{i}_{j} \langle Y, \Delta\rangle$. Because the right-side expressions of Δ and Ψ are linear, applying (31) there exists Z and k such that $\langle X, \Delta\rangle \overset{\alpha}{\underset{\Psi}{\longmapsto}} \langle Z, \Delta\rangle \overset{s}{\underset{\Psi}{\longmapsto}}{}^{k}_{j} \langle Y, \Delta\rangle$. There are two cases:

- $\langle X, \Delta\rangle \xrightarrow{\alpha} \langle Z, \Delta\rangle$ where $k = 0$. Because Ψ has the proper effect on Δ, this transition can be only inferred from $\langle [\![\Delta]\!](X), \Delta\rangle \xrightarrow{\alpha} \langle Z, \Delta\rangle$. Thus we have $\langle [\![\Gamma_n]\!](X_i), \Gamma_n\rangle \xrightarrow{\alpha} \langle Z_0, \Gamma_n\rangle$ and by induction:

$$\widetilde{\mathcal{T}}_i(\langle X, \Delta\rangle, \Psi, n) \xrightarrow{\alpha} \widetilde{\mathcal{T}}_0(\langle Z, \Delta\rangle, \Psi, n) \xrightarrow{s} \widetilde{\mathcal{T}}_j(\langle Y, \Delta\rangle, \Psi, n)$$

- $\langle X, \Delta\rangle \overset{\alpha}{\underset{\Psi}{\cdots\!\!\rightarrow}} \langle Z, \Delta\rangle$ where $i \neq n$ and $k = i + 1$. It can be only inferred from $\langle [\![\Psi]\!](X), \Delta\rangle \overset{\alpha}{\underset{\Psi}{\longmapsto}} \langle Z, \Delta\rangle$ and so $\langle [\![\Psi]\!](X), \Delta\rangle \xrightarrow{\alpha} \langle Z, \Delta\rangle$ (Ψ is linear). Thus we have $\langle [\![\Gamma_n]\!](X_i), \Gamma_n\rangle \xrightarrow{\alpha} \langle Z_{i+1}, \Gamma_n\rangle$ because $i \neq n$ and by induction

$$\widetilde{\mathcal{T}}_i(\langle X, \Delta\rangle, \Psi, n) \xrightarrow{\alpha} \widetilde{\mathcal{T}}_{i+1}(\langle Z, \Delta\rangle, \Psi, n) \xrightarrow{s} \widetilde{\mathcal{T}}_j(\langle Y, \Delta\rangle, \Psi, n)$$

($\Leftarrow$) For $s = \varepsilon$ it is immediate. Let $\widetilde{\mathcal{T}}_i(\langle X, \Delta\rangle, \Psi, n) \xrightarrow{\alpha:s} \widetilde{\mathcal{T}}_j(\langle Y, \Delta\rangle, \Psi, n)$. Applying (31), $\widetilde{\mathcal{T}}_i(\langle X, \Delta\rangle, \Psi, n) \xrightarrow{\alpha} \widetilde{\mathcal{T}}_k(\langle Z, \Delta\rangle, \Psi, n) \xrightarrow{s} \widetilde{\mathcal{T}}_j(\langle Y, \Delta\rangle, \Psi, n)$ for some Z and k. The first transition can be only inferred from $\langle [\![\Gamma_n]\!](X_i), \Gamma_n\rangle \xrightarrow{\alpha} \langle Z_k, \Gamma_n\rangle$ and there are two possible cases:

- $\langle [\![\Delta_n]\!](X_i), \Gamma_n\rangle \xrightarrow{\alpha} \langle Z_k, \Gamma_n\rangle$ where $k = 0$. Then $\langle [\![\Delta]\!](X), \Delta\rangle \xrightarrow{\alpha} \langle Z, \Delta\rangle$ and by induction we have $\langle X, \Delta\rangle \xrightarrow{\alpha} \langle Z, \Delta\rangle \overset{s}{\underset{\Psi}{\longmapsto}}{}^{0}_{j} \langle Y, \Delta\rangle$ and $\langle X, \Delta\rangle \overset{\alpha:s}{\underset{\Psi}{\longmapsto}}{}^{i}_{j} \langle Y, \Delta\rangle$.
- $\langle [\![\Psi_n]\!](X_i), \Gamma_n\rangle \xrightarrow{\alpha} \langle Z_k, \Gamma_n\rangle$ where $i \neq n$ and $k = i + 1$. Then we have $\langle [\![\Psi]\!](X), \Delta\rangle \xrightarrow{\alpha} \langle Z, \Delta\rangle$ so $\langle X, \Delta\rangle \overset{\alpha}{\underset{\Psi}{\cdots\!\!\rightarrow}} \langle Z, \Delta\rangle$ (Ψ has the proper effect on Δ). Thus by induction, $\langle X, \Delta\rangle \overset{\alpha}{\underset{\Psi}{\cdots\!\!\rightarrow}} \langle Z, \Delta\rangle \overset{s}{\underset{\Psi}{\longmapsto}}{}^{i+1}_{j} \langle Y, \Delta\rangle$ so $\langle X, \Delta\rangle \overset{\alpha:s}{\underset{\Psi}{\longmapsto}}{}^{i}_{j} \langle Y, \Delta\rangle$. □

Proof of Proposition 7:
It is enough to show that for all $i \in [0, n]$ and $M \in \mathcal{M}$:

$$\langle X, \Delta\rangle \models_i M \text{ iff } \widetilde{\mathcal{T}}_i(\langle X, \Delta\rangle, \Psi, n) \models M$$

The proof proceeds by induction on the structure of M.
($\Rightarrow$) Let $\langle X, \Delta\rangle \models_i \langle\beta\rangle M_1$. Then $\langle X, \Delta\rangle \overset{s}{\underset{\Psi}{\longmapsto}}{}^{i}_{j} Q'$ where $\widehat{s} = \beta$, $Q' \models_j M_1$ and by (31), $Q' \equiv \langle Y, \Delta\rangle$ for some $Y \in \mathcal{X}$. Thus $\widetilde{\mathcal{T}}_i(\langle X, \Delta\rangle, \Psi, n) \xrightarrow{s} \widetilde{\mathcal{T}}_j(\langle Y, \Delta\rangle, \Psi, n)$ applying Lemma 6 and by induction, $\widetilde{\mathcal{T}}_j(\langle Y, \Delta\rangle, \Psi, n) \models M_1$. Then because $\widehat{s} = \beta$, we finally have $\widetilde{\mathcal{T}}_i(\langle X, \Delta\rangle, \Psi, n) \models \langle\beta\rangle M_1$. In the remaining cases ($true$, $\neg M_1$ and $M_1 \wedge M_2$) the proof is obvious.
($\Leftarrow$) Let $\widetilde{\mathcal{T}}_i(\langle X, \Delta\rangle, \Psi, n) \models \langle\beta\rangle M_1$. Then $\widetilde{\mathcal{T}}_i(\langle X, \Delta\rangle, \Psi, n) \xrightarrow{s} R'$ where $\widehat{s} = \beta$, $R' \models M_1$ and by (31), $R' \equiv \widetilde{\mathcal{T}}_j(\langle Y, \Delta\rangle, \Psi, n)$ for some $Y \in \mathcal{X}$ and $j \in [0, n]$. Thus we have $\langle X, \Delta\rangle \overset{s}{\underset{\Psi}{\longmapsto}}{}^{i}_{j} \langle Y, \Delta\rangle$ applying Lemma 6, and by induction $\langle Y, \Delta\rangle \models_j M_1$. Finally we get $\langle X, \Delta\rangle \models_i \langle\beta\rangle M_1$ ($\widehat{s} = \beta$). The remaining cases are simple. □

Proof of Proposition 8:
($\Rightarrow$) Let $P \mathrel{\mathbb{L}}_n \langle X, \Delta \rangle$. Then $(P, \langle X, \Delta \rangle) \in B_0$ where $\{B_i\}_{i=0}^n$ is the smallest conditional must-bisimulation with this property, and it is enough to show that

$$\widetilde{\mathfrak{T}}(\{B_i\}_{i=0}^n, \Psi, n) =_{def} \bigcup_{i=0}^n \{(P, \widetilde{T}_i(\langle X, \Delta \rangle, \Psi, n)) \mid (P, \langle X, \Delta \rangle) \in B_i\}$$

is a bisimulation. Let $i \in [0, n]$ and $(P, \widetilde{T}_i(\langle X, \Delta \rangle, \Psi, n)) \in \widetilde{\mathfrak{T}}(\{B_i\}_{i=0}^n, \Psi, n)$ where $(P, \langle X, \Delta \rangle) \in B_i$. We have:

- If $P \xrightarrow{\alpha} P'$ then $\langle X, \Delta \rangle \overset{s}{\underset{\Psi}{\longmapsto}}{}_j^i Q'$ where $\widehat{s} = \widehat{\alpha}$, $j \in [0, n]$, $(P', Q') \in B_j$ and $Q' \equiv \langle Y, \Delta \rangle$ applying (31). Thus $(P', \widetilde{T}_j(\langle Y, \Delta \rangle, \Psi, n)) \in \widetilde{\mathfrak{T}}(\{B_i\}_{i=0}^n, \Psi, n)$ and $\widetilde{T}_i(\langle X, \Delta \rangle, \Psi, n) \xrightarrow{s} \widetilde{T}_j(\langle Y, \Delta \rangle, \Psi, n)$ from Lemma 6.
- If $\widetilde{T}_i(\langle X, \Delta \rangle, \Psi, n) \xrightarrow{\alpha} R'$ then $R' \equiv \widetilde{T}_j(\langle Y, \Delta \rangle, \Psi, n)$ for some Y and j (31) and applying Lemma 6 we have $\langle X, \Delta \rangle \overset{\alpha}{\underset{\Psi}{\longmapsto}}{}_j^i \langle Y, \Delta \rangle$. Because $(P, \langle X, \Delta \rangle) \in B_i$, there exists P' and s $(\widehat{s} = \widehat{\alpha})$ such that $P \xrightarrow{s} P'$, $(P', \langle Y, \Delta \rangle) \in B_j$ and finally $(P', \widetilde{T}_j(\langle Y, \Delta \rangle, \Psi, n)) \in \widetilde{\mathfrak{T}}(\{B_i\}_{i=0}^n, \Psi, n)$.

($\Leftarrow$) Let $P \approx \widetilde{T}(\langle X, \Delta \rangle, \Psi, n)$ and $(P, \widetilde{T}(\langle X, \Delta \rangle, \Psi, n)) \in B$ where B is the smallest bisimulation with this property. Applying (31), it is easy to see that there exists $\{B_i\}_{i=0}^n$ such that $B = \widetilde{\mathfrak{T}}(\{B_i\}_{i=0}^n, \Psi, n)$, and it is enough to show that $\{B_i\}_{i=0}^n$ is a conditional must-bisimulation. Let $(P, \langle X, \Delta \rangle) \in B_i$ where $i \in [0, n]$. Then $(P, \widetilde{T}_i(\langle X, \Delta \rangle, \Psi, n)) \in B$ and we have:

- If $P \xrightarrow{\alpha} P'$ then $\widetilde{T}_i(\langle X, \Delta \rangle, \Psi, n) \xrightarrow{s} R'$ where $(P', R') \in B$, $\widehat{s} = \widehat{\alpha}$ and $R' \equiv \widetilde{T}_j(\langle Y, \Delta \rangle, \Psi, n)$ for some Y and j (31). Thus $(P', \langle Y, \Delta \rangle) \in B_j$ and applying Lemma 6 we have $\langle X, \Delta \rangle \overset{s}{\underset{\Psi}{\longmapsto}}{}_j^i \langle Y, \Delta \rangle$.
- If $\langle X, \Delta \rangle \overset{\alpha}{\underset{\Psi}{\longmapsto}}{}_j^i Q'$ then $Q' \equiv \langle Y, \Delta \rangle$ (31) and applying Lemma 6 we have $\widetilde{T}_i(\langle X, \Delta \rangle, \Psi, n) \xrightarrow{\alpha} \widetilde{T}_j(\langle Y, \Delta \rangle, \Psi, n)$. Thus there exists P' and s $(\widehat{s} = \widehat{\alpha})$ such that $P \xrightarrow{s} P'$ and $(P', \widetilde{T}_j(\langle Y, \Delta \rangle, \Psi, n)) \in B$. Then it is enough to note that $(P', \langle Y, \Delta \rangle) \in B_j$. □

This article was processed using the LaTeX macro package with LLNCS style