

Symmetric and Asymmetric Asynchronous Interaction

Rob van Glabbeek

NICTA, Sydney, Australia
University of New South Wales, Sydney, Australia
rvg@cs.stanford.edu

Ursula Goltz

Institute for Programming and Reactive Systems
TU Braunschweig, Braunschweig, Germany
goltz@ips.cs.tu-bs.de

Jens-Wolfhard Schicke*

Institute for Programming and Reactive Systems
TU Braunschweig, Braunschweig, Germany
drahflow@gmx.de

We investigate classes of systems based on different interaction patterns with the aim of achieving distributability. As our system model we use Petri nets. In Petri nets, an inherent concept of simultaneity is built in, since when a transition has more than one preplace, it can be crucial that tokens are removed instantaneously. When modelling a system which is intended to be implemented in a distributed way by a Petri net, this built-in concept of synchronous interaction may be problematic. To investigate the problem we assume that removing tokens from places can no longer be considered as instantaneous. We model this by inserting silent (unobservable) transitions between transitions and their preplaces. We investigate three different patterns for modelling this type of asynchronous interaction. *Full asynchrony* assumes that every removal of a token from a place is time consuming. For *symmetric asynchrony*, tokens are only removed slowly in case of backward branched transitions, hence where the concept of simultaneous removal actually occurs. Finally we consider a more intricate pattern by allowing to remove tokens from preplaces of backward branched transitions asynchronously in sequence (*asymmetric asynchrony*).

We investigate the effect of these different transformations of instantaneous interaction into asynchronous interaction patterns by comparing the behaviours of nets before and after insertion of the silent transitions. We exhibit for which classes of Petri nets we obtain equivalent behaviour with respect to failures equivalence.

It turns out that the resulting hierarchy of Petri net classes can be described by semi-structural properties. In case of full asynchrony and symmetric asynchrony, we obtain precise characterisations; for asymmetric asynchrony we obtain lower and upper bounds.

We briefly comment on possible applications of our results to Message Sequence Charts.

1 Introduction

In this paper, we investigate classes of systems based on different asynchronous interaction patterns with the aim of achieving distributability, i.e. the possibility to execute a system on spatially distributed locations, which do not share a common clock. As our system model we use Petri nets. The main reason for this choice is the detailed way in which a Petri net represents a concurrent system, including the interaction between the components it may consist of. In an interleaving based model of concurrency such as labelled transition systems modulo bisimulation semantics, a system representation as such cannot be said to display synchronous or asynchronous interaction; at best these are properties of composition operators, or communication primitives, defined in terms of such a model. A Petri net on the other hand displays enough detail of a concurrent system to make the presence of synchronous communication discernible. This makes it possible to study asynchronous communication without digressing to the realm of composition operators.

*This paper was partially written during a four month stay of J.-W. Schicke at NICTA, during which he was supported by DAAD (Deutscher Akademischer Austauschdienst) and NICTA.




Figure 1: Transformation to the symmetrically asynchronous implementation

In a Petri net, a transition interacts with its preplaces by consuming tokens. An inherent concept of simultaneity is built in, since when a transition has more than one preplace, it can be crucial that tokens are removed instantaneously, depending on the surrounding structure or—more elaborately—the behaviour of the net.

When modelling a distributed system by a Petri net, this built-in concept of synchronous interaction may become problematic. Assume a transition t on a location l models an activity involving another location l' , for example by receiving a message. This can be modelled by a preplace s of t such that s and t are situated in different locations. We assume that taking a token can in this situation not be considered as instantaneous; rather the interaction between s and t takes time. We model this effect by inserting silent (unobservable) transitions between transitions and their preplaces. We call the effect of such a transformation of a net N an *asynchronous implementation* of N .

An example of such an implementation is shown in Figure 1. Note that a can be disabled in the implementation before any visible behaviour has taken place. This difference will cause non-equivalence between the original and the implementation under branching time equivalences.

Our asynchronous implementation allows a token to start its journey from a place to a transition even when not all preplaces of the transition contain a token. This design decision is motivated by the observation that it is fundamentally impossible to check in an asynchronous way whether all preplaces of a transition are marked—it could be that a token moves back and forth between two such places.

We investigate different interaction patterns for the asynchronous implementation of nets. The simplest pattern (*full asynchrony*) assumes that every removal of a token from a place is time consuming. For the next pattern (*symmetric asynchrony*), tokens are only removed slowly when they are consumed by a backward branched transition, hence where the concept of simultaneous removal actually occurs. Finally we consider a more intricate pattern by allowing to remove tokens from preplaces of backward branched transitions asynchronously in sequence (*asymmetric asynchrony*).

Given a choice of interaction pattern, we call a net N *asynchronous* when there is no essential behavioural difference between N and its asynchronous implementation $I(N)$. In order to formally define this concept, we wish to compare the behaviours of N and $I(N)$ using a semantic equivalence that fully preserves branching time, causality and their interplay, whilst of course abstracting from silent transitions. By choosing the most discriminating equivalence possible, we obtain the smallest possible class of asynchronous nets, thus excluding nets that might be classified as asynchronous merely because a less discriminating equivalence would fail to see the differences between such a net and its asynchronous implementation. To simplify the exposition, here we merely compare the behaviours of N and $I(N)$ up to *failures equivalence* [6]. This interleaving equivalence abstracts from causality and respects branching time only to some degree. However, we conjecture that our results are in fact largely independent of this choice and that more discriminating equivalences, such as the history preserving ST-bisimulation of [20],

would yield the same classes of asynchronous nets. Using a linear time equivalence would give rise to larger classes; this possibility is investigated in [18].

Thus we investigate the effect of our three transformations of instantaneous interaction into asynchronous interaction patterns by comparing the behaviours of nets before and after insertion of the silent transitions up to failures equivalence. We show that in the case of full asynchrony, we obtain equivalent behaviour exactly for conflict-free Petri nets. Further we establish that symmetric asynchrony is a valid concept for N -free Petri nets and asymmetric asynchrony for M -free Petri nets, where N and M stand for certain structural properties; the reachability of such structures is crucial. For symmetric asynchrony we obtain a precise characterisation of the class of nets which is asynchronously implementable. For asymmetric asynchrony we obtain lower and upper bounds.

In the concluding section, we discuss the use of our results for Message Sequence Charts, as an example how they may be useful for other models than Petri nets. When interpreting basic Message Sequence Chart as Petri nets, the resulting Petri nets lie within the class of conflict-free and hence N -free Petri nets. The more expressive classes give insights in the effect of choices in non-basic MSCs.

The paper is structured as follows. In Section 2 we establish the necessary basic notions. In Section 3 we introduce the fully asynchronous transformation and give a semi-structural characterisation of the resulting net class. In Section 4 we repeat those steps for the symmetrically asynchronous transformation. Furthermore we describe how the resulting net class relates to the classes of free-choice and extended free choice nets. In Section 5 we introduce the asymmetrically asynchronous transformation. We give semi-structural upper and lower bounds for the resulting net class and relate it to simple and extended simple nets. In the conclusion in Section 6 we compare our findings to similar results in the literature.

An extended abstract of this paper will be presented at the first *Interaction and Concurrency Experience (ICE'08) on Synchronous and Asynchronous Interactions in Concurrent Distributed Systems*, and will appear in *Electronic Notes in Theoretical Computer Science*, Elsevier.

2 Basic Notions

We consider here 1-safe net systems, i.e. places never carry more than one token, but a transition can fire even if pre- and postset intersect. To represent unobservable behaviour, which we use to model asynchrony, the set of transitions is partitioned into observable and silent (unobservable) ones.

Definition 2.1

A net with silent transitions is a tuple $N = (S, O, U, F, M_0)$ where

- S is a set (of places),
- O is a set (of observable transitions),
- U is a set (of silent transitions),
- $F \subseteq S \times T \cup T \times S$ (the flow relation) with $T := O \cup U$ (transitions) and
- $M_0 \subseteq S$ (the initial marking).

Petri nets are depicted by drawing the places as circles, the transitions as boxes, and the flow relation as arrows (*arcs*) between them. When a Petri net represents a concurrent system, a global state of such a system is given as a *marking*, a set of places, the initial state being M_0 . A marking is depicted by

placing a dot (*token*) in each of its places. The dynamic behaviour of the represented system is defined by describing the possible moves between markings. A marking M may evolve into a marking M' when a nonempty set of transitions G fires. In that case, for each arc $(s, t) \in F$ leading to a transition t in G , a token moves along that arc from s to t . Naturally, this can happen only if all these tokens are available in M in the first place. These tokens are consumed by the firing, but also new tokens are created, namely one for every outgoing arc of a transition in G . These end up in the places at the end of those arcs. A problem occurs when as a result of firing G multiple tokens end up in the same place. In that case M' would not be a marking as defined above. In this paper we restrict attention to nets in which this never happens. Such nets are called *1-safe*. Unfortunately, in order to formally define this class of nets, we first need to correctly define the firing rule without assuming 1-safety. Below we do this by forbidding the firing of sets of transitions when this might put multiple tokens in the same place.

Definition 2.2 Let $N = (S, O, U, F, M_0)$ be a net. Let $M_1, M_2 \subseteq S$.

We denote the preset and postset of a net element x by $\bullet x := \{y \mid (y, x) \in F\}$ and $x^\bullet := \{y \mid (x, y) \in F\}$ respectively. A nonempty set of transitions $G \subseteq (O \cup U), G \neq \emptyset$, is called a *step from M_1 to M_2* , notation $M_1 [G]_N M_2$, iff

- all transitions contained in G are *enabled*, that is

$$\forall t \in G. \bullet t \subseteq M_1 \wedge (M_1 \setminus \bullet t) \cap t^\bullet = \emptyset,$$

- all transitions of G are *independent*, that is *not conflicting*:

$$\forall t, u \in G, t \neq u. \bullet t \cap \bullet u = \emptyset \wedge t^\bullet \cap u^\bullet = \emptyset,$$

- in M_2 all tokens have been removed from the *preplaces* of G and new tokens have been inserted at the *postplaces* of G :

$$M_2 = \left(M_1 \setminus \bigcup_{t \in G} \bullet t \right) \cup \bigcup_{t \in G} t^\bullet.$$

To simplify statements about possible behaviours of nets, we use some abbreviations.

Definition 2.3 Let $N = (S, O, U, F, M_0)$ be a net with silent transitions.

- $\longrightarrow_N \subseteq \mathcal{P}(S) \times \mathcal{P}(O) \times \mathcal{P}(S)$ is defined by $M_1 \xrightarrow{G}_N M_2 \Leftrightarrow G \subseteq O \wedge M_1 [G]_N M_2$

- $\xrightarrow{\tau}_N \subseteq \mathcal{P}(S) \times \mathcal{P}(S)$ is defined by $M_1 \xrightarrow{\tau}_N M_2 \Leftrightarrow \exists t \in U. M_1 [\{t\}]_N M_2$

- $\Longrightarrow_N \subseteq \mathcal{P}(S) \times O^* \times \mathcal{P}(S)$ is defined by $M_1 \xrightarrow{t_1 t_2 \dots t_n}_N M_2 \Leftrightarrow$

$$M_1 \xrightarrow{\tau}_N^* \xrightarrow{\{t_1\}}_N \xrightarrow{\tau}_N^* \xrightarrow{\{t_2\}}_N \xrightarrow{\tau}_N^* \dots \xrightarrow{\tau}_N^* \xrightarrow{\{t_n\}}_N \xrightarrow{\tau}_N^* M_2$$

where $\xrightarrow{\tau}_N^*$ denotes the reflexive and transitive closure of $\xrightarrow{\tau}_N$.

We write $M_1 \xrightarrow{G}_N$ for $\exists M_2. M_1 \xrightarrow{G}_N M_2$, $M_1 \not\xrightarrow{G}_N$ for $\nexists M_2. M_1 \xrightarrow{G}_N M_2$ and similar for the other two relations.

A marking M_1 is said to be *reachable* iff there is a $\sigma \in O^*$ such that $M_0 \xrightarrow{\sigma} M_1$. The set of all reachable markings is denoted by $[M_0]_N$.

We omit the subscript N if clear from context.

As said before, here we only want to consider 1-safe nets. Formally, we restrict ourselves to *contact-free nets* where in every reachable marking $M_1 \in [M_0]$ for all $t \in O \cup U$ with $\bullet t \subseteq M_1$

$$(M_1 \setminus \bullet t) \cap t^\bullet = \emptyset.$$

For such nets, in Definition 2.2 we can just as well consider a transition t to be enabled in M iff $\bullet t \subseteq M$, and two transitions to be independent when $\bullet t \cap \bullet u = \emptyset$. In this paper we furthermore restrict attention to nets for which $\bullet t \neq \emptyset$, and $\bullet t$ and t^\bullet are finite for all $t \in O \cup U$. We also require the initial marking M_0 to be finite. A consequence of these restrictions is that all reachable markings are finite, and it can never happen that infinitely many independent transitions are enabled. Henceforth, we employ the name τ -nets for nets with silent transitions obeying the above restrictions, and *plain nets* for τ -nets without silent transitions, i.e. with $U = \emptyset$.

Plain nets have the nice property of being deterministic, i.e. the marking obtained after firing a sequence of transitions is uniquely determined by the sequence of transitions fired.

Lemma 2.1 Let $N = (S, O, \emptyset, F, M_0)$ be a plain net, $\sigma \in O^*$ and $M \subseteq S$.

If $M \xrightarrow{\sigma} M_1 \wedge M \xrightarrow{\sigma} M_2$ then $M_1 = M_2$.

Proof Let $t \in O$, $\sigma \in O^*$ and $M' \subseteq S$.

Then $M \xrightarrow{\{t\}} M' \Leftrightarrow M \xrightarrow{\{t\}} M'$ and $M \xrightarrow{\{t\}} M'$ implies $M' = (M \setminus \bullet t) \cup t^\bullet$.

Hence $M \xrightarrow{\{t\}} M_1 \wedge M \xrightarrow{\{t\}} M_2$ implies $M_1 = M_2$.

The result follows for a trace σ by induction on the length of σ . □

Our nets with silent transitions can be regarded as special *labelled nets*, defined as in Definition 2.1, but without the split of T into O and U , and instead equipped with a *labelling function* $\ell : T \rightarrow Act \cup \{\tau\}$, where Act is a set of *visible actions* and $\tau \notin Act$ an invisible one. Nets with silent transitions correspond to labelled nets in which no two different transitions are labelled by the same visible actions, which can be formalised by taking $\ell(t) = t$ for $t \in O$ and $\ell(t) = \tau$ for $t \in U$.

To describe which nets are ‘‘asynchronous’’, we will compare their behaviour to that of their asynchronous implementations using a suitable equivalence relation. As explained in the introduction, we consider here branching time semantics. Technically, we use failures equivalence, as defined below.

Definition 2.4 Let $N = (S, O, U, F, M_0)$ be a τ -net, $\sigma \in O^*$ and $X \subseteq O$.

$\langle \sigma, X \rangle$ is a *failure pair* of N iff

$$\exists M_1. M_0 \xrightarrow{\sigma} M_1 \wedge M_1 \not\xrightarrow{\tau} \wedge \forall t \in X. M_1 \not\xrightarrow{\{t\}}.$$

We define $\mathcal{F}(N) := \{ \langle \sigma, X \rangle \mid \langle \sigma, X \rangle \text{ is a failure pair of } N \}$.

Two τ -nets N and N' are *failures equivalent*, $N \approx_{\mathcal{F}} N'$, iff $\mathcal{F}(N) = \mathcal{F}(N')$.

A τ -net $N = (S, O, U, F, M_0)$ is called *divergence free* iff there are no infinite chains of markings $M_1 \xrightarrow{\tau} M_2 \xrightarrow{\tau} \dots$ with $M_1 \in [M_0]$.

3 Full Asynchrony

As explained in the introduction, we will examine in this paper different possible assumptions of how asynchronous interaction between transitions and their preplaces takes place. In this section, we start with the simple and intuitive assumption that the removal of any token by a transition takes time. This is implemented by inserting silent transitions between visible ones and their preplaces.

Definition 3.1 Let $N = (S, O, \emptyset, F, M_0)$ be a plain net.

The *fully asynchronous implementation* of N is defined as the net $FI(N) := (S \cup S^\tau, O, U', F', M_0)$ with

$$\begin{aligned} S^\tau &:= \{s_t \mid t \in O, s \in \bullet t\}, \\ U' &:= \{t_s \mid t \in O, s \in \bullet t\} \text{ and} \\ F' &:= (F \cap (O \times S)) \cup \{(s, t_s), (t_s, s_t), (s_t, t) \mid t \in O, s \in \bullet t\}. \end{aligned}$$

For better readability we will use the abbreviations ${}^\circ x := \{y \mid (y, x) \in F'\}$ and $x^\circ := \{y \mid (x, y) \in F'\}$ instead of $\bullet x$ or x^\bullet when making assertions about the flow relation of an implementation.

The following lemma shows how the fully asynchronous implementation of a plain net N simulates the behaviour of N .

Lemma 3.1 Let $N = (S, O, \emptyset, F, M_0)$ be a plain net, $G \subseteq O$, $\sigma \in O^*$ and $M_1, M_2 \subseteq S$.

1. If $M_1 \xrightarrow{G}_N M_2$ then $M_1 \xrightarrow{\tau^*}_{FI(N)} \xrightarrow{G}_{FI(N)} M_2$.
2. If $M_1 \xrightarrow{\sigma}_N M_2$ then $M_1 \xrightarrow{\sigma}_{FI(N)} M_2$.

Proof Assume $M_1 \xrightarrow{G}_N M_2$. Then, by construction of $FI(N)$,

$$M_1 [\{t_s \mid t \in G, s \in \bullet t\}]_{FI(N)} [\{t \mid t \in G\}]_{FI(N)} M_2.$$

The first part of that execution can be split into a sequence of singletons.

The second statement follows by a straightforward induction on the length of σ . □

This lemma uses the fact that any marking of N is also a marking on $FI(N)$. The reverse does not hold, so in order to describe the degree to which the behaviour of $FI(N)$ is simulated by N we need to explicitly relate markings of $FI(N)$ to those of N . This is in fact not so hard, as any reachable marking of $FI(N)$ can be obtained from a reachable marking of N by moving some tokens into the newly introduced buffering places s_t . To establish this formally, we define a function which transforms implementation markings into the related original markings, by shifting these tokens back.

Definition 3.2 Let $N = (S, O, \emptyset, F, M_0)$ be a plain net and let $FI(N) = (S \cup S^\tau, O, U', F', M_0)$.

$\tau^\leftarrow : S \cup S^\tau \rightarrow S$ is the function defined by

$$\tau^\leftarrow(p) := \begin{cases} s & \text{iff } p = s_t \text{ with } s_t \in S^\tau, s \in S, t \in O \\ p & \text{otherwise } (p \in S) \end{cases}$$

Where necessary we extend functions to sets elementwise. So for any $M \subseteq S \cup S^\tau$ we have $\tau^\leftarrow(M) = \{\tau^\leftarrow(s) \mid s \in M\} = (M \cap S) \cup \{s \mid s_t \in M\}$. In particular, $\tau^\leftarrow(M) = M$ when $M \subseteq S$.

We now introduce a predicate α on the markings of $FI(N)$ that holds for a marking iff it can be obtained from a reachable marking of N (which is also a marking of $FI(N)$) by firing some unobservable transitions. Each of these unobservable transitions moves a token from a place s into a buffering place s_t . Later, we will show that α exactly characterises the reachable markings of $FI(N)$. Furthermore, as every token can be moved only once, we can also give an upper bound on how many such movements can still take place.

Definition 3.3 Let $N = (S, O, \emptyset, F, M_0)$ be a plain net and $FI(N) = (S \cup S^\tau, O, U', F', M_0)$.

The predicate $\alpha \subseteq \mathcal{P}(S \cup S^\tau)$ is given by

$$\alpha(M) \quad :\Leftrightarrow \quad \tau^\leftarrow(M) \in [M_0]_N \wedge \forall p, q \in M. \tau^\leftarrow(p) = \tau^\leftarrow(q) \Rightarrow p = q.$$

The function $d : \mathcal{P}(S \cup S^\tau) \rightarrow \mathbb{N} \cup \{\infty\}$ is given by $d(M) := |M \cap \{s \mid s \in S, s^\bullet \neq \emptyset\}|$, where we choose not to distinguish between different degrees of infinity.

Note that $\alpha(M)$ implies $|M| = |\tau^\leftarrow(M)|$, and reachable markings of N are always finite (thanks to our definition of a plain net). Hence $\alpha(M)$ implies $d(M) \in \mathbb{N}$. The following lemma confirms that our informal description of α matches its formal definition.

Lemma 3.2 Let N and $FI(N)$ be as above and $M \subseteq S \cup S^\tau$, with M finite.

Then $\forall p, q \in M. \tau^\leftarrow(p) = \tau^\leftarrow(q) \Rightarrow p = q$ iff $\tau^\leftarrow(M) \xrightarrow{\tau}_{FI(N)}^* M$.

Proof Given that $\tau^\leftarrow(M) \subseteq S$, “if” follows directly from the construction of $FI(N)$.

For “only if”, assume $\forall p, q \in M. \tau^\leftarrow(p) = \tau^\leftarrow(q) \Rightarrow p = q$. Then $\tau^\leftarrow(M) [\{t_s \mid s_t \in M\}]_{FI(N)} M$. \square

Now we can describe how any net simulates the behaviour of its fully asynchronous implementation.

Lemma 3.3 Let N and $FI(N)$ be as above, $G \subseteq O$, $\sigma \in O^*$ and $M, M' \subseteq S \cup S^\tau$.

1. $\alpha(M_0)$.
2. If $\alpha(M) \wedge M \xrightarrow{G}_{FI(N)} M'$ then $\tau^\leftarrow(M) \xrightarrow{G}_N \tau^\leftarrow(M') \wedge \alpha(M')$.
3. If $\alpha(M) \wedge M \xrightarrow{\tau}_{FI(N)} M'$ then $d(M) > d(M') \wedge \tau^\leftarrow(M) = \tau^\leftarrow(M') \wedge \alpha(M')$.
4. If $M_0 \xrightarrow{\sigma}_{FI(N)} M'$ then $M_0 \xrightarrow{\sigma}_N \tau^\leftarrow(M') \wedge \alpha(M')$.

Proof (1): $M_0 \in [M_0]_N$ and $\forall s \in M_0 \subseteq S. \tau^\leftarrow(s) = s$.

(2): Suppose $\alpha(M)$ and $M \xrightarrow{G}_{FI(N)} M'$ with $G \subseteq O$. So $\tau^\leftarrow(M)$ is a reachable marking of N .

Let $t \in G$. Since t is enabled in M , we have ${}^\circ t \subseteq M$ and hence $\tau^\leftarrow({}^\circ t) \subseteq \tau^\leftarrow(M)$. By construction, ${}^\circ t = \{s_t \mid s \in {}^\bullet t\}$ so $\tau^\leftarrow({}^\circ t) = {}^\bullet t$. Given that N is contact-free, it follows that t is enabled in $\tau^\leftarrow(M)$.

Now let $t, u \in G$ with $t \neq u$. If $s \in {}^\bullet t \cap {}^\bullet u$ then $s_t \in {}^\circ t$ and $s_u \in {}^\circ u$, so $s_t, s_u \in M$. However, $\tau^\leftarrow(s_t) = \tau^\leftarrow(s_u)$, contradicting $\alpha(M)$. Hence ${}^\bullet t \cap {}^\bullet u = \emptyset$. Given that ${}^\bullet t \cup {}^\bullet u \subseteq \tau^\leftarrow(M)$ and N is contact-free, it follows that also $t^\bullet \cap u^\bullet = \emptyset$ and hence t and u are independent.

We will now show that $\left(\tau^{\leftarrow}(M) \setminus \bigcup_{t \in G} \bullet t \right) \cup \bigcup_{t \in G} t^{\bullet} = \tau^{\leftarrow}(M')$.

$$\begin{aligned} M' &= (M \setminus \{s \mid s \in {}^{\circ}t, t \in G\}) \cup \{s \mid s \in t^{\circ}, t \in G\} \\ &= (M \setminus \{s_t \mid s \in \bullet t, t \in G\}) \cup \{s \mid s \in t^{\bullet}, t \in G\}. \end{aligned}$$

Therefore $\tau^{\leftarrow}(M') = \tau^{\leftarrow}(M \setminus \{s_t \mid s \in \bullet t, t \in G\}) \cup \tau^{\leftarrow}(\{s \mid s \in t^{\bullet}, t \in G\})$.

Take any $t \in G$ and any $s \in \bullet t$. Then $s_t \in M$ and $\alpha(M)$ implies $s \notin M \wedge \nexists u \in O. u \neq t \wedge s_u \in M$. Hence $\tau^{\leftarrow}(M \setminus \{s_t \mid s \in \bullet t, t \in G\}) = \tau^{\leftarrow}(M) \setminus \{s \mid s \in \bullet t, t \in G\}$. Thus we find

$$\tau^{\leftarrow}(M') = \tau^{\leftarrow}(M) \setminus \{s \mid s \in \bullet t, t \in G\} \cup \{s \mid s \in t^{\bullet}, t \in G\}$$

and conclude that $\tau^{\leftarrow}(M) \xrightarrow{G}_N \tau^{\leftarrow}(M')$.

Next we establish $\alpha(M')$. To this end, we may assume that G is a singleton set, for G must be finite—this follows from our definition of a plain net—and when $M[\{t_0, t_1, \dots, t_n\}]M'$ for some $n \geq 0$ then there are M_1, M_2, \dots, M_n with $M[\{t_0\}]M_1[\{t_1\}]M_2 \cdots M_n[\{t_n\}]M'$, allowing us to obtain the general case by induction. So let $G = \{t\}$ with $t \in T$.

Above we have shown that $\tau^{\leftarrow}(M') \in [M_0]_N$. We still need to prove that $\forall p, q \in M'. p \neq q \Rightarrow \tau^{\leftarrow}(p) \neq \tau^{\leftarrow}(q)$. Assume the contrary, i.e. there are $p, q \in M'$ with $p \neq q \wedge \tau^{\leftarrow}(p) = \tau^{\leftarrow}(q)$. Since $\alpha(M)$, at least one of p and q —say p —must not be present in M . Then $p \in t^{\circ} = t^{\bullet} \subseteq S$. As $\tau^{\leftarrow}(q) = \tau^{\leftarrow}(p) = p$ and $q \neq p$, it must be that $q \in S^{\tau}$. Hence $q \notin t^{\circ}$, so $q \in M$, and $p = \tau^{\leftarrow}(q) \in \tau^{\leftarrow}(M)$. As shown above, t is enabled in $\tau^{\leftarrow}(M)$. By the contact-freeness of N , $(\tau^{\leftarrow}(M) \setminus \bullet t) \cap t^{\bullet} = \emptyset$, so $p \in \bullet t$. Hence $p_t \in {}^{\circ}t \subseteq M$. As by construction ${}^{\circ}t \cap t^{\circ} = \emptyset$, we have $p_t \notin M'$, so $q \neq p_t$. Yet $\tau^{\leftarrow}(q) = \tau^{\leftarrow}(p_t)$, contradicting $\alpha(M)$.

(3): Let $t_s \in U'$ such that $M[\{t_s\}]_{FI(N)}M'$. Then, by construction of $FI(N)$, ${}^{\circ}t_s = \{s\} \wedge t_s^{\circ} = \{s_t\}$. Hence $M' = M \setminus \{s\} \cup \{s_t\}$ and $d(M') = d(M) - 1 \wedge \tau^{\leftarrow}(M') = \tau^{\leftarrow}(M)$. Moreover, $\alpha(M') \Leftrightarrow \alpha(M)$.

(4): Using (1–3), this follows by a straightforward induction on the number of transitions in the derivation $M_0 \xrightarrow{\sigma}_{FI(N)} M'$. \square

It follows that α exactly characterises the reachable markings of $FI(N)$. Using this it is not hard to check that implementations of contact-free nets are contact-free, and hence τ -nets.

Proposition 3.1 Let N and $FI(N)$ be as before and $M \subseteq S \cup S^{\tau}$.

1. $M \in [M_0]_{FI(N)}$ iff $\alpha(M)$.
2. $FI(N)$ is contact-free.
3. $FI(N)$ is a τ -net.

Proof (1): “Only if” follows from Lemma 3.3(4), and “if” follows by Lemmas 3.1 and 3.2.

(2): Let $M \in [M_0]_{FI(N)}$. Then $\alpha(M)$, and hence $\tau^{\leftarrow}(M) \in [M_0]_N$.

Consider any $t \in O$ with ${}^{\circ}t \subseteq M$. Assume $(M \setminus {}^{\circ}t) \cap t^{\circ} \neq \emptyset$. Since $t^{\circ} = t^{\bullet} \subseteq S$ let $p \in S$ be such that $p \in M \cap t^{\circ}$ and $p \notin {}^{\circ}t$. As N is contact-free we have $(\tau^{\leftarrow}(M) \setminus \bullet t) \cap t^{\bullet} = \emptyset$, so since $p \in \tau^{\leftarrow}(M) \cap t^{\bullet}$ it must be that $p \in \bullet t$. Hence $p_t \in {}^{\circ}t \subseteq M$ and we have $p \neq p_t$ yet $\tau^{\leftarrow}(p) = p = \tau^{\leftarrow}(p_t)$, violating $\alpha(M)$.

Now consider any $t_p \in U'$ with ${}^{\circ}t_p \subseteq M$. As ${}^{\circ}t_p = \{p\}$ and $t_p^{\circ} = \{p_t\}$ we have that $(M \setminus {}^{\circ}t_p) \cap t_p^{\circ} \neq \emptyset$ only if $p \in M \wedge p_t \in M$. However, $\tau^{\leftarrow}(p) = p = \tau^{\leftarrow}(p_t)$ which would violate $\alpha(M)$.

(3): By construction, M_0 is finite, ${}^{\circ}t \neq \emptyset$, and ${}^{\circ}t$ and t° are finite for all $t \in O \cup U'$. \square




Figure 2: A net which is not failures equivalent to its fully asynchronous implementation

By Lemma 3.3(3) implementations are always divergence free:

Proposition 3.2 Let N be a plain net. Then $FI(N)$ is divergence free. \square

Whereas in a plain net N for any sequence of observable transitions $\sigma \in O^*$ there is at most one marking M with $M_0 \xrightarrow{\sigma} M$, in its fully asynchronous implementation $FI(N)$ there can be several such markings. These markings M' differ from M in that some tokens may have wandered off into the added buffer places on the incoming arcs of visible transitions. As a consequence, a visible transition t that is enabled in M need not be enabled in M' —we say that in $FI(N)$ t can be refused after σ . This may occur for instance for the net N of Figure 2, namely with $\sigma = \varepsilon$ (the empty sequence), M the initial marking of N , M' the marking of $FI(N)$ obtained by firing the rightmost invisible transition, and $t = a$.

When this happens, we have $\langle \sigma, \{t\} \rangle \in \mathcal{F}(FI(N)) \setminus \mathcal{F}(N)$, so the nets N and $FI(N)$ are not failures equivalent. The direction from implementation to original is nicer however as every transition enabled in the implementation must also have been enabled in the original net. Hence the only difference in behaviour between original and implementation can consist of additional failures in the implementation.

Proposition 3.3 Let N and $FI(N)$ be as before. Then $\mathcal{F}(N) \subseteq \mathcal{F}(FI(N))$.

Proof Let $\langle \sigma, X \rangle \in \mathcal{F}(N)$. Applying Lemma 2.1, let $M_1 \subseteq S$ be the unique marking of N such that $M_0 \xrightarrow{\sigma}_N M_1$. By Lemma 3.1 also $M_0 \xrightarrow{\sigma}_{FI(N)} M_1$. So $\alpha(M_1)$. As $M_1 \subseteq S$ we have $\tau^{\leftarrow}(M_1) = M_1$. By Proposition 3.2 there exists a marking M_2 with $M_1 \xrightarrow{\varepsilon}_{FI(N)} M_2 \wedge M_2 \xrightarrow{\tau}_{FI(N)}$. Lemma 3.3(3) yields $\tau^{\leftarrow}(M_2) = \tau^{\leftarrow}(M_1) \wedge \alpha(M_2)$.

Suppose $\langle \sigma, X \rangle \notin \mathcal{F}(FI(N))$. Then $M_2 \xrightarrow{\{t\}}_{FI(N)} M_3$ for some $t \in X$ and marking M_3 of $FI(N)$. Lemma 3.3(2) yields $M_1 = \tau^{\leftarrow}(M_1) = \tau^{\leftarrow}(M_2) \xrightarrow{\{t\}} \tau^{\leftarrow}(M_3)$, which is a contradiction. \square

If the wandering off of tokens into τ -transitions never disables a transition that would be enabled otherwise, then there is no essential behavioural difference between N and $FI(N)$, and they are equivalent in any reasonable behavioural equivalence that abstracts from silent transition firings. In that case, N could be called *fully asynchronous*.

Definition 3.4

The class of *fully asynchronous nets respecting branching time equivalence* is defined as $FA(B) := \{N \mid FI(N) \approx_{\mathcal{F}} N\}$.

As for any plain net N we have $\mathcal{F}(N) \subseteq \mathcal{F}(FI(N))$, the class of nets $FA(B)$ can equivalently be defined as $FA(B) := \{N \mid \mathcal{F}(FI(N)) \subseteq \mathcal{F}(N)\}$.

It turns out that there exists a quite structural characterisation of those nets which are failures equivalent to their fully asynchronous implementation.

Definition 3.5 Let $N = (S, O, \emptyset, F, M_0)$ be a plain net

N has a partially reachable conflict iff $\exists t, u \in O. t \neq u \wedge \bullet t \cap \bullet u \neq \emptyset$ and $\exists M \in [M_0]. \bullet t \subseteq M$.

The nets N of Figures 2 and 3, for instance, have a partially reachable conflict.

Theorem 3.1 A plain net N is in $FA(B)$ iff N has no partially reachable conflict.

Proof Let $N = (S, O, \emptyset, F, M_0)$ and $FI(N) = (S \cup S^\tau, O, U', F', M_0)$.

“ \Rightarrow ”: Assume N has a partially reachable conflict. Then there exist $t, u \in O, t \neq u, \sigma \in O^*$ and $M_1 \subseteq S$ such that $M_0 \xrightarrow{\sigma}_N M_1, \bullet t \cap \bullet u \neq \emptyset$ and $\bullet t \subseteq M_1$. By Lemma 2.1 we know that $\langle \sigma, \{t\} \rangle \notin \mathcal{F}(N)$.

On the other hand, $M_0 \xrightarrow{\sigma}_{FI(N)} M_1$ by Lemma 3.1. Let $p \in \bullet t \cap \bullet u$. Then, by construction of $FI(N)$, there exists an $M_2 \subseteq S \cup S^\tau$ with $M_1[\{u_p\}]M_2, p \notin M_2$ and since $t \neq u$ also $p_t \notin M_2$. Now let $M_3 \subseteq S \cup S^\tau$ such that $M_2 \xrightarrow{\tau}_{FI(N)} M_3 \wedge M_3 \xrightarrow{\tau}_{FI(N)}$ (which exists according to Proposition 3.2). Since $\forall v \in U'. p \notin v \bullet \wedge (p_t \in v \bullet \Rightarrow p \in \bullet v)$ we know that $p_t \notin M_3$. Thus $M_3 \xrightarrow{\{t\}}$ and there exists a failure pair $\langle \sigma, \{t\} \rangle \in \mathcal{F}(FI(N))$. Hence $\mathcal{F}(FI(N)) \neq \mathcal{F}(N)$, so $N \notin FA(B)$.

“ \Leftarrow ”: Assume $N \notin FA(B)$. Then $\mathcal{F}(FI(N)) \neq \mathcal{F}(N)$ and hence $\mathcal{F}(FI(N)) \setminus \mathcal{F}(N) \neq \emptyset$ by Proposition 3.3. Let $\langle \sigma, X \rangle \in \mathcal{F}(FI(N)) \setminus \mathcal{F}(N)$. Then there exists an $M_1 \subseteq S \cup S^\tau$ such that $M_0 \xrightarrow{\sigma}_{FI(N)} M_1 \wedge M_1 \xrightarrow{\tau} \wedge \forall t \in X. M_1 \xrightarrow{\{t\}}$. By Lemma 3.3(4) we have $M_0 \xrightarrow{\sigma}_N \tau^{\leftarrow}(M_1)$.

Let $t \in X$ such that $\tau^{\leftarrow}(M_1) \xrightarrow{\{t\}}_N$ (which exists, otherwise $\langle \sigma, X \rangle \in \mathcal{F}(N)$). Let $p \in \bullet t$ such that $p_t \notin M_1$ (such p_t exists, otherwise $M_1 \xrightarrow{\{t\}}_{FI(N)}$). Since $\tau^{\leftarrow}(M_1) \xrightarrow{\{t\}}_N$ it follows that $p \in \tau^{\leftarrow}(M_1)$. But $p \notin M_1$, for otherwise $M_1 \xrightarrow{\tau}_{FI(N)}$, which would be a contradiction. Hence there must exist some $u \in O$ with $p_u \in M_1$ and $u \neq t$. By construction of $FI(N)$ we have $p \in \bullet u$. Thus $t, u \in O \wedge t \neq u \wedge \bullet t \cap \bullet u \neq \emptyset \wedge \tau^{\leftarrow}(M_1) \in [M_0]_N \wedge \bullet t \subseteq \tau^{\leftarrow}(M_1)$ and N has a partially reachable conflict. \square

4 Symmetric Asynchrony


For investigating the next interaction pattern, we change our notion of asynchronous implementation of a net. We only insert silent transitions wherever a transition has multiple preplaces. These are the situations where the synchronous removal of tokens is really essential.

Definition 4.1 Let $N = (S, O, \emptyset, F, M_0)$ be a net. Let $O^b = \{t \mid t \in O, |\bullet t| > 1\}$.

The *symmetrically asynchronous implementation* of N is defined as the net $SI(N) := (S \cup S^\tau, O, U', F', M_0)$ with

$$\begin{aligned} S^\tau &:= \{s_t \mid t \in O^b, s \in \bullet t\}, \\ U' &:= \{t_s \mid t \in O^b, s \in \bullet t\} \text{ and} \\ F' &:= F \cap \left((O \times S) \cup (S \times (O \setminus O^b)) \right) \\ &\quad \cup \{(s, t_s), (t_s, s_t), (s_t, t) \mid t \in O^b, s \in \bullet t\}. \end{aligned}$$

An example is shown in Figure 3.

Figure 3: The transition a can be refused in $SI(N)$ by firing the left τ .

Similar to Section 3, we use ${}^\circ x$ and x° when describing the flow relation of the implementation.

As Definition 4.1 is only a slight variation of Definition 3.1, the lemmas and propositions about FI in Section 3 apply to SI as well, with minimal changes in the proofs. We will again begin with how the implementation can simulate the original net.

Lemma 4.1 Let $N = (S, O, \emptyset, F, M_0)$ be a plain net, $G \subseteq O$, $\sigma \in O^*$ and $M_1, M_2 \subseteq S$.

1. If $M_1 \xrightarrow{G}_N M_2$ then $M_1 \xrightarrow{\tau^*}_{SI(N)} \xrightarrow{G}_{SI(N)} M_2$.
2. If $M_1 \xrightarrow{\sigma}_N M_2$ then $M_1 \xrightarrow{\sigma}_{SI(N)} M_2$.

Proof Let $O^b = \{t \mid t \in O, |\bullet t| > 1\}$. Assume $M_1 \xrightarrow{G}_N M_2$. Then, by construction of $SI(N)$,

$$M_1 [\{t_s \mid t \in G \cap O^b, s \in \bullet t\}]_{SI(N)} [\{t \mid t \in G\}]_{SI(N)} M_2.$$

The rest of the proof is identical to the proof of Lemma 3.1. □

Also similar to the fully asynchronous case, we wish to undo the effect of firing extraneous τ -transitions. The function doing so is the same τ^\leftarrow defined earlier. We also reuse the predicate α and the distance function d . However, $d(M)$ is no longer a *strict* upper bound, or exact measure, on the number of silent transitions that need to be fired from the marking M before no further silent transitions are possible. Optionally, strictness can be ensured by replacing it by the function e , defined by

$$e(M) := |M \cap \{s \mid s \in S, \exists t \in s^\bullet. |\bullet t| > 1\}|.$$

Again $\alpha(M)$ implies $d(M) \in \mathbb{N}$.

Lemma 4.2 Let N and $SI(N)$ be as above and $M \subseteq S \cup S^\tau$, with M finite.

$$\text{Then } \forall p, q \in M. \tau^\leftarrow(p) = \tau^\leftarrow(q) \Rightarrow p = q \text{ iff } \tau^\leftarrow(M) \xrightarrow{\tau^*}_{SI(N)} M.$$

Proof This is Lemma 3.2 applied to $SI(N)$ rather than $FI(N)$. The proof is identical. □

Lemma 4.3 Let N and $SI(N)$ be as above, $G \subseteq O$, $\sigma \in O^*$ and $M, M' \subseteq S \cup S^\tau$.

1. $\alpha(M_0)$.
2. If $\alpha(M) \wedge M \xrightarrow{G}_{SI(N)} M'$ then $\tau^\leftarrow(M) \xrightarrow{G}_N \tau^\leftarrow(M') \wedge \alpha(M')$.
3. If $\alpha(M) \wedge M \xrightarrow{\tau}_{SI(N)} M'$ then $d(M) > d(M') \wedge \tau^\leftarrow(M) = \tau^\leftarrow(M') \wedge \alpha(M')$.
4. If $M_0 \xrightarrow{\sigma}_{SI(N)} M'$ then $M_0 \xrightarrow{\sigma}_N \tau^\leftarrow(M') \wedge \alpha(M')$.

Proof This is Lemma 3.3 applied to $SI(N)$ rather than $FI(N)$; the proofs of (1), (3) and (4) are identical.

(2): Suppose $\alpha(M)$ and $M \xrightarrow{G}_{SI(N)} M'$ with $G \subseteq O$. So $\tau^{\leftarrow}(M)$ is a reachable marking of N .

For any $t \in O$ and $s \in \bullet t$ we set $\hat{s}_t := s_t$ if $|\bullet t| > 1$ and $\hat{s}_t := s$ otherwise.

Let $t \in G$. Since t is enabled in M , we have ${}^\circ t \subseteq M$ and hence $\tau^{\leftarrow}({}^\circ t) \subseteq \tau^{\leftarrow}(M)$. By construction, ${}^\circ t = \{\hat{s}_t \mid s \in \bullet t\}$ so $\tau^{\leftarrow}({}^\circ t) = \bullet t$. Given that N is contact-free, it follows that t is enabled in $\tau^{\leftarrow}(M)$.

Now let $t, u \in G$ with $t \neq u$. If $s \in \bullet t \cap \bullet u$ then $\hat{s}_t \in {}^\circ t$ and $\hat{s}_u \in {}^\circ u$, so $\hat{s}_t, \hat{s}_u \in M$. As t and u are independent in $SI(N)$, we have $\hat{s}_t \neq \hat{s}_u$. However, $\tau^{\leftarrow}(\hat{s}_t) = s = \tau^{\leftarrow}(\hat{s}_u)$, contradicting $\alpha(M)$. Hence $\bullet t \cap \bullet u = \emptyset$. Given that $\bullet t \cup \bullet u \subseteq \tau^{\leftarrow}(M)$ and N is contact-free, it follows that also $t^\bullet \cap u^\bullet = \emptyset$ and hence t and u are independent.

We will now show that $\left(\tau^{\leftarrow}(M) \setminus \bigcup_{t \in G} \bullet t \right) \cup \bigcup_{t \in G} t^\bullet = \tau^{\leftarrow}(M')$.

$$\begin{aligned} M' &= (M \setminus \{s \mid s \in {}^\circ t, t \in G\}) \cup \{s \mid s \in t^\bullet, t \in G\} \\ &= (M \setminus \{\hat{s}_t \mid s \in \bullet t, t \in G\}) \cup \{s \mid s \in t^\bullet, t \in G\}. \end{aligned}$$

Therefore $\tau^{\leftarrow}(M') = \tau^{\leftarrow}(M \setminus \{\hat{s}_t \mid s \in \bullet t, t \in G\}) \cup \tau^{\leftarrow}(\{s \mid s \in t^\bullet, t \in G\})$.

Take any $t \in G$ and any $s \in \bullet t$. Then $\hat{s}_t \in M$, $\tau^{\leftarrow}(\hat{s}_t) = s$ and $\alpha(M)$ implies $\nexists p \in M. p \neq \hat{s}_t \wedge \tau^{\leftarrow}(p) = s$. Hence $\tau^{\leftarrow}(M \setminus \{\hat{s}_t \mid s \in \bullet t, t \in G\}) = \tau^{\leftarrow}(M) \setminus \{s \mid s \in \bullet t, t \in G\}$. Thus we find

$$\tau^{\leftarrow}(M') = \tau^{\leftarrow}(M) \setminus \{s \mid s \in \bullet t, t \in G\} \cup \{s \mid s \in t^\bullet, t \in G\}$$

and conclude that $\tau^{\leftarrow}(M) \xrightarrow{G}_N \tau^{\leftarrow}(M')$.

That $\alpha(M')$ holds is established in exactly the same way as in the proof of Lemma 3.3(2), noting that in deriving $p_t \in {}^\circ t \subseteq M$ we use $\hat{p}_t \in {}^\circ t \subseteq M$ and $p \notin M$. \square

Proposition 4.1 Let N and $SI(N)$ be as before and $M \subseteq S \cup S^\tau$.

1. $M \in [M_0]_{SI(N)}$ iff $\alpha(M)$.
2. $SI(N)$ is contact-free.
3. $SI(N)$ is a τ -net.

Proof Identical to the proof of Proposition 3.1, using the lemmas of Section 4. \square

Proposition 4.2 Let N be a plain net. Then $SI(N)$ is divergence free.

Proof This follows immediately from Lemma 4.3(3). \square

Proposition 4.3 Let N and $SI(N)$ be as before. Then $\mathcal{F}(N) \subseteq \mathcal{F}(SI(N))$.

Proof Identical to that of Proposition 3.3, using the lemmas of Section 4. \square

Again, the only difference in behaviour between the original net and its implementation is that observable transitions can potentially be refused in the implementation, as in Figure 3. This yields a concept of a *symmetrically asynchronous net*.

Definition 4.2

The class of *symmetrically asynchronous nets respecting branching time equivalence* is defined as $SA(B) := \{N \mid SI(N) \approx_{\mathcal{F}} N\}$.

We now show that plain nets can be implemented symmetrically asynchronously with respect to failure equivalence exactly when they do not contain reachable structures of the form shown in Figure 3.

Definition 4.3

A plain net $N = (S, O, \emptyset, F, M_0)$ has a *partially reachable N* iff $\exists t, u \in O. t \neq u \wedge \bullet t \cap \bullet u \neq \emptyset \wedge |\bullet u| > 1 \wedge \exists M \in [M_0]_N. \bullet t \subseteq M \vee \bullet u \subseteq M$.

Theorem 4.1 A plain net N is in $SA(B)$ iff N has no partially reachable N.

Proof Let $N = (S, O, \emptyset, F, M_0)$ and $SI(N) = (S \cup S^\tau, O, U', F', M_0)$.

“ \Rightarrow ”: Assume N has a partially reachable N. Then there exist $t, u \in O, t \neq u, \sigma \in O^*$ and $M_1 \subseteq S$ such that $M_0 \xrightarrow{\sigma}_N M_1, \bullet t \cap \bullet u \neq \emptyset, |\bullet u| > 1$ and $\bullet t \subseteq M_1 \vee \bullet u \subseteq M_1$. We will show that $SI(N) \not\approx_{\mathcal{F}} N$.

There are two cases:

Case 1, $\bullet t \subseteq M$: We will show that $\langle \sigma, \{t\} \rangle \in \mathcal{F}(SI(N))$ but $\langle \sigma, \{t\} \rangle \notin \mathcal{F}(N)$. As N has no silent transitions, by Lemma 2.1 we have $M_0 \xrightarrow{\sigma}_N M'$ only if $M' = M_1$. Since $M_1 \xrightarrow{\{t\}}_N$ it follows that $\langle \sigma, \{t\} \rangle \notin \mathcal{F}(N)$.

On the other hand, $M_0 \xrightarrow{\sigma}_{SI(N)} M_1$ by Lemma 4.1. Let $p \in \bullet t \cap \bullet u$. Then, by construction of $SI(N)$, there exists an $M_2 \subseteq S \cup S^\tau$ with $M_1 \{u_p\} M_2, p \notin M_2$ and since $t \neq u$ also $\hat{p}_t \notin M_2$. Now let $M_3 \subseteq S \cup S^\tau$ such that $M_2 \xrightarrow{\tau}_* M_3 \wedge M_3 \xrightarrow{\tau}_* M_3$ (which exists according to Proposition 4.2). Since $\forall v \in U'. p \notin v \bullet \wedge (p_t \in v \bullet \Rightarrow p \in v \bullet)$ we have $\hat{p}_t \notin M_3$. Thus $M_3 \xrightarrow{\{t\}}$ and $\langle \sigma, \{t\} \rangle \in \mathcal{F}(SI(N))$.

Case 2, $\bullet t \not\subseteq M$: Then $\bullet u \subseteq M$. Thus $\exists q \in \bullet t \setminus \bullet u$, so $|\bullet t| > 1$. This case proceeds as case 1 with the roles of t and u exchanged.

“ \Leftarrow ”: Assume $N \notin SA(B)$. Then $\mathcal{F}(SI(N)) \neq \mathcal{F}(N)$ and hence $\mathcal{F}(SI(N)) \setminus \mathcal{F}(N) \neq \emptyset$ by Proposition 4.3. Let $\langle \sigma, X \rangle \in \mathcal{F}(SI(N)) \setminus \mathcal{F}(N)$. Then there exists an $M_1 \subseteq S \cup S^\tau$ such that $M_0 \xrightarrow{\sigma}_{SI(N)} M_1, M_1 \xrightarrow{\tau}$ and $\forall t \in X. M_1 \xrightarrow{\{t\}}$. By Lemma 4.3(4) we have $M_0 \xrightarrow{\sigma}_N \tau^{\leftarrow}(M_1)$.

Let $t \in X$ such that $\tau^{\leftarrow}(M_1) \xrightarrow{\{t\}}_N$ (which exists, otherwise $\langle \sigma, X \rangle \in \mathcal{F}(N)$). Let $p \in \bullet t$ such that $\hat{p}_t \notin M_1$ (such a p exists, otherwise $M_1 \xrightarrow{\{t\}}_{SI(N)}$). Since $\tau^{\leftarrow}(M_1) \xrightarrow{\{t\}}_N$ it follows that $p \in \tau^{\leftarrow}(M_1)$. But $p \notin M_1$, for otherwise $p \neq \hat{p}_t$ and $M_1 \xrightarrow{\tau}_{SI(N)}$, which would be a contradiction. Hence there must exist some $u \in O$ with $p_u \in M_1$ and $u \neq t$. By construction of $SI(N)$ we have $p \in \bullet u$ and $|\bullet u| > 1$. Thus $t, u \in O \wedge t \neq u \wedge \bullet t \cap \bullet u \neq \emptyset \wedge |\bullet u| > 1 \wedge \tau^{\leftarrow}(M_1) \in [M_0]_N \wedge \bullet t \subseteq \tau^{\leftarrow}(M_1)$, so N has a partially reachable N. \square

The following proposition shows that the current class of nets strictly extends the one from the previous section.

Proposition 4.4 $FA(B) \subsetneq SA(B)$.

Proof A net without partially reachable conflict surely has no partially reachable N. The inequality follows from the example in Figure 2. \square

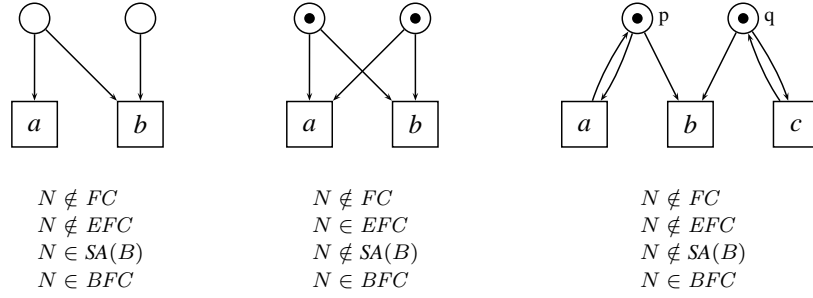


Figure 4: Differences between various classes of free-choice-like nets

It turns out that our class of nets $SA(B)$ is strongly related to the following established net classes [2, 3].

Definition 4.4 Let $N = (S, O, \emptyset, F, M_0)$ be a plain net.

1. N is *free choice*, $N \in FC$, iff $\forall p, q \in S. p \neq q \wedge p^\bullet \cap q^\bullet \neq \emptyset \Rightarrow |p^\bullet| = |q^\bullet| = 1$.
2. N is *extended free choice*, $N \in EFC$, iff $\forall p, q \in S. p^\bullet \cap q^\bullet \neq \emptyset \Rightarrow p^\bullet = q^\bullet$.
3. N is *behaviourally free choice*, $N \in BFC$, iff $\forall u, v \in O. \bullet u \cap \bullet v \neq \emptyset \Rightarrow (\forall M_1 \in [M_0]. \bullet u \subseteq M_1 \Leftrightarrow \bullet v \subseteq M_1)$.

The above definition of a free choice net is in terms of places, but the notion can equivalently be defined in terms of transitions:

$$N \in FC \text{ iff } \forall t, u \in T. t \neq u \wedge \bullet t \cap \bullet u \neq \emptyset \Rightarrow |\bullet t| = |\bullet u| = 1.$$

Both conditions are equivalent to the requirement that N must be N -free, where N is defined as in Definition 4.3 but without the reachability clause. Also the notion of an extended free choice net can equivalently be defined in terms of transitions:

$$N \in EFC \text{ iff } \forall t, u \in T. \bullet t \cap \bullet u \neq \emptyset \Rightarrow \bullet t = \bullet u.$$

This condition says that N may not contain what we call a *pure N*: places p, q and transitions t, u such that $p \in \bullet t \cap \bullet u, q \in \bullet u$ and $q \notin \bullet t$.

In [3] it has been established that $FC \subsetneq EFC \subsetneq BFC$. In fact, the inclusions follow directly from the definitions, and Figure 4 displays counterexamples to strictness.

The class of free choice nets is strictly smaller than the class of symmetrically asynchronous nets respecting branching time equivalence, which in turn is strictly smaller than the class of behavioural free choice nets. The class of extended free choice nets and the class of symmetrically asynchronous nets respecting branching time equivalence are incomparable.

Proposition 4.5 $FC \subsetneq SA(B) \subsetneq BFC, EFC \not\subseteq SA(B)$ and $SA(B) \not\subseteq EFC$.

Proof The first inclusion follows because a partially reachable N is surely an N , and also the second inclusion follows directly from the definitions. The four inequalities follow from the examples in Figure 4. The first net is unmarked and thus trivially in $SA(B)$. The second ones symmetrically asynchronous implementation has the additional failure $\langle \varepsilon, \{a, b\} \rangle$ and hence this net is not in $SA(B)$. \square

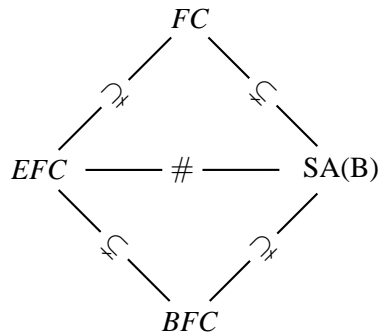


Figure 5: Overview of free-choice-like net classes

In Figure 5 the relations between our semantically defined net class $SA(B)$, the structurally defined classes FC , EFC , and the more behaviourally defined class BFC are summarised. These relations may be interpreted as follows.

Starting at the top of the diagram, free choice nets are characterised structurally, enforcing that for every place, a token therein can choose freely (i.e. without inquiring about the existence of tokens in any other places) which outgoing arc to take. This property makes it possible to implement the system asynchronously. In particular, the component which holds the information represented by a token can choose arbitrarily when and into which of multiple asynchronous output channels to forward said information, without further knowledge about the rest of the system. As this decision is solely in the discretion of the sending component and not based upon any knowledge of the rest of the system, no synchronisation with other components is necessary.

The difference between $SA(B)$ and FC is that in $SA(B)$ the quantification over the places is dropped, making the requirement more straightforward: Every token can choose freely which outgoing arc to follow. Thus, $SA(B)$ allows for non-free-choice structures as long as these never receive any tokens.

This also explains why BFC includes $SA(B)$. Since $SA(B)$ guarantees that all transitions of a problematic structure are never enabled, transitions in such structures are never enabled while others are disabled.

The incomparability between the left and the right side of the diagram stems from the conceptual allowance of slight transformations of the net before evaluating whether it is free choice or not. Extended free choice nets and behavioural free choice nets were proposed as nets that are easily seen to be behaviourally equivalent to free choice nets, and hence share some of their desirable properties: in [2, 3] constructions can be found to turn any extended free choice net into an equivalent free choice net, and any behavioural free choice net into an extended free choice net.¹ Applied on the last two nets in Figure 4 these constructions yield:



Figure 6: Transformed nets from Figure 4




Figure 7: Transformation to asymmetric asynchrony; g such that $p \stackrel{b}{<}_g s \stackrel{b}{<}_g q$.

For the second net of Figure 4, a τ -transition is introduced, which collects both tokens and then marks a single postplace from which the two original transitions are enabled. Hence the choice between the two transitions is centralised in the newly introduced place and thus free again. In the definition of our symmetrically asynchronous implementation *SI*, we do not allow any insertion of such “helping” τ -transitions, as it seems unclear to us how much computing power should be allowed in possibly larger networks of such transitions. This becomes especially problematic if these networks somehow track part of the global status of the net inside themselves and thus make quite informed decisions about what outgoing transition to enable.

5 Asymmetric Asynchrony

As seen in the previous section, the class of symmetrically asynchronous nets is quite small. It precludes the implementation of many real-world behaviours, like waiting for one of multiple inputs to become readable, a Petri net representation of which will always include non free-choice structures.

Therefore we propose a less strict definition of asynchrony such that actions may depend synchronously on a single predetermined condition. In a hardware implementation the places which earlier could always forward a token into some silent transitions must now wait until they receive an explicit token removal signal from their posttransitions.

To this end we introduce a static priority over the preplaces of each transition. Every transition first removes the token from the most prioritised preplace and then continues along decreasing priority. To formalise this behaviour in a Petri net we insert a silent transition for each incoming arc of every transition. These silent transitions are forced to execute in sequence by newly introduced buffer places between them. In the final position of this chain, the original visible transition is executed. An example of this transformation is given in Figure 7.

¹ In [2, 3] the nature of the equivalence between the original and transformed net is not precisely specified. However, it can be argued that whereas the transformation from *EFC*-nets to *FC*-nets preserves branching time as well as causality, the transformation from *BFC*-nets to *EFC*-nets preserves branching time only: the third net of Figure 4 is interleaving bisimulation equivalent with its *EFC*-counterpart in Figure 6, but whereas the original net can perform the transitions *a* and *c* concurrently (in one step), the transformed net cannot.

Definition 5.1 Let $N = (S, O, \emptyset, F, M_0)$ be a plain net.

Let $g \subseteq (S \times O) \times (S \times O)$ be a relation on $F \cap (S \times O)$ such that for each $t \in O$ we have that $g \cap (\bullet t \times \{t\})$ is a total order on $\bullet t \times \{t\}$. Let \leq_g^t be the total order on $\bullet t$ given by $p \leq_g^t s$ iff $((p, t), (s, t)) \in g$.

We write \min_g^t for the \leq_g^t -minimal element of $\bullet t$ and $(s - 1)_g^t$ for the next place in $\bullet t$ that is \leq_g^t -smaller than s .

We define a set of silent transitions as $X := \{t_s \mid t \in O, s \in \bullet t\}$.

Let $h : X \rightarrow X \cup O$ be the function

$$h(t_s) = \begin{cases} t & \text{iff } s = \min_g^t \\ t_s & \text{otherwise} \end{cases}$$

The *asymmetrically asynchronous implementation with respect to g* of N is defined as the net $AI_g := (S \cup S^\tau, O, U', F', M_0)$ with

$$\begin{aligned} S^\tau &:= \{s_t \mid t \in O, s \in \bullet t, s \neq \min_g^t\}, \\ U' &:= h(X) \setminus O = \{t_s \mid t \in O, s \in \bullet t, s \neq \min_g^t\} \text{ and} \\ F' &:= F \cap (O \times S) \\ &\quad \cup \{(s, h(t_s)) \mid t \in O, s \in \bullet t\} \\ &\quad \cup \{(t_s, s_t) \mid t \in O, s \in \bullet t, s \neq \min_g^t\} \\ &\quad \cup \{(s_t, h(t_p)) \mid t \in O, s \in \bullet t, s \neq \min_g^t, p = (s - 1)_g^t\}. \end{aligned}$$

As before, we are interested in the relationship between nets and their possible implementations. The definition of asymmetric asynchrony however allows different implementations for the same net. We show that the lemmas and propositions from the previous sections carry over for all possible implementations. As in the earlier sections, we start by showing how the implementation simulates the original.

Lemma 5.1 Let $N = (S, O, \emptyset, F, M_0)$ be a plain net, $G \subseteq O$, $\sigma \in O^*$, $M_1, M_2 \subseteq S$ and g as above.

1. If $M_1 \xrightarrow{G}_N M_2$ then $M_1 \xrightarrow{\tau}_{AI_g(N)}^* \xrightarrow{G}_{AI_g(N)} M_2$.
2. If $M_1 \xrightarrow{\sigma}_N M_2$ then $M_1 \xrightarrow{\sigma}_{AI_g(N)} M_2$.

Proof Let $AI_g(N) = (S \cup S^\tau, O, U', F', M_0)$. Assume $M_1 \xrightarrow{G}_N M_2$. Then, due to the restrictions on g , there exists a sequence of pairwise disjoint nonempty sets $G_1, G_2, \dots, G_n \subseteq U'$ such that $\forall t \in G, s \in \bullet t, s \neq \min_g^t \exists i, 1 \leq i \leq n. t_s \in G_i$ and $\forall t_p \in G_i, t_q \in G_j, i < j. (t, p) \geq_g^t (t, q)$. By the construction of $AI_g(N)$ then

$$M_1 [G_1]_{AI_g(N)} [G_2]_{AI_g(N)} \cdots [G_n]_{AI_g(N)} [G]_{AI_g(N)} M_2.$$

All non-final steps of that execution can be split into a sequence of singletons.

The second statement follows by a straightforward induction on the length of σ . □

As for the symmetrical case, we wish to push back all tokens on S^τ in a marking to their roots in S . This time, however, multiple silent transitions need to be undone.

Definition 5.2 Let $N = (S, O, \emptyset, F, M_0)$ be a plain net and let $AI_g(N) = (S \cup S^\tau, O, U', F', M_0)$.

$\tau^{\leftarrow} : \mathcal{P}(S \cup S^\tau) \rightarrow \mathcal{P}(S)$ is the function defined by

$$\tau^{\leftarrow}(M) := (M \cap S) \cup \{s \mid \exists t \in O. s \in \bullet t \wedge \exists p_t \in M \cap S^\tau. p \leq_g^t s\}.$$

Given a reachable marking M of the implementation, τ^{\leftarrow} will produce a reachable marking of the original net, which by Lemma 5.1 is also a reachable marking of the implementation, from which M could have arisen by firing some of the added unobservable transitions. Note that $\tau^{\leftarrow}(\circ t) = \bullet t$ for any $t \in O$. The application of τ^{\leftarrow} is only meaningful for markings where no two elements of S^τ have originated from the same transition. However, implementations of contact-free nets produce only reachable markings which fulfil this condition, as we will show below.

We now give the invariant predicate γ that characterises the markings of an implementation that can be obtained from a reachable marking of the original net by firing some unobservable transitions.

Definition 5.3 Let $N = (S, O, \emptyset, F, M_0)$ be a plain net and $AI_g(N) = (S \cup S^\tau, O, U', F', M_0)$.

The predicate $\gamma \subseteq \mathcal{P}(S \cup S^\tau)$ is given by

$$\gamma(M) \quad :\Leftrightarrow \quad \tau^{\leftarrow}(M) \in [M_0]_N \wedge \forall p, q \in M. \tau^{\leftarrow}(\{p\}) \cap \tau^{\leftarrow}(\{q\}) \neq \emptyset \Rightarrow p = q.$$

Note that $\gamma(M)$ implies $f(M) \in \mathbb{N}$.

Lemma 5.2 Let N and $AI_g(N)$ be as above and $M \subseteq S \cup S^\tau$, with M finite.

Then $\forall p, q \in M. \tau^{\leftarrow}(\{p\}) \cap \tau^{\leftarrow}(\{q\}) \neq \emptyset \Rightarrow p = q$ iff $\tau^{\leftarrow}(M) \xrightarrow{\tau}_{AI_g}^* M$.

Proof Given that $\tau^{\leftarrow}(M) \subseteq S$, “if” follows directly from the construction of $AI_g(N)$.

For “only if”, assume $\forall p, q \in M. \tau^{\leftarrow}(p) \cap \tau^{\leftarrow}(q) \neq \emptyset \Rightarrow p = q$. Let $G_1 := \{t_s \mid s_t \in M\}$ and $G_{i+1} := \{t_p \mid \exists t_q \in G_i. q = (p-1)_g^t\}$ for $i > 1$. The assumption guarantees that all G_i are disjoint. Since M is finite, and $\bullet t$ is finite for all $t \in O$, there must be an $n \geq 0$ such that $G_i = \emptyset$ iff $i > n$. Now $\tau^{\leftarrow}(M) [G_n]_{AI_g(N)} [G_{n-1}]_{AI_g(N)} \cdots [G_1]_{AI_g(N)} M$. \square

Lemma 5.3 Let N and $AI_g(N)$ be as above, $G \subseteq O$, $\sigma \in O^*$ and $M, M' \subseteq S \cup S^\tau$.

1. $\gamma(M_0)$.
2. If $\gamma(M) \wedge M \xrightarrow{G}_{AI_g(N)} M'$ then $\tau^{\leftarrow}(M) \xrightarrow{G}_N \tau^{\leftarrow}(M') \wedge \gamma(M')$.
3. If $\gamma(M) \wedge M \xrightarrow{\tau}_{AI_g(N)} M'$ then $d(M) > d(M') \wedge \tau^{\leftarrow}(M) = \tau^{\leftarrow}(M') \wedge \gamma(M')$.
4. If $M_0 \xrightarrow{\sigma}_{AI_g(N)} M'$ then $M_0 \xrightarrow{\sigma}_N \tau^{\leftarrow}(M') \wedge \gamma(M')$.

Proof (1): $M_0 \in [M_0]_N$ and $\forall s \in M_0 \subseteq S. \tau^{\leftarrow}(\{s\}) = \{s\}$.

(2): Suppose $\gamma(M)$ and $M \xrightarrow{G}_{AI_g} M'$ with $G \subseteq O$. So $\tau^{\leftarrow}(M)$ is a reachable marking of N .

Let $t \in G$. Since t is enabled in M , we have $\circ t \subseteq M$ and hence $\tau^{\leftarrow}(\circ t) \subseteq \tau^{\leftarrow}(M)$. By construction, $\tau^{\leftarrow}(\circ t) = \bullet t$. Given that N is contact-free it follows that t is enabled in $\tau^{\leftarrow}(M)$.

Now let $t, u \in G$ with $t \neq u$. If $s \in \bullet t \cap \bullet u$ then $s \in \tau^{\leftarrow}(\circ t) \cap \tau^{\leftarrow}(\circ u)$ but since $t, u \in G, \circ t \cap \circ u = \emptyset$ and $\circ t \cup \circ u \subseteq M$, contradicting $\gamma(M)$. Hence $\bullet t \cap \bullet u = \emptyset$. Given that $\bullet t \cup \bullet u \subseteq \tau^{\leftarrow}(M)$ and N is contact-free, it follows that also $t^\bullet \cap u^\bullet = \emptyset$ and hence t and u are independent.

We will now show that $\left(\tau^{\leftarrow}(M) \setminus \bigcup_{t \in G} \bullet t \right) \cup \bigcup_{t \in G} t^{\circ} = \tau^{\leftarrow}(M')$.

By Definition 5.2 we have $\tau^{\leftarrow}(M) = \bigcup_{s \in M} \tau^{\leftarrow}(\{s\})$ for any $M \subseteq S \cup S^{\tau}$. Moreover, when M satisfies $\gamma(M)$ this union is disjoint. In that case, for any set $Y \subseteq M$ we have $\tau^{\leftarrow}(M \setminus Y) = \tau^{\leftarrow}(M) \setminus \tau^{\leftarrow}(Y)$.

$$M' = (M \setminus \{s \mid s \in {}^{\circ}t, t \in G\}) \cup \{s \mid s \in t^{\circ}, t \in G\}.$$

Therefore

$$\begin{aligned} \tau^{\leftarrow}(M') &= (\tau^{\leftarrow}(M) \setminus \tau^{\leftarrow}(\{s \mid s \in {}^{\circ}t, t \in G\})) \cup \tau^{\leftarrow}(\{s \mid s \in t^{\circ}, t \in G\}) \\ &= (\tau^{\leftarrow}(M) \setminus \{s \mid s \in \bullet t, t \in G\}) \cup \{s \mid s \in t^{\circ}, t \in G\}. \end{aligned}$$

Next we establish $\gamma(M')$. As in the proof of Lemma 3.3(2) we may assume that G is a singleton set $\{t\}$. Above we have shown that $\tau^{\leftarrow}(M') \in [M_0]_N$. We still need to prove that $\tau^{\leftarrow}(\{p\}) \cap \tau^{\leftarrow}(\{q\}) \neq \emptyset \Rightarrow p = q$ for all $p, q \in M'$. Assume the contrary, i.e. there are $p, q \in M'$ with $\tau^{\leftarrow}(\{p\}) \cap \tau^{\leftarrow}(\{q\}) \neq \emptyset$ but $p \neq q$. Since $\gamma(M)$, at least one of p and q —say p —must not be present in M . Then $p \in t^{\circ} = t^{\bullet} \subseteq S$. As $\tau^{\leftarrow}(\{q\}) \cap \tau^{\leftarrow}(\{p\}) \neq \emptyset$, $\tau^{\leftarrow}(\{p\}) = \{p\}$ and $q \neq p$ it must be that $q \in S^{\tau}$. Hence $q \notin t^{\circ}$, so $q \in M$, and $p \in \tau^{\leftarrow}(\{q\}) \subseteq \tau^{\leftarrow}(M)$. As shown above, t is enabled in $\tau^{\leftarrow}(M)$. By the contact-freeness of N , $(\tau^{\leftarrow}(M) \setminus \bullet t) \cap t^{\bullet} = \emptyset$, so $p \in \bullet t$. Since $p \notin M$, there exists a place $r_t \in ({}^{\circ}t \cap S^{\tau}) \subseteq M$ with $p \in \tau^{\leftarrow}(\{r_t\})$. By construction, $t^{\circ} \cap S^{\tau} = \emptyset$, so we have $r_t \notin M'$, hence $q \neq r_t$. However, $p \in \tau^{\leftarrow}(\{q\}) \cap \tau^{\leftarrow}(\{r_t\})$, contradicting $\gamma(M)$.

(3): Let $t_s \in U'$ such that $M[\{t_s\}]_{AI_g(N)} M'$. Then ${}^{\circ}t_s \cap S = \{s\}$ and $s^{\bullet} \neq \emptyset$. As $t_s^{\circ} \cap S = \emptyset$, no element of t_s° contributes to $d(M')$ and hence $d(M') = d(M) - 1$.

If ${}^{\circ}t_s \subseteq S$ then $\tau^{\leftarrow}(M') = \tau^{\leftarrow}((M \setminus {}^{\circ}t_s) \cup t_s^{\circ}) = \tau^{\leftarrow}((M \setminus \{s\}) \cup \{s_t\}) = \tau^{\leftarrow}(M)$. Otherwise let $p \in S$ such that $p_t \in {}^{\circ}t_s$. Then $\tau^{\leftarrow}(M') = \tau^{\leftarrow}((M \setminus {}^{\circ}t_s) \cup t_s^{\circ}) = \tau^{\leftarrow}((M \setminus \{s, p_t\}) \cup \{s_t\}) = \tau^{\leftarrow}(M)$.

Moreover, $\gamma(M') \Leftrightarrow \gamma(M)$.

(4): As in Lemma 3.3. □

Proposition 5.1 Let N and $AI_g(N)$ be as before and $M \subseteq S \cup S^{\tau}$.

1. $M \in [M_0]_{AI_g(N)}$ iff $\gamma(M)$.
2. $AI_g(N)$ is contact-free.
3. $AI_g(N)$ is a τ -net.

Proof (1): Identical to the proof of Proposition 3.1(1), using the lemmas of Section 5.

(2): Let $M \in [M_0]_{AI_g(N)}$. Then $\gamma(M)$, and hence $\tau^{\leftarrow}(M) \in [M_0]_N$.

Consider any $t \in O$ with ${}^{\circ}t \subseteq M$. Assume $(M \setminus {}^{\circ}t) \cap t^{\circ} \neq \emptyset$. Since $t^{\circ} = t^{\bullet} \subseteq S$ let $p \in S$ be such that $p \in M \cap t^{\circ}$ and $p \notin {}^{\circ}t$. As N is contact-free we have $(\tau^{\leftarrow}(M) \setminus \bullet t) \cap t^{\bullet} = \emptyset$, so since $p \in \tau^{\leftarrow}(M) \cap t^{\bullet}$ it must be that $p \in \bullet t$. Hence, using that $p \notin {}^{\circ}t$ there must be an $s_t \in {}^{\circ}t \subseteq M$ with $s \leq_g^t p$, and hence $p \in \tau^{\leftarrow}(\{s_t\})$. We have $p \neq s_t$ yet $p \in \tau^{\leftarrow}(\{p\})$, violating $\gamma(M)$.

Now consider any $t_p \in U'$ with ${}^{\circ}t_p \subseteq M$. As $p \in {}^{\circ}t_p$ and $t_p^{\circ} = \{p_t\}$ we have that $(M \setminus {}^{\circ}t_p) \cap t_p^{\circ} \neq \emptyset$ only if $p \in M \wedge p_t \in M$. However, $\tau^{\leftarrow}(\{p\}) = \{p\} \subseteq \tau^{\leftarrow}(\{p_t\})$ which would violate $\gamma(M)$.

(3): Identical to the proof of Proposition 3.1(3). □

Proposition 5.2 Let N be a plain net and g as before. Then $AI_g(N)$ is divergence free.

Proof This follows immediately from Lemma 5.3(3). □

Proposition 5.3 Let N and $AI_g(N)$ as before. Then $\mathcal{F}(N) \subseteq \mathcal{F}(AI_g(N))$.

Proof Identical to that of Proposition 3.3, using the lemmas of Section 5. \square

We define a net to be *asymmetrically asynchronous* if any of the possible implementations simulates the net sufficiently.

Definition 5.4

The class of *asymmetrically asynchronous nets respecting branching time equivalence* is defined as $AA(B) := \{N \mid \exists g. AI_g(N) \approx_{\mathcal{F}} N\}$.

As before, we would like to obtain a semi-structural characterisation of $AA(B)$ in the spirit of Theorems 3.1 and 4.1. Unfortunately we didn't succeed in this, but we obtained structural upper and lower bounds for this net class.

Definition 5.5

A net $N = (S, O, \emptyset, F, M_0)$ has a left and right reachable M iff $\exists t, u, v \in O \exists p \in \bullet t \cap \bullet u \exists q \in \bullet u \cap \bullet v. t \neq u \wedge u \neq v \wedge p \neq q \wedge \exists M_1, M_2 \in [M_0]. \bullet t \cup \bullet u \subseteq M_1 \wedge \bullet v \cup \bullet u \subseteq M_2$.

A net $N = (S, O, \emptyset, F, M_0)$ has a left and right border reachable M iff $\exists t, u, v \in O \exists p \in \bullet t \cap \bullet u \exists q \in \bullet u \cap \bullet v. t \neq u \wedge u \neq v \wedge p \neq q \wedge \exists M_1, M_2 \in [M_0]. \bullet t \subseteq M_1 \wedge \bullet v \subseteq M_2$.

Theorem 5.1

A plain net N in $AA(B)$ has no left and right reachable M .

Proof Let $N = (S, O, \emptyset, F, M_0)$. Assume N has a left and right reachable M . Then there exist $t, u, v \in O$ and $p, q \in S$ such that $p \in \bullet t \cap \bullet u, q \in \bullet u \cap \bullet v, t \neq u, u \neq v, p \neq q$ and there are reachable markings $M_1, M_2 \in [M_0]$ such that $\bullet t \cup \bullet u \subseteq M_1$ and $\bullet v \cup \bullet u \subseteq M_2$. We will show that $AI_g(N) \not\approx_{\mathcal{F}} N$, regardless of the choice of g .


The problematic transition will be u . Either $p >_g^u q$ or $q >_g^u p$. Due to symmetry we can assume the former without loss of generality. So $p \neq \min_g^u$. We know that there is some $\sigma \in O^*$ such that $M_0 \xrightarrow{\sigma}_N M_1 \wedge \bullet t \subseteq M_1$. By Lemma 2.1 it follows that $\langle \sigma, \{t\} \rangle \notin \mathcal{F}(N)$.

By Lemma 5.1 also $M_0 \xrightarrow{\sigma}_{AI_g(N)} M_1$. Let $p_0, \dots, p_n \in S$ such that $\{s \in \bullet u \mid p \leq_g^u s\} = \{p_0, \dots, p_n\}$ with $p_0 = p$ and $p_{i-1} = (p_i - 1)_g^u$ for $2 \leq i \leq n$. Since $\bullet u \subseteq M_1$ there thus exists an M'_1 with $M_1[\{u_{p_n}\}]_{AI_g(N)}[\{u_{p_{n-1}}\}]_{AI_g(N)} \cdots [\{u_{p_0}\}]_{AI_g(N)} M'_1$. Note that $p_u \in M'_1$. By Propositions 5.2 there exists an M''_1 with $M'_1 \xrightarrow{\tau}_{AI_g(N)} M''_1 \wedge M'_1 \xrightarrow{\tau}_{AI_g(N)}$, and Proposition 5.1(1) yields $\gamma(M''_1)$.

From the construction of $AI_g(N)$, using that $p_u \in M'_1$, it follows that $\exists s \in \bullet u. s \leq_g^u p \wedge s_u \in M''_1$. Moreover, $p \in \tau^{\Leftarrow}(\{s_u\})$. We also have $p \in \bullet t = \tau^{\Leftarrow}(\circ t)$, so $\exists r \in \circ t. p \in \tau^{\Leftarrow}(\{r\})$. As $\gamma(M''_1)$, we have $r \notin M''_1$, and thus $\circ t \not\subseteq M''_1$. Therefore $\langle \sigma, \{t\} \rangle \in \mathcal{F}(AI_g(N))$. Hence N is not in $AA(B)$. \square

Theorem 5.2

A plain net N which has no left and right border reachable M is in $AA(B)$.

Figure 8: Nets which have a left and right border reachable M , but no left and right reachable M

Proof Let $N = (S, O, \emptyset, F, M_0)$. Given a transition $u \in O$, we say that a u -conflict occurs in a preplace $p \in \bullet u$ when $\exists t \in p^\bullet. t \neq u \wedge (\exists M_1 \in [M_0]_N. \bullet t \subseteq M_1)$.

Assume N has no left and right border reachable M . This means that every $u \in O$ has at most one preplace where an u -conflict occurs. Now choose $g \subseteq (S \times O) \times (S \times O)$ such that for all $u \in O$, \min_g^u is that single place, if it exists. Let $AI_g(N) = (S \cup S^\tau, O, U', F', M_0)$.

We prove that $\mathcal{F}(N) = \mathcal{F}(AI_g(N))$. From Proposition 5.3 we already have that $\mathcal{F}(N) \subseteq \mathcal{F}(AI_g(N))$. Therefore consider a failure pair $\langle \sigma, X \rangle \in \mathcal{F}(AI_g(N))$. We need to show that $\langle \sigma, X \rangle \in \mathcal{F}(N)$.

There exists some $M_1 \subseteq S \cup S^\tau$ with $M_0 \xrightarrow{\sigma}_{AI_g(N)} M_1 \wedge M_1 \xrightarrow{\tau} \wedge \forall t \in X. M_1 \xrightarrow{\{t\}}$. By Lemma 5.3 $M_0 \xrightarrow{\sigma}_N \tau^{\leftarrow}(M_1)$. Now take any $t \in X$. Assume $\tau^{\leftarrow}(M_1) \xrightarrow{\{t\}}_N$. Then $\circ t \not\subseteq M_1$ but $\bullet t \subseteq \tau^{\leftarrow}(M_1)$. By construction of τ^{\leftarrow} we have $\forall p \in \bullet t. p \in M_1 \vee \exists u \in O. p \in \bullet u \wedge \exists s_u \in M_1 \cap S^\tau. s \leq_g^u p$. Now suppose we had $\forall p \in \bullet t. p \in M_1 \vee \exists s_t \in M_1 \cap S^\tau. s \leq_g^t p$. Then $\circ t \subseteq M_1$, using $M_1 \xrightarrow{\tau}$. \nexists Hence $\exists p \in \bullet t \exists u \in O. u \neq t \wedge p \in \bullet u \wedge \exists s_u \in M_1 \cap S^\tau. s \leq_g^u p$.


But then $t \in p^\bullet \wedge t \neq u \wedge \tau^{\leftarrow}(M_1) \in [M_0]_N \wedge \bullet t \subseteq \tau^{\leftarrow}(M_1)$, so a u -conflict occurs in $p \in \bullet u$. Yet $\exists s \leq_g^u p. s_u \in S^\tau$ implies that $s \neq \min_g^u$ and hence $p \neq \min_g^u$, by the construction of $AI_g(N)$. This however contradicts our construction for g given above. Hence $\tau^{\leftarrow}(M_1) \xrightarrow{\{t\}}_N$. Applying this argument for all $t \in X$ yields $\langle \sigma, X \rangle \in \mathcal{F}(N)$ and thereby $\mathcal{F}(AI_g(N)) \subseteq \mathcal{F}(N)$. Thus $N \in AA(B)$. \square

Figure 8 shows two nets, each with a left and right border reachable M but no left and right reachable M , that thus fall in the grey area between our structural upper and lower bounds for the class $AA(B)$. In this case the first net falls outside $AA(B)$, whereas the second net falls inside. The crucial difference between these two examples is the information available to u about the execution of y .

There exists an implementation for the right net, namely by u taking the tokens from r , q and s in that order. The first token (from r) conveys the information that y was executed, and thus t is not enabled. Collecting the last token (from s) could fail, due to v removing it earlier. Even so, removing the tokens from r and q did not disable any transition that could fire in the original net. In the left net such an implementation will not work.

The following proposition says that our class of symmetrically asynchronous nets strictly extends the corresponding class of asymmetrically asynchronous nets.

Proposition 5.4 $SA(B) \subsetneq AA(B)$.

Figure 9: $N \in AA(B)$, $N \notin ESPL$

Proof A net which has no partially reachable N also has no left or right border reachable M. The inequality follows from the example in Figure 3. \square

As before, our class $AA(B)$ is related to some known net classes [3].

Definition 5.6 Let $N = (S, O, \emptyset, F, M_0)$ be a plain net.

1. N is *simple*, $N \in SPL$, iff $\forall p, q \in S. p \neq q \wedge p^\bullet \cap q^\bullet \neq \emptyset \Rightarrow |p^\bullet| = 1 \vee |q^\bullet| = 1$.
2. N is *extended simple*, $N \in ESPL$, iff $\forall p, q \in S. p^\bullet \cap q^\bullet \neq \emptyset \Rightarrow p^\bullet \subseteq q^\bullet \vee q^\bullet \subseteq p^\bullet$.

Extended simple nets appear in [2] under the name *asymmetric choice systems*. Note that simple is equivalent to M-free, where M is as in Definition 5.5 but without the reachability clauses. Clearly, we have $FC \subsetneq SPL \subsetneq ESPL$ and $EFC \subsetneq ESPL$, whereas $EFC \not\subseteq SPL$ and $SPL \not\subseteq EFC$: the inclusions follow immediately from the definitions, and the first two nets of Figure 4 provide counterexamples to the inequalities.

The class of asymmetrically asynchronous nets respecting branching time equivalence strictly extends the class of simple nets, whereas it is incomparable with the class of extended simple nets.

Proposition 5.5 $SPL \subsetneq AA(B)$, $AA(B) \not\subseteq ESPL$ and $ESPL \not\subseteq AA(B)$.

Proof The inclusion is straightforward, and the inequalities follow from the counterexamples in Figure 4 (the second one) and Figure 9. The missing tokens in the latter example are intended. As no action is possible there will not be any additional implementation failures. \square

The relations between the classes SPL , $ESPL$ and $AA(B)$ are summarised in Figure 10. Similarly to what we did in Section 4, we now try to translate Figure 10 into an intuitive description.

The basic intuition behind SPL is that for every transition there is only one preplace where conflict can possibly occur. Whereas in SPL that possibility is determined by the static net structure, in $AA(B)$ reachability is also considered.

Similar to the difference between FC and EFC there exists a difference between $ESPL$ and SPL which originates from the fact that $ESPL$ allows small transformations to a net before testing whether it lies in SPL . Again our class $AA(B)$ does not allow such “helping” transformations.




Figure 10: Overview of asymmetric-choice-like net classes

6 Conclusion and Related Work

We have investigated the effect of different types of asynchronous interaction, using Petri nets as our system model. We propose three different interaction patterns: fully asynchronous, symmetrically asynchronous and asymmetrically asynchronous. An asynchronous implementation of a net is then obtained by inserting silent (unobservable) transitions according to the respective pattern. The pattern for asymmetric asynchrony is parametric in the sense that the actual asynchronous implementation of a net depends on a chosen priority function on the input places of a transition. For each of these cases, we investigated for which types of nets the asynchronous implementation of a net changes its behaviour with respect to failures equivalence (in the case of asymmetric asynchrony, the ‘best’ priority function may be used). It turns out that we obtain a hierarchy of Petri net classes, where each class contains those nets which do not change their behaviour when transformed into the asynchronous version according to one of the interaction patterns. This is not surprising because later constructions allow a more fine-grained control over the interactions than earlier ones.

We did not consider connections from transitions to their postplaces as relevant to determine asynchrony and distributability. This is because we only discussed contact-free nets, where no synchronisation by postplaces is necessary. In the spirit of Definition 3.1 we could insert τ -transitions on any or all arcs from transitions to their postplaces, and the resulting net would always be equivalent to the original.

Although we compare the behaviour of a net and its asynchronous implementations in terms of failures equivalence, we believe that the very same classes of nets are obtained when using any other reasonable behavioural equivalence that respects branching time to some degree and abstracts from silent transitions—no matter if this is an interleaving equivalence, or one that respects causality. We would get larger classes of nets, for example for the case of full asynchrony including the net of Figure 2, if we merely required a net N and its implementation to be equivalent under a suitably chosen linear time equivalence. This option is investigated in [18].

The central results of the paper give semi-structural characterisations of our semantically defined classes of nets. Moreover, we relate these classes to well-known and well-understood structurally defined classes of nets, like free choice nets, extended free choice nets and simple nets.

To illustrate the potential interpretation of our results in other models of distributed systems, we give an example.

Message sequence charts (MSCs), also contained in UML 2.0 under the name sequence diagrams, are a model for specifying interactions between components (*instances*) of a system. A simple kind are *basic message sequence charts* (BMSCs) as defined in [12], where choices are not allowed. A Petri net semantics of BMSCs with asynchronous communication and a unique sending and receiving event for each message will yield Petri nets with unbranched places (see for instance [9]). Hence in this case the resulting Petri nets are conflict-free and therefore fully asynchronously implementable according to Theorem 3.1.

However in extended versions of MSCs, e.g. in UML 2.0 or in live sequence charts (LSCs, see [10]), inline expressions allow to describe choices between possible behaviours in MSCs. Consider for example the MSC given in Figure 11 and a naive Petri net representation. The instances i_1 and i_2 can either communicate or execute their local actions. Obviously, this requires some mechanism in order to make sure that the choice is performed in a coherent way (see e.g. [7] for a discussion of this type of problem). In the Petri net representation, we find a reachable N , hence with Theorem 4.1 the net does not belong to the class $SA(B)$ of symmetrically asynchronously implementable nets. However, the net is M -free, and thus does belong to the class $AA(B)$ of asymmetrically asynchronously implementable nets. By




Figure 11: An MSC and a potential implementation as Petri net, which has an N.

giving priority to the collection of the message token (choosing the appropriate function g in our notion of implementation), it can be assured that instance $i2$ does not make the wrong choice and gets stuck (however it is still not clear whether the message will actually be consumed).

The obvious question is whether the naive Petri net interpretation we have given is conform with the intended semantics of the *alt*-construct (according to the informal UML semantics the alternatives always have to be executed completely; in LSCs it is specified explicitly whether messages are assured to arrive). However, on basis of a maybe more elaborate Petri nets semantics, it could be discussed what types of MSCs can be used to describe physically distributed systems, in particular which type of construct for choices is reasonable in this case.

Another model of reactive systems where we can transfer our results to are process algebras. When giving Petri net semantics to process algebras, it is an interesting question to investigate which classes of nets in our classification are obtained for certain types of operators or restricted languages, and to compare the results with results on language hierarchies (as summarised below).

We now give an overview on related work. A more extensive discussion is contained in [18]. We start by commenting on related work in Petri net theory.

The structural net classes we compare our constructions to were all taken from [3], where Eike Best and Mike Shields introduce various transformations between free choice nets, simple nets and extended variants thereof. They use “essential equivalence” to compare the behaviour of different nets, which they only give informally. This equivalence is insensitive to divergence, which is also relied upon in their transformations. As observed in Footnote 1, it also does not preserve concurrency. They continue to show conditions under which liveness can be guaranteed for some of the classes.

In [1], Wil van der Aalst, Ekkart Kindler and Jörg Desel introduce two extensions to extended simple nets, by allowing self-loops to ignore the discipline imposed by the *ESPL*-requirement. This however assumes a kind of “atomicity” of self-loops, which we did not allow in this paper. In particular we do not implicitly assume that a transition will not change the state of a place it is connected to by a self-loop, since in case of deadlock, the temporary removal of a token from such a place might not be temporary indeed.

In [17] Wolfgang Reisig introduces a class of systems which communicate using buffers and where the relative speeds of different components are guaranteed to be irrelevant. The resulting nets are simple nets. He then proceeds introducing a decision procedure for the problem whether a marking exists which makes the complete system live.

The most similar work to our approach we have found is [11], where Richard Hopkins introduces the

concept of *distributable* Petri Nets. These are defined in terms of *locality functions*, which assign to every transition t a set of possible machines or locations $L(t)$ on which t may be executed, subject to the restriction that a set of transitions with a common preplace must share a common machine. A plain net N is distributable iff for every locality function L that can be imposed on it, it has a “distributed implementation”, a τ -net N' with the same set of visible transitions, in which each transition is assigned a specific location, subject to three restrictions:

- the location of a visible transition t is chosen from $L(t)$,
- transitions with a common preplace must have the same location
- and there exists a weak bisimulation between N and N' , such that all τ -transitions involved in simulating a transition t from N reside on one of the locations $L(t)$.

The last clause enforces both a behavioural correspondence between N and N' and a structural one (through the requirement on locations). Thus, as in our work, the implementation is a τ -net that is required to be behaviourally equivalent to the original net. However, whereas we enforce particular implementations of an original net, Hopkins allows implementations which are quite elaborate and make informed decisions based upon global knowledge of the net. Consequently, his class of distributable nets is larger than our asynchronous net classes. As Hopkins notes, due to his use of interleaving semantics, his distributed implementations do not always display the same concurrent behaviour as the original nets, namely they add concurrency in some cases. This does not happen in our asynchronous implementations.

Another branch of related work is in the context of distributed algorithms. In [5] Luc Bougé considers the problem of implementing symmetric leader election in the sublanguages of CSP obtained by either allowing all guards, only input guards or no communication guards at all in guarded choice. He finds that the possibility of implementing it depends heavily on the structure of the communication graphs, while truly symmetric schemes are only possible in CSP with input and output guards.

Quite a number of papers consider the question of synchronous versus asynchronous interaction in the realm of process algebras and the π -calculus. In [4] Frank de Boer and Catuscia Palamidessi consider various dialects of CSP with differing degrees of asynchrony. In particular, they consider CSP without output guards and CSP without any communication based guards. They also consider explicitly asynchronous variants of CSP where output actions cannot block, i.e. asynchronous sending is assumed. Similar work is done for the π -calculus in [16] by Catuscia Palamidessi, in [15] by Uwe Nestmann and in [8] by Dianele Gorla. A rich hierarchy of asynchronous π -calculi has been mapped out in these papers. Again mixed-choice, i.e. the ability to combine input and output guards in a single choice, plays a central role in the implementation of truly synchronous behaviour. It would be interesting to explore the possible connections between these languages and our net classes.

In [19], Peter Selinger considers labelled transition systems whose visible actions are partitioned into input and output actions. He defines asynchronous implementations of such a system by composing it with in- and output queues, and then characterises the systems that are behaviourally equivalent to their asynchronous implementations. The main difference with our approach is that we focus on asynchrony within a system, whereas Selinger focusses on the asynchronous nature of the communications of a system with the outside world.

Finally, there are approaches on hardware design where asynchronous interaction is an intriguing feature due to performance issues. For this, see the papers [13] and [14] by Leslie Lamport. In [14] he considers arbitration in hardware and outlines various arbitration-free “wait/signal” registers. He notes that non-determinism is thought to require arbitration, but no proof is known. He concludes that only marked graphs can be implemented using these registers. Lamport then introduces “Or-Waiting”, i.e. waiting for any of two signals, but has no model available to characterise the resulting processes. The used communication primitives bear a striking similarity to our symmetrically asynchronous nets.

References

- [1] W.M.P. van der Aalst, E. Kindler & J. Desel (1998): *Beyond asymmetric choice: A note on some extensions*. *Petri Net Newsletter* 55, pp. 3–13.
- [2] E. Best (1987): *Structure theory of Petri nets: The free choice hiatus*. In W. Brauer, W. Reisig & G. Rozenberg, editors: *Advances in Petri Nets 1986*, LNCS 254, Springer, pp. 168–206.
- [3] E. Best & M.W. Shields (1983): *Some equivalence results for free choice nets and simple nets and on the periodicity of live free choice nets*. In G. Ausiello & M. Protasi, editors: *Proceedings 8th Colloquium on Trees in Algebra and Programming (CAAP '83)*, LNCS 159, Springer, pp. 141–154.
- [4] F.S. de Boer & C. Palamidessi (1991): *Embedding as a tool for language comparison: On the CSP hierarchy*. In J.C.M. Baeten & J.F. Groote, editors: *Proceedings 2nd International Conference on Concurrency Theory (CONCUR '91)*, Amsterdam, The Netherlands, LNCS 527, Springer, pp. 127–141.
- [5] L. Bougé (1988): *On the existence of symmetric algorithms to find leaders in networks of communicating sequential processes*. *Acta Informatica* 25(2), pp. 179–201.
- [6] S.D. Brookes, C.A.R. Hoare & A.W. Roscoe (1984): *A theory of communicating sequential processes*. *Journal of the ACM* 31(3), pp. 560–599.
- [7] T. Gehrke, U. Goltz & H. Wehrheim (1999): *Zur semantischen Analyse der dynamischen Modelle von UML mit Petri-Netzen*. In E. Schnieder, editor: *Proceedings 6th Symposium on Development and Operation of Complex Automation Systems*.
- [8] D. Gorla (2006): *On the relative expressive power of asynchronous communication primitives*. In L. Aceto & A. Ingólfssdóttir, editors: *Proceedings 9th International Conference on Foundations of Software Science and Computation Structures (FoSSaCS '06)*, LNCS 3921, Springer, pp. 47–62.
- [9] P. Graubmann, E. Rudolph & J. Grabowski (1993): *Towards a petri net based semantics definition for message sequence charts*. In *Proceedings 6th SDL Forum (SDL '93)*.
- [10] D. Harel & R. Marelly (2003): *Come, Let's Play*. Springer.
- [11] R.P. Hopkins (1991): *Distributable nets*. In *Advances in Petri Nets 1991*, LNCS 524, Springer, pp. 161–187.
- [12] International Telecommunication Union (1996): *Message sequence chart*. Standard ITU-T Z.120.
- [13] L. Lamport (1978): *Time, clocks, and the ordering of events in a distributed system*. *Communications of the ACM* 21(7), pp. 558–565.
- [14] L. Lamport (2003): *Arbitration-free synchronization*. *Distributed Computing* 16(2-3), pp. 219–237.
- [15] U. Nestmann (2000): *What is a 'good' encoding of guarded choice?* *Information and Computation* 156, pp. 287–319.
- [16] C. Palamidessi (1997): *Comparing the expressive power of the synchronous and the asynchronous pi-calculus*. In *Conference Record of the 24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '97)*, ACM Press, pp. 256–265.

- [17] W. Reisig (1982): *Deterministic buffer synchronization of sequential processes*. *Acta Informatica* 18, pp. 115–134.
- [18] J.-W. Schicke (2008): *Studienarbeit: Asynchronous Petri net classes*.
- [19] P. Selinger (1997): *First-order axioms for asynchrony*. In Proceedings 8th International Conference on Concurrency Theory (CONCUR '97), Warsaw, Poland, LNCS 1243, Springer, pp. 376–390.
- [20] W. Vogler (1993): *Bisimulation and action refinement*. *Theoretical Computer Science* 114(1), pp. 173–200.