

ON CHARACTERISING DISTRIBUTABILITY *

ROB VAN GLABBEK ^a, URSULA GOLTZ ^b, AND JENS-WOLFHARD SCHICKE-UFFMANN ^c

^a NICTA, Sydney, Australia
School of Computer Science and Engineering, Univ. of New South Wales, Sydney, Australia
e-mail address: rvg@cs.stanford.edu

^{b,c} Institute for Programming and Reactive Systems, TU Braunschweig, Germany
e-mail address: goltz@ips.cs.tu-bs.de, draflow@gmx.de

ABSTRACT. We formalise a general concept of distributed systems as sequential components interacting asynchronously. We define a corresponding class of Petri nets, called LSGA nets, and precisely characterise those system specifications which can be implemented as LSGA nets up to branching ST-bisimilarity with explicit divergence.

1. INTRODUCTION

The aim of this paper is to contribute to a fundamental understanding of the concept of a distributed reactive system and the paradigms of synchronous and asynchronous interaction. We start by giving an intuitive characterisation of the basic features of distributed systems. In particular we assume that distributed systems consist of components that reside on different locations, and that any signal from one component to another takes time to travel. Hence the only interaction mechanism between components is asynchronous communication.

Our aim is to characterise which system specifications may be implemented as distributed systems. In many formalisms for system specification or design, synchronous communication is provided as a basic notion; this happens for example in process algebras. Hence a particular challenge is that it may be necessary to simulate synchronous communication by asynchronous communication.

2012 ACM CCS: [Theory of computation]: Models of computation—Concurrency—Distributed computing models.

Key words and phrases: Concurrency, Petri nets, distributed systems, reactive systems, asynchronous interaction, semantic equivalences.

* This paper is adapted from [GGS12]; it characterises distributability for a slightly larger range of semantic equivalence relations, and incorporates various remarks stemming from [GGS08]. An extended abstract appeared in L. Birkedal, ed.: Proc. 15th Int. Conf. on Foundations of Software Science and Computation Structures (FoSSaCS 2012), LNCS 7213, Springer, 2012, pp. 331–345, doi:10.1007/978-3-642-2872-9_22.

^a NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Centre of Excellence program.

This work was partially supported by the DFG (German Research Foundation).

Trivially, any system specification may be implemented distributedly by locating the whole system on one single component. Hence we need to pose some additional requirements. One option would be to specify locations for system activities and then to ask for implementations satisfying this distribution and still preserving the behaviour of the original specification. This is done in [BCD02]. Here we pursue a different approach. We add another requirement to our notion of a distributed system, namely that its components only allow sequential behaviour. We then ask whether an arbitrary system specification may be implemented as a distributed system consisting of sequential components in an optimal way, that is without restricting the concurrency of the original specification. This is a particular challenge when synchronous communication interacts with concurrency in the specification of the original system. We will give a precise characterisation of the class of distributable systems, which answers in particular under which conditions synchronous communication may be implemented in a distributed setting.

For our investigations we need a model which is expressive enough to represent concurrency. It is also useful to have an explicit representation of the distributed state space of a distributed system, showing in particular the local control states of components. We choose Petri nets, which offer these possibilities and additionally allow finite representations of infinite behaviours. We focus on the class of *structural conflict nets* [GGS11]—a proper generalisation of the class of one-safe place/transition systems, where conflict and concurrency are clearly separated.

For comparing the behaviour of systems with their distributed implementation we need a suitable equivalence notion. Since we think of open systems interacting with an environment, and since we do not want to restrict concurrency in applications, we need an equivalence that respects branching time and concurrency to some degree. Our implementations use transitions which are invisible to the environment, and this should be reflected in the equivalence by abstracting from such transitions. However, we do not want implementations to introduce divergence. In the light of these requirements we work with two semantic equivalences. *Step failures equivalence* is one of the weakest equivalences that captures branching time, concurrency and divergence to some degree; whereas *branching ST-bisimilarity with explicit divergence* fully captures branching time, divergence, and those aspects of concurrency that can be represented by concurrent actions overlapping in time. We obtain the same characterisation for both notions of equivalence, and thus implicitly for all notions in between these extremes.

We model distributed systems consisting of sequential components as an appropriate class of Petri nets, called *LSGA nets*. These are obtained by composing nets with sequential behaviour by means of an asynchronous parallel composition. We show that this class corresponds exactly to a more abstract notion of distributed systems, formalised as *distributed nets* [GGS08].

We then consider distributability of system specifications which are represented as structural conflict nets. A net N is *distributable* if there exists a distributed implementation of N , that is a distributed net which is semantically equivalent to N . In the implementation we allow unobservable transitions, and labellings of transitions, so that single actions of the original system may be implemented by multiple transitions. However, the system specifications for which we search distributed implementations are *plain* nets without these features. This restriction is motivated in the conclusion.

We give a precise characterisation of distributable nets in terms of a semi-structural property. This characterisation provides a formal proof that the interplay between choice and synchronous communication is a key issue for distributability.

To establish the correctness of our characterisation we develop a new method for rigorously proving the equivalence of two Petri nets, one of which known to be plain, up to branching ST-bisimilarity with explicit divergence.

2. BASIC NOTIONS

In this paper we employ *signed multisets*, which generalise multisets by allowing elements to occur in it with a negative multiplicity.

Definition 2.1. Let X be a set.

- A *signed multiset* over X is a function $A: X \rightarrow \mathbb{Z}$, i.e. $A \in \mathbb{Z}^X$.
It is a *multiset* iff $A \in \mathbb{N}^X$, i.e. iff $A(x) \geq 0$ for all $x \in X$.
- $x \in X$ is an *element of* a signed multiset $A \in \mathbb{Z}^X$, notation $x \in A$, iff $A(x) \neq 0$.
- For signed multisets A and B over X we write $A \leq B$ iff $A(x) \leq B(x)$ for all $x \in X$;
 $A \cup B$ denotes the signed multiset over X with $(A \cup B)(x) := \max(A(x), B(x))$,
 $A \cap B$ denotes the signed multiset over X with $(A \cap B)(x) := \min(A(x), B(x))$,
 $A + B$ denotes the signed multiset over X with $(A + B)(x) := A(x) + B(x)$,
 $A - B$ denotes the signed multiset over X with $(A - B)(x) := A(x) - B(x)$, and
for $k \in \mathbb{Z}$ the signed multiset $k \cdot A$ is given by $(k \cdot A)(x) := k \cdot A(x)$.
- The function $\emptyset: X \rightarrow \mathbb{N}$, given by $\emptyset(x) := 0$ for all $x \in X$, is the *empty* multiset over X .
- If A is a signed multiset over X and $Y \subseteq X$ then $A \upharpoonright Y$ denotes the signed multiset over Y defined by $(A \upharpoonright Y)(x) := A(x)$ for all $x \in Y$.
- The cardinality $|A|$ of a signed multiset A over X is given by $|A| := \sum_{x \in X} |A(x)|$.
- A signed multiset A over X is *finite* iff $|A| < \infty$, i.e., iff the set $\{x \mid x \in A\}$ is finite.
We write $A \in_F \mathbb{Z}^X$ or $A \in_F \mathbb{N}^X$ to indicate that A is a finite (signed) multiset over X .
- Any function $f: X \rightarrow \mathbb{Z}$ or $f: X \rightarrow \mathbb{Z}^Y$ from X to either the integers or the signed multisets over some set Y extends to the finite signed multisets A over X by $f(A) = \sum_{x \in X} A(x) \cdot f(x)$.

Two signed multisets $A: X \rightarrow \mathbb{Z}$ and $B: Y \rightarrow \mathbb{Z}$ are *extensionally equivalent* iff $A \upharpoonright (X \cap Y) = B \upharpoonright (X \cap Y)$, $A \upharpoonright (X \setminus Y) = \emptyset$, and $B \upharpoonright (Y \setminus X) = \emptyset$. In this paper we often do not distinguish extensionally equivalent signed multisets. This enables us, for instance, to use $A + B$ even when A and B have different underlying domains. A multiset A with $A(x) \in \{0, 1\}$ for all x is identified with the set $\{x \mid A(x) = 1\}$. A signed multiset with elements x and y , having multiplicities -2 and 3 , is denoted as $-2 \cdot \{x\} + 3 \cdot \{y\}$.

We consider here general labelled place/transition systems with arc weights. Arc weights are not necessary for the results of the paper, but are included for the sake of generality.

Definition 2.2. Let Act be a set of *visible actions* and $\tau \notin \text{Act}$ be an *invisible action*. Let $\text{Act}_\tau := \text{Act} \dot{\cup} \{\tau\}$. A (*labelled*) *Petri net* (over Act_τ) is a tuple $N = (S, T, F, M_0, \ell)$ where

- S and T are disjoint sets (of *places* and *transitions*, together called the *elements* of N),
- $F: (S \times T \cup T \times S) \rightarrow \mathbb{N}$ (the *flow relation* including *arc weights*),
- $M_0: S \rightarrow \mathbb{N}$ (the *initial marking*), and
- $\ell: T \rightarrow \text{Act}_\tau$ (the *labelling function*).

Petri nets are depicted by drawing the places as circles and the transitions as boxes, containing their label. Identities of places and transitions are displayed next to the net element. When $F(x, y) > 0$ for $x, y \in S \cup T$ there is an arrow (*arc*) from x to y , labelled with the *arc weight* $F(x, y)$. Weights 1 are elided. When a Petri net represents a concurrent system, a global state of this system is given as a *marking*, a multiset M of places, depicted by placing $M(s)$ dots (*tokens*) in each place s . The initial state is M_0 .

The behaviour of a Petri net is defined by the possible moves between markings M and M' , which take place when a finite multiset G of transitions *fires*. In that case, each occurrence of a transition t in G consumes $F(s, t)$ tokens from each place s . Naturally, this can happen only if M makes all these tokens available in the first place. Next, each t produces $F(t, s)$ tokens in each s . Definition 2.4 formalises this notion of behaviour.

Definition 2.3. Let $N = (S, T, F, M_0, \ell)$ be a Petri net and $x \in S \cup T$.

The multisets $\bullet x, x^\bullet : S \cup T \rightarrow \mathbb{N}$ are given by $\bullet x(y) = F(y, x)$ and $x^\bullet(y) = F(x, y)$ for all $y \in S \cup T$. If $x \in T$, the elements of $\bullet x$ and x^\bullet are called *pre-* and *postplaces* of x , respectively, and if $x \in S$ we speak of *pre-* and *posttransitions*. The *token replacement function* $\llbracket _ \rrbracket : T \rightarrow \mathbb{Z}^S$ is given by $\llbracket t \rrbracket = t^\bullet - \bullet t$ for all $t \in T$. These functions extend to finite signed multisets as usual (see Definition 2.1).

Definition 2.4. Let $N = (S, T, F, M_0, \ell)$ be a Petri net, $G \in \mathbb{N}^T$, G non-empty and finite, and $M, M' \in \mathbb{N}^S$.

G is a *step* from M to M' , written $M [G]_N M'$, iff

- $\bullet G \leq M$ (G is *enabled*) and
- $M' = (M - \bullet G) + G^\bullet = M + \llbracket G \rrbracket$.

Note that steps are (finite) multisets, thus allowing self-concurrency, i.e. the same transition can occur multiple times in a single step. We write $M [t]_N M'$ for $M [\{t\}]_N M'$, whereas $M [G]_N$ abbreviates $\exists M'. M [G]_N M'$. We may omit the subscript N if clear from context.

In our nets transitions are labelled with *actions* drawn from a set $\text{Act} \dot{\cup} \{\tau\}$. This makes it possible to see these nets as models of *reactive systems* that interact with their environment. A transition t can be thought of as the occurrence of the action $\ell(t)$. If $\ell(t) \in \text{Act}$, this occurrence can be observed and influenced by the environment—we call such transitions *external* or *visible*, but if $\ell(t) = \tau$, it cannot and t is an *internal* or *silent* transition. Transitions whose occurrences cannot be distinguished by the environment carry the same label. In particular, since the environment cannot observe the occurrence of internal transitions at all, they are all labelled τ .

The labelling function ℓ extends to finite signed multisets of transitions $G \in \mathbb{Z}^T$ by $\ell(G) := \sum_{t \in T} G(t) \cdot \{\ell(t)\}$. For $A, B \in \mathbb{Z}^{\text{Act} \dot{\cup} \tau}$ we write $A \equiv B$ iff $\ell(A)(a) = \ell(B)(a)$ for all $a \in \text{Act}$, i.e. iff A and B contain the same (numbers of) visible actions, allowing $\ell(A)(\tau) \neq \ell(B)(\tau)$. Hence $\ell(G) \equiv \emptyset$ indicates that $\ell(t) = \tau$ for all transitions $t \in T$ with $G(t) \neq 0$.

Definition 2.5. Let $N = (S, T, F, M_0, \ell)$ be a Petri net.

- The set $[M_0]_N$ of *reachable markings* of N is defined as the smallest set containing M_0 that is closed under $[G]_N$, meaning that if $M \in [M_0]_N$ and $M [G]_N M'$ then $M' \in [M_0]_N$.
- N is *one-safe* iff $M \in [M_0]_N \Rightarrow \forall s \in S. M(s) \leq 1$.
- The *concurrency relation* $\smile \subseteq T^2$ is given by $t \smile u \Leftrightarrow \exists M \in [M_0]. M[\{t\} + \{u\}]$.
- N is a *structural conflict net* iff for all $t, u \in T$ with $t \smile u$ we have $\bullet t \cap \bullet u = \emptyset$.

We use the term *plain nets* for Petri nets where ℓ is injective and no transition has the label τ , i.e. essentially unlabelled nets.

This paper first of all aims at studying finite Petri nets: nets with finitely many places and transitions. Additionally, our work also applies to infinite nets with the properties that $\bullet t \neq \emptyset$ for all transitions $t \in T$, and any reachable marking (a) is finite, and (b) enables only finitely many transitions. Henceforth, we call such nets *finitary*. Finitariness can be ensured by requiring $|M_0| < \infty \wedge \forall t \in T. \bullet t \neq \emptyset \wedge \forall x \in S \cup T. |x^\bullet| < \infty$, i.e. that the initial marking is finite, no transition has an empty set of preplaces, and each place and transition has only finitely many outgoing arcs. Our characterisation of distributability pertains to finitary plain structural conflict nets, and our distributed implementations are again structural conflict nets, but they need not be finitary (nor plain). However, our distributed implementations of finite nets are again finite.

3. SEMANTIC EQUIVALENCES

In this section, we give an overview on some semantic equivalences for reactive systems. Most of these may be defined formally for Petri nets in a uniform way, by first defining equivalences for transition systems and then associating different transition systems with a Petri net. This yields in particular different non-interleaving equivalences for Petri nets.

Definition 3.1. Let \mathcal{Act} be a set of *visible actions* and $\tau \notin \mathcal{Act}$ be an *invisible action*. Let $\mathcal{Act}_\tau := \mathcal{Act} \dot{\cup} \{\tau\}$. A *labelled transition system (LTS)* (over \mathcal{Act}_τ) is a triple $(\mathfrak{S}, \mathfrak{T}, \mathfrak{M}_o)$ with

- \mathfrak{S} a set of *states*,
- $\mathfrak{T} \subseteq \mathfrak{S} \times \mathcal{Act}_\tau \times \mathfrak{S}$ a *transition relation*
- and $\mathfrak{M}_o \in \mathfrak{S}$ the *initial state*.

Given an LTS $(\mathfrak{S}, \mathfrak{T}, \mathfrak{M}_o)$ with $\mathfrak{M}, \mathfrak{M}' \in \mathfrak{S}$ and $\alpha \in \mathcal{Act}_\tau$, we write $\mathfrak{M} \xrightarrow{\alpha} \mathfrak{M}'$ for $(\mathfrak{M}, \alpha, \mathfrak{M}') \in \mathfrak{T}$. We write $\mathfrak{M} \xrightarrow{\alpha}$ for $\exists \mathfrak{M}'. \mathfrak{M} \xrightarrow{\alpha} \mathfrak{M}'$ and $\mathfrak{M} \not\xrightarrow{\alpha}$ for $\nexists \mathfrak{M}'. \mathfrak{M} \xrightarrow{\alpha} \mathfrak{M}'$. Furthermore, $\mathfrak{M} \xrightarrow{(\alpha)} \mathfrak{M}'$ denotes $\mathfrak{M} \xrightarrow{\alpha} \mathfrak{M}' \vee (\alpha = \tau \wedge \mathfrak{M} = \mathfrak{M}')$, meaning that in case $\alpha = \tau$ performing a τ -transition is optional. For $a_1 a_2 \cdots a_n \in \mathcal{Act}^*$ we write $\mathfrak{M} \xrightarrow{a_1 a_2 \cdots a_n} \mathfrak{M}'$ when

$$\mathfrak{M} \implies \xrightarrow{a_1} \implies \xrightarrow{a_2} \implies \cdots \implies \xrightarrow{a_n} \implies \mathfrak{M}'$$

where \implies denotes the reflexive and transitive closure of $\xrightarrow{\tau}$. A state $\mathfrak{M} \in \mathfrak{S}$ is said to be *reachable* iff there is a $\sigma \in \mathcal{Act}^*$ such that $\mathfrak{M}_o \xrightarrow{\sigma} \mathfrak{M}$. The set of all reachable states is denoted by $[\mathfrak{M}_o]$. In case there is an infinite sequence of states $(\mathfrak{M}^k)_{k \in \mathbb{N}}$ such that $\mathfrak{M}^0 \in [\mathfrak{M}_o]$ and $\mathfrak{M}^k \xrightarrow{\tau} \mathfrak{M}^{k+1}$ for all $k \in \mathbb{N}$, the LTS is said to display *divergence*.

Many semantic equivalences on LTSs that in some way abstract from internal transitions are defined in the literature; an overview can be found in [vG93]. On divergence-free LTSs, the most discriminating semantics in the spectrum of equivalences of [vG93], and the only one that fully respects the branching structure of related systems, is *branching bisimilarity*, proposed in [GW89].

Definition 3.2. Two LTSs $(\mathfrak{S}_1, \mathfrak{T}_1, \mathfrak{M}_{o1})$ and $(\mathfrak{S}_2, \mathfrak{T}_2, \mathfrak{M}_{o2})$ are *branching bisimilar* iff there exists a relation $\mathcal{B} \subseteq \mathfrak{S}_1 \times \mathfrak{S}_2$ —a *branching bisimulation*—such that, for all $\alpha \in \mathcal{Act}_\tau$:

1. $\mathfrak{M}_{o1} \mathcal{B} \mathfrak{M}_{o2}$;
2. if $\mathfrak{M}_1 \mathcal{B} \mathfrak{M}_2$ and $\mathfrak{M}_1 \xrightarrow{\alpha} \mathfrak{M}'_1$ then $\exists \mathfrak{M}_2^\dagger, \mathfrak{M}'_2$ such that $\mathfrak{M}_2 \implies \mathfrak{M}_2^\dagger \xrightarrow{(\alpha)} \mathfrak{M}'_2$, $\mathfrak{M}_1 \mathcal{B} \mathfrak{M}_2^\dagger$ and $\mathfrak{M}'_1 \mathcal{B} \mathfrak{M}'_2$;
3. if $\mathfrak{M}_1 \mathcal{B} \mathfrak{M}_2$ and $\mathfrak{M}_2 \xrightarrow{\alpha} \mathfrak{M}'_2$ then $\exists \mathfrak{M}_1^\dagger, \mathfrak{M}'_1$ such that $\mathfrak{M}_1 \implies \mathfrak{M}_1^\dagger \xrightarrow{(\alpha)} \mathfrak{M}'_1$, $\mathfrak{M}_1^\dagger \mathcal{B} \mathfrak{M}_2$ and $\mathfrak{M}'_1 \mathcal{B} \mathfrak{M}'_2$.

Branching bisimilarity with explicit divergence [vG93, GW96, GLT09] is a variant of branching bisimilarity that fully respects the diverging behaviour of related systems. It is the most discriminating semantics in the spectrum of equivalences of [vG93].

Definition 3.3. Two LTSs $(\mathfrak{S}_1, \mathfrak{T}_1, \mathfrak{M}_{\circ 1})$ and $(\mathfrak{S}_2, \mathfrak{T}_2, \mathfrak{M}_{\circ 2})$ are branching bisimilar *with explicit divergence* iff there exists a branching bisimulation $\mathcal{B} \subseteq \mathfrak{S}_1 \times \mathfrak{S}_2$ such that furthermore

4. if $\mathfrak{M}_1 \mathcal{B} \mathfrak{M}_2$ and there is an infinite sequence of states $(\mathfrak{M}_1^k)_{k \in \mathbb{N}}$ such that $\mathfrak{M}_1 = \mathfrak{M}_1^0$, $\mathfrak{M}_1^k \xrightarrow{\tau} \mathfrak{M}_1^{k+1}$ and $\mathfrak{M}_1^k \mathcal{B} \mathfrak{M}_2$ for all $k \in \mathbb{N}$, then there exists an infinite sequence of states $(\mathfrak{M}_2^\ell)_{\ell \in \mathbb{N}}$ such that $\mathfrak{M}_2 = \mathfrak{M}_2^0$, $\mathfrak{M}_2^\ell \xrightarrow{\tau} \mathfrak{M}_2^{\ell+1}$ for all $\ell \in \mathbb{N}$, and $\mathfrak{M}_1^k \mathcal{B} \mathfrak{M}_2^\ell$ for all $k, \ell \in \mathbb{N}$;
5. if $\mathfrak{M}_1 \mathcal{B} \mathfrak{M}_2$ and there is an infinite sequence of states $(\mathfrak{M}_2^\ell)_{\ell \in \mathbb{N}}$ such that $\mathfrak{M}_2 = \mathfrak{M}_2^0$, $\mathfrak{M}_2^\ell \xrightarrow{\tau} \mathfrak{M}_2^{\ell+1}$ and $\mathfrak{M}_1 \mathcal{B} \mathfrak{M}_2^\ell$ for all $\ell \in \mathbb{N}$, then there exists an infinite sequence of states $(\mathfrak{M}_1^k)_{k \in \mathbb{N}}$ such that $\mathfrak{M}_1 = \mathfrak{M}_1^0$, $\mathfrak{M}_1^k \xrightarrow{\tau} \mathfrak{M}_1^{k+1}$ for all $k \in \mathbb{N}$, and $\mathfrak{M}_1^k \mathcal{B} \mathfrak{M}_2^\ell$ for all $k, \ell \in \mathbb{N}$.

Since in this paper we mainly compare systems of which one admits no divergence at all, the definition simplifies to the requirement that the other system may not diverge either.

Proposition 3.4. *Let $\mathfrak{L}_1, \mathfrak{L}_2$ be two LTSs, of which \mathfrak{L}_2 does not display divergence. Then \mathfrak{L}_1 and \mathfrak{L}_2 are branching bisimilar with explicit divergence iff \mathfrak{L}_1 and \mathfrak{L}_2 are branching bisimilar and \mathfrak{L}_1 does not display divergence either.*

Proof. “If”: In case neither \mathfrak{L}_1 nor \mathfrak{L}_2 display divergence, any branching bisimulation \mathcal{B} between \mathfrak{L}_1 and \mathfrak{L}_2 , when restricted to the reachable states of \mathfrak{L}_1 and \mathfrak{L}_2 , trivially satisfies Clauses 4 and 5 above.

“Only if”: Suppose that \mathcal{B} is a branching bisimulation between $\mathfrak{L}_1 = (\mathfrak{S}_1, \mathfrak{T}_1, \mathfrak{M}_{\circ 1})$ and $\mathfrak{L}_2 = (\mathfrak{S}_2, \mathfrak{T}_2, \mathfrak{M}_{\circ 2})$ that satisfies Clauses 4 and 5 above, and suppose \mathfrak{L}_1 displays divergence, i.e. there is an infinite sequence of states $(\mathfrak{M}_1^k)_{k \in \mathbb{N}}$ such that $\mathfrak{M}_{\circ 1} \xrightarrow{\sigma} \mathfrak{M}_1^0$ for some $\sigma \in \mathcal{Act}^*$ and $\mathfrak{M}_1^k \xrightarrow{\tau} \mathfrak{M}_1^{k+1}$ for all $k \in \mathbb{N}$. By Definition 3.2, Clauses 1 and 2, there exists an infinite sequence of states $(\mathfrak{M}_2^k)_{k \in \mathbb{N}}$ such that $\mathfrak{M}_{\circ 2} \xrightarrow{\sigma} \mathfrak{M}_2^0$, $\mathfrak{M}_2^k \xrightarrow{\tau} \mathfrak{M}_2^{k+1}$ and $\mathfrak{M}_1^k \mathcal{B} \mathfrak{M}_2^k$ for all $k \in \mathbb{N}$. In case infinitely many of those \mathfrak{M}_2^k are different, this sequence constitutes a divergence of \mathfrak{L}_2 . Otherwise, there is an $k_0 \geq 0$ such that all \mathfrak{M}_2^k for $k \geq k_0$ are equal, and then \mathfrak{L}_2 has a divergence by Clause 4. \square

One of the semantics reviewed in [vG93] that respects branching time and divergence only to a minimal extent, is (*stable*) *failures equivalence*, proposed in [BKO87] and further elaborated in [Ro98]. It is a variant of the failures equivalence of [BHR84], only differing in the treatment of divergence.¹

Definition 3.5. Let $\mathfrak{L} = (\mathfrak{S}, \mathfrak{T}, \mathfrak{M}_{\circ})$ be an LTS, $\sigma \in \mathcal{Act}^*$ and $X \subseteq \mathcal{Act}$, X finite.²

σ is a *trace* of \mathfrak{L} iff $\exists \mathfrak{M}. \mathfrak{M}_{\circ} \xrightarrow{\sigma} \mathfrak{M}$.

$\langle \sigma, X \rangle$ is a *failure pair* of \mathfrak{L} iff $\exists \mathfrak{M}. \mathfrak{M}_{\circ} \xrightarrow{\sigma} \mathfrak{M} \wedge \mathfrak{M} \not\xrightarrow{\tau} \wedge \forall a \in X. \mathfrak{M} \not\xrightarrow{a}$.

¹When comparing two systems without divergence, the stable failure equivalence coincides with the failures equivalence of [BHR84]. When comparing systems of which one is known to be divergence-free—as we will do in this paper—the stable failures semantics is strictly less discriminating than the failures equivalence of [BHR84]—only the latter guarantees that the other system is divergence-free as well. As a less discriminating equivalence will give rise to stronger results about the absence of distributed implementations of certain systems, we will use a version of the stable failures equivalence, rather than of the failures equivalence from [BHR84].

²Although the version without the restriction that X be finite has arguably better properties, we here use the version with this restriction—the *finite failures equivalence* of [vG93]—since it is less discriminating.

We write $\mathfrak{T}(\mathfrak{L})$ for the set of all traces, and $\mathfrak{F}(\mathfrak{L})$ for the set of all failure pairs of \mathfrak{L} . Two LTSs \mathfrak{L}_1 and \mathfrak{L}_2 are *failures equivalent* iff $\mathfrak{T}(\mathfrak{L}_1) = \mathfrak{T}(\mathfrak{L}_2)$ and $\mathfrak{F}(\mathfrak{L}_1) = \mathfrak{F}(\mathfrak{L}_2)$.

As indicated in [GG01], see in particular the diagram on Page 317 (or 88), equivalences on LTSs have been ported to Petri nets and other causality respecting models of concurrency chiefly in five ways: we distinguish *interleaving semantics*, *step semantics*, *split semantics*, *ST-semantics* and *causal semantics*. Causal semantics fully respect the causal relationships between the actions of related systems, whereas interleaving semantics fully abstract from this information. Step semantics differ from interleaving semantics by taking into account the possibility of multiple actions to occur simultaneously (in *one step*); this carries a minimal amount of causal information. ST-semantics respect causality to the extent that it can be expressed in terms of the possibility of durational actions to overlap in time. They are formalised by executing a visible action a in two phases: its start a^+ and its termination a^- . Moreover, terminating actions are properly matched with their starts. Split semantics are a simplification of ST-semantics in which the matching of starts and terminations is dropped.

Interleaving semantics on Petri nets can be formalised by associating to each net $N = (S, T, F, M_0, \ell)$ the LTS $(\mathfrak{S}, \mathfrak{T}, M_0)$ with \mathfrak{S} the set of markings of N and \mathfrak{T} given by

$$M_1 \xrightarrow{\alpha} M_2 :\Leftrightarrow \exists t \in T. \ell(t) = \alpha \wedge M_1 [t] M_2.$$

Here we take $\mathfrak{Act} := \text{Act}$. Now each equivalence on LTSs from [vG93] induces a corresponding interleaving equivalence on nets by declaring two nets equivalent iff the associated LTSs are. For example, *interleaving branching bisimilarity* is the relation of Definition 3.2 with the \mathfrak{M} 's denoting markings, and the α 's actions from Act_τ .

Step semantics on Petri nets can be formalised by associating another LTS to each net. Again we take \mathfrak{S} to be the markings of the net, and \mathfrak{M}_0 the initial marking, but this time \mathfrak{Act} consists of the *steps* over Act , the non-empty, finite multisets A of visible actions from Act , and the transition relation \mathfrak{T} is given by

$$M_1 \xrightarrow{A} M_2 :\Leftrightarrow \exists G \in_F \mathbb{N}^T. \ell(G) = A \wedge \tau \notin \ell(G) \wedge M_1 [G] M_2$$

with τ -transitions defined just as in the interleaving case:

$$M_1 \xrightarrow{\tau} M_2 :\Leftrightarrow \exists t \in T. \ell(t) = \tau \wedge M_1 [t] M_2.$$

In particular, the step version of failures equivalence would be the relation of Definition 3.5 with the \mathfrak{M} 's denoting markings, the a 's steps over Act , the X 's sets of steps, and the σ 's sequences of steps. This form of step failures semantics, but based on the failures semantics of [BHR84] rather than the stable failures semantics of Definition 3.5, has been studied in [TV89]. However, variations in this type of definition are possible. In this paper we employ a form of step failures semantics that is a bit closer to interleaving semantics, thereby coarsening the equivalence and strengthening the final result: σ is a sequence of single actions, whereas the set X of impossible continuations after σ is a set of steps. Moreover, we drop the comparison of the sets of traces. We define this notion directly on Petri nets, without using intermediate LTSs.

Definition 3.6. Let $N = (S, T, F, M_0, \ell)$ be a Petri net, $\sigma \in \text{Act}^*$ and $X \subseteq \mathbb{N}^{\text{Act}}$, X finite. $\langle \sigma, X \rangle$ is a *step failure pair* of N iff

$$\exists M. M_0 \xrightarrow{\sigma} M \wedge M \not\xrightarrow{\tau} \wedge \forall A \in X. M \not\xrightarrow{A}.$$

We write $\mathcal{F}(N)$ for the set of all step failure pairs of N .

Two Petri nets N_1 and N_2 are *step failures equivalent*, $N_1 \approx_{\mathcal{F}} N_2$, iff $\mathcal{F}(N_1) = \mathcal{F}(N_2)$.

Next we propose a general definition on Petri nets of ST-versions of each of the semantics of [vG93]. Again we do this through a mapping from nets to a suitable LTS. An *ST-marking* of a net (S, T, F, M_0, ℓ) is a pair $(M, U) \in \mathbb{N}^S \times T^*$ of a normal marking, together with a sequence of visible transitions *currently firing*. The *initial* ST-marking is $\mathfrak{M}_0 := (M_0, \epsilon)$. The elements of $\text{Act}^\pm := \{a^+, a^{-n} \mid a \in \text{Act}, n > 0\}$ are called *visible action phases*, and $\text{Act}_\tau^\pm := \text{Act}^\pm \dot{\cup} \{\tau\}$. For $U \in T^*$, we write $t \in^{(n)} U$ if t is the n^{th} element of U . Furthermore U^{-n} denotes U after removal of the n^{th} transition.

Definition 3.7. Let $N = (S, T, F, M_0, \ell)$ be a Petri net, labelled over Act_τ .

The *ST-transition relations* $\xrightarrow{\eta}$ for $\eta \in \text{Act}_\tau^\pm$ between ST-markings are given by

$$(M, U) \xrightarrow{a^+} (M', U') \text{ iff } \exists t \in T. \ell(t) = a \wedge M[t] \wedge M' = M - \bullet t \wedge U' = Ut.$$

$$(M, U) \xrightarrow{a^{-n}} (M', U') \text{ iff } \exists t \in^{(n)} U. \ell(t) = a \wedge U' = U^{-n} \wedge M' = M + t \bullet.$$

$$(M, U) \xrightarrow{\tau} (M', U') \text{ iff } M \xrightarrow{\tau} M' \wedge U' = U.$$

Now the ST-LTS associated to a net N is $(\mathfrak{S}, \mathfrak{T}, \mathfrak{M}_0)$ with \mathfrak{S} the set of ST-markings of N , $\mathfrak{Act} := \text{Act}^\pm$, \mathfrak{T} as defined in Definition 3.7, and \mathfrak{M}_0 the initial ST-marking. Again, each equivalence on LTSs from [vG93] induces a corresponding ST-equivalence on nets by declaring two nets equivalent iff their associated LTSs are. In particular, *branching ST-bisimilarity* is the relation of Definition 3.2 with the \mathfrak{M} 's denoting ST-markings, and the α 's action phases from Act_τ^\pm . We write $N_1 \approx_{bSTb}^\Delta N_2$ iff N_1 and N_2 are branching ST-bisimilar with explicit divergence.

ST-bisimilarity was originally proposed in [GV87]. It was extended to a setting with internal actions in [Vo93], based on the notion of *weak bisimilarity* of [Mi89], which is a bit less discriminating than branching bisimilarity. The above can be regarded as a reformulation of the same idea; the notion of weak ST-bisimilarity defined according to the recipe above agrees with the ST-bisimilarity of [Vo93].

The next proposition says that branching ST-bisimilarity with explicit divergence is more discriminating than (i.e. *stronger* than, *finer* than, or included in) step failures equivalence.

Proposition 3.8. *Let N_1 and N_2 be Petri nets. If $N_1 \approx_{bSTb}^\Delta N_2$ then $N_1 \approx_{\mathcal{F}} N_2$.*

Proof. Suppose $N_1 \approx_{bSTb}^\Delta N_2$ and $\langle \sigma, X \rangle \in \mathcal{F}(N_1)$. By symmetry it suffices to show that $\langle \sigma, X \rangle \in \mathcal{F}(N_2)$.

Since $N_1 \approx_{bSTb}^\Delta N_2$, there must be a branching bisimulation \mathcal{B} between the ST-markings of $N_1 = (S_1, T_1, F_1, M_{01}, \ell_1)$ and $N_2 = (S_2, T_2, F_2, M_{02}, \ell_2)$. In particular, $(M_{01}, \epsilon) \mathcal{B} (M_{02}, \epsilon)$. Let $\sigma =: a_1 a_2 \cdots a_n \in \text{Act}^*$. Then $M_{01} \xRightarrow{a_1} \xRightarrow{a_2} \xRightarrow{\cdots} \xRightarrow{a_n} M'_1$ for a marking $M'_1 \in \mathbb{N}^{S_1}$ with $M'_1 \not\xrightarrow{\tau}$ and $\forall A \in X. M'_1 \not\xrightarrow{A}$. So $(M_{01}, \epsilon) \xRightarrow{a_1^+} \xRightarrow{a_1^{-1}} \xRightarrow{a_2^+} \xRightarrow{a_2^{-1}} \xRightarrow{\cdots} \xRightarrow{a_n^+} \xRightarrow{a_n^{-1}} \xRightarrow{\cdots} \xRightarrow{a_n^+} \xRightarrow{a_n^{-1}} \xRightarrow{\cdots} \xRightarrow{a_n^+} \xRightarrow{a_n^{-1}} \xRightarrow{\cdots} (M'_1, \epsilon)$. Thus, using the properties of a branching bisimulation on the ST-LTSs associated to N_1 and N_2 , there must be a marking $M'_2 \in \mathbb{N}^{S_2}$ such that $(M_{02}, \epsilon) \xRightarrow{a_1^+} \xRightarrow{a_1^{-1}} \xRightarrow{a_2^+} \xRightarrow{a_2^{-1}} \xRightarrow{\cdots} \xRightarrow{a_n^+} \xRightarrow{a_n^{-1}} \xRightarrow{\cdots} (M'_2, \epsilon)$ and $(M'_1, \epsilon) \mathcal{B} (M'_2, \epsilon)$. Since $(M'_1, \epsilon) \not\xrightarrow{\tau}$, the ST-marking (M'_1, ϵ) admits no divergence. As \approx_{bSTb}^Δ respects this property (cf. the proof of Proposition 3.4), also (M'_2, ϵ) admits no divergence, and there must be an $M''_2 \in \mathbb{N}^{S_2}$ with $M''_2 \not\xrightarrow{\tau}$ and $(M'_2, \epsilon) \xRightarrow{\sigma} (M''_2, \epsilon)$. Clause 3. of a branching bisimulation gives $(M'_1, \epsilon) \mathcal{B} (M''_2, \epsilon)$, and Definition 3.7 yields $M_{02} \xRightarrow{\sigma} M''_2$. Here we use that if $(M, U) \xrightarrow{a^{-1}} \xrightarrow{\tau} (M', U')$ then $(M, U) \xrightarrow{\tau} \xrightarrow{a^{-1}} (M', U')$.

Now let $B = \{b_1, \dots, b_m\} \in X$. Then $M'_1 \not\stackrel{B}{\rightarrow}$. Suppose, towards a contradiction, that $M''_2 \stackrel{B}{\rightarrow}$. Then $(M''_2, \epsilon) \xrightarrow{b_1^+} \xrightarrow{b_2^+} \dots \xrightarrow{b_m^+}$. Property 2. of a branching bisimulation implies $(M'_1, \epsilon) \xrightarrow{b_1^+} \xrightarrow{b_2^+} \dots \xrightarrow{b_m^+}$ and hence $M'_1 \stackrel{B}{\rightarrow}$. This is a contradiction, so $M''_2 \not\stackrel{B}{\rightarrow}$. It follows that $\langle \sigma, X \rangle \in \mathcal{F}(N_2)$. \square

In this paper we employ both step failures equivalence and branching ST-bisimilarity with explicit divergence. Fortunately it will turn out that for our purposes the latter equivalence coincides with its split version (since always one of the compared nets is plain, see Proposition 3.15).

A *split marking* of a net $N = (S, T, F, M_0, \ell)$ is a pair $(M, U) \in \mathbb{N}^S \times \mathbb{N}^T$ of a normal marking M , together with a multiset of visible transitions currently firing. The *initial* split marking is $\mathfrak{M}_0 := (M_0, \emptyset)$. A split marking can be regarded as an abstraction from an ST-marking, in which the total order on the (finite) multiset of transitions that are currently firing has been dropped. Let $\text{Act}_{\text{split}}^\pm := \{a^+, a^- \mid a \in \text{Act}\}$.

Definition 3.9. Let $N = (S, T, F, M_0, \ell)$ be a Petri net, labelled over Act_τ .

The *split transition relations* $\xrightarrow{\zeta}$ for $\zeta \in \text{Act}_{\text{split}}^\pm \dot{\cup} \{\tau\}$ between split markings are given by

$$\begin{aligned} (M, U) &\xrightarrow{a^+} (M', U') \text{ iff } \exists t \in T. \ell(t) = a \wedge M[t] \wedge M' = M - \bullet t \wedge U' = U + \{t\}. \\ (M, U) &\xrightarrow{a^-} (M', U') \text{ iff } \exists t \in U. \ell(t) = a \wedge U' = U - \{t\} \wedge M' = M + t^\bullet. \\ (M, U) &\xrightarrow{\tau} (M', U') \text{ iff } M \xrightarrow{\tau} M' \wedge U' = U. \end{aligned}$$

Note that $(M, U) \xrightarrow{a^+}$ iff $M \xrightarrow{a}$, whereas $(M, U) \xrightarrow{a^-}$ iff $a \in \ell(U)$. With induction on reachability of markings it is furthermore easy to check that $(M, U) \in [\mathfrak{M}_0]$ iff $\tau \notin \ell(U)$ and $M + \bullet U \in [M_0]$.

The split LTS associated to a net N is $(\mathfrak{S}, \mathfrak{T}, \mathfrak{M}_0)$ with \mathfrak{S} the set of split markings of N , $\mathfrak{Act} := \text{Act}^\pm$, \mathfrak{T} as defined in Definition 3.9, and \mathfrak{M}_0 the initial split marking. Again, each equivalence on LTSs from [vG93] induces a corresponding split equivalence on nets by declaring two nets equivalent iff their associated LTSs are. In particular, *branching split bisimilarity* is the relation of Definition 3.2 with the \mathfrak{M} 's denoting split markings, and the α 's action phases from $\text{Act}_{\text{split}}^\pm \dot{\cup} \{\tau\}$.

For $\mathfrak{M} = (M, U) \in \mathbb{N}^S \times T^*$ an ST-marking, let $\overline{\mathfrak{M}} = (M, \overline{U}) \in \mathbb{N}^S \times \mathbb{N}^T$ be the split marking obtained by converting the sequence U into the multiset \overline{U} , where $\overline{U}(t)$ is the number of occurrences of the transition $t \in T$ in U . Moreover, define $\ell(\mathfrak{M})$ by $\ell(M, U) := \ell(U)$ and $\ell(t_1 t_2 \dots t_k) := \ell(t_1) \ell(t_2) \dots \ell(t_k)$. Furthermore, for $\eta \in \text{Act}_\tau^\pm$, let $\overline{\eta} \in \text{Act}_{\text{split}}^\pm \dot{\cup} \{\tau\}$ be given by $\overline{a^+} := a^+$, $\overline{a^-} := a^-$ and $\overline{\tau} := \tau$.

Observation 3.10. Let $\mathfrak{M}, \mathfrak{M}'$ be ST-markings, \mathfrak{M}^\dagger a split marking, $\eta \in \text{Act}_\tau^\pm$ and $\zeta \in \text{Act}_{\text{split}}^\pm \dot{\cup} \{\tau\}$. Then

- (1) $\mathfrak{M} \in \mathbb{N}^S \times T^*$ is the initial ST-marking of N iff $\overline{\mathfrak{M}} \in \mathbb{N}^S \times \mathbb{N}^T$ is the initial split marking of N ;
- (2) if $\mathfrak{M} \xrightarrow{\eta} \mathfrak{M}'$ then $\overline{\mathfrak{M}} \xrightarrow{\overline{\eta}} \overline{\mathfrak{M}'}$;
- (3) if $\overline{\mathfrak{M}} \xrightarrow{\zeta} \mathfrak{M}^\dagger$ then there is a $\mathfrak{M}' \in \mathbb{N}^S \times T^*$ and $\eta \in \text{Act}_\tau^\pm$ such that $\mathfrak{M} \xrightarrow{\eta} \mathfrak{M}'$, $\overline{\eta} = \zeta$ and $\overline{\mathfrak{M}'} = \mathfrak{M}^\dagger$;
- (4) if $\mathfrak{M} \xrightarrow{(\eta)} \mathfrak{M}'$ then $\overline{\mathfrak{M}} \xrightarrow{(\overline{\eta})} \overline{\mathfrak{M}'}$;
- (5) if $\overline{\mathfrak{M}} \xrightarrow{(\zeta)} \mathfrak{M}^\dagger$ then there is a $\mathfrak{M}' \in \mathbb{N}^S \times T^*$ and $\eta \in \text{Act}_\tau^\pm$ such that $\mathfrak{M} \xrightarrow{(\eta)} \mathfrak{M}'$, $\overline{\eta} = \zeta$ and $\overline{\mathfrak{M}'} = \mathfrak{M}^\dagger$;
- (6) if $\mathfrak{M} \Longrightarrow \mathfrak{M}'$ then $\overline{\mathfrak{M}} \Longrightarrow \overline{\mathfrak{M}'}$;
- (7) if $\overline{\mathfrak{M}} \Longrightarrow \mathfrak{M}^\dagger$ then there is a $\mathfrak{M}' \in \mathbb{N}^S \times T^*$ such that $\mathfrak{M} \Longrightarrow \mathfrak{M}'$ and $\overline{\mathfrak{M}'} = \mathfrak{M}^\dagger$. \square

Lemma 3.11. *Let $N_1 = (S_1, T_1, F_1, M_{01}, \ell)$ and $N_2 = (S_2, T_2, F_2, M_{02}, \ell_2)$ be two nets, N_2 being plain; let $\mathfrak{M}_1, \mathfrak{M}'_1$ be ST-markings of N_1 , and $\mathfrak{M}_2, \mathfrak{M}'_2$ ST-markings of N_2 . If $\ell(\mathfrak{M}_2) = \ell(\mathfrak{M}_1)$, $\mathfrak{M}_1 \xrightarrow{\eta} \mathfrak{M}'_1$ and $\mathfrak{M}_2 \xrightarrow{(\eta')} \mathfrak{M}'_2$ with $\overline{\eta'} = \overline{\eta}$, then there is an \mathfrak{M}''_2 with $\mathfrak{M}_2 \xrightarrow{(\eta)} \mathfrak{M}''_2$, $\ell(\mathfrak{M}''_2) = \ell(\mathfrak{M}'_1)$, and $\overline{\mathfrak{M}''_2} = \overline{\mathfrak{M}'_2}$.*

Proof. If $\mathfrak{M} \xrightarrow{\eta} \mathfrak{M}'$ or $\mathfrak{M} \xrightarrow{(\eta)} \mathfrak{M}'$ then $\ell(\mathfrak{M}')$ is completely determined by $\ell(\mathfrak{M})$ and η . For this reason the requirement $\ell(\mathfrak{M}''_2) = \ell(\mathfrak{M}'_1)$ will hold as soon as the other requirements are met.

First suppose η is of the form τ or a^+ . Then $\overline{\eta} = \eta$ and moreover $\overline{\eta'} = \overline{\eta}$ implies $\eta' = \eta$. Thus we can take $\mathfrak{M}''_2 := \mathfrak{M}'_2$.

Now suppose $\eta = a^{-n}$ for some $n > 0$. Then $\eta' = a^{-m}$ for some $m > 0$. As $\mathfrak{M}_1 \xrightarrow{\eta}$, the n^{th} element of $\ell(\mathfrak{M}_1)$ must (exist and) be a . Since $\ell(\mathfrak{M}_2) = \ell(\mathfrak{M}_1)$, also the n^{th} element of $\ell(\mathfrak{M}_2)$ must be a , so there is an \mathfrak{M}''_2 with $\mathfrak{M}_2 \xrightarrow{(\eta)} \mathfrak{M}''_2$. Let $\mathfrak{M}_2 := (M_2, U_2)$. Then U_2 is a sequence of transitions of which the n^{th} and the m^{th} elements are both labelled a . Since the net N_2 is plain, those two transitions must be equal. Let $\mathfrak{M}'_2 := (M'_2, U'_2)$ and $\mathfrak{M}''_2 := (M''_2, U''_2)$. We find that $M'_2 = M''_2$ and $\overline{U'_2} = \overline{U''_2}$. It follows that $\overline{\mathfrak{M}'_2} = \overline{\mathfrak{M}''_2}$. \square

Observation 3.12. If $\mathfrak{M} \implies \mathfrak{M}'$ for ST-markings $\mathfrak{M}, \mathfrak{M}'$ then $\ell(\mathfrak{M}') = \ell(\mathfrak{M})$.

Observation 3.13. If $\ell(\mathfrak{M}_1) = \ell(\mathfrak{M}_2)$ and $\mathfrak{M}_2 \xrightarrow{a^{-n}}$ for some $a \in \text{Act}$ and $n > 0$, then $\mathfrak{M}_1 \xrightarrow{a^{-n}}$.

Observation 3.14. If $\mathfrak{M} \xrightarrow{a^{-n}} \mathfrak{M}'$ and $\mathfrak{M} \xrightarrow{a^{-n}} \mathfrak{M}''$ for some $a \in \text{Act}$ and $n > 0$, then $\mathfrak{M}' = \mathfrak{M}''$.

Proposition 3.15. *Let $N_1 = (S_1, T_1, F_1, M_{01}, \ell)$ and $N_2 = (S_2, T_2, F_2, M_{02}, \ell_2)$ be two nets, N_2 being plain. Then N_1 and N_2 are branching ST-bisimilar (with explicit divergence) iff they are branching split bisimilar (with explicit divergence).*

Proof. Suppose \mathcal{B} is a branching ST-bisimulation between N_1 and N_2 . Then, by Observation 3.10, the relation $\mathcal{B}_{\text{split}} := \{(\overline{\mathfrak{M}_1}, \overline{\mathfrak{M}_2}) \mid (\mathfrak{M}_1, \mathfrak{M}_2) \in \mathcal{B}\}$ is a branching split bisimulation between N_1 and N_2 .

Now let \mathcal{B} be a branching split bisimulation between N_1 and N_2 . Then, using Observation 3.10, the relation $\mathcal{B}_{\text{ST}} := \{(\mathfrak{M}_1, \mathfrak{M}_2) \mid \ell_1(\mathfrak{M}_1) = \ell_2(\mathfrak{M}_2) \wedge (\overline{\mathfrak{M}_1}, \overline{\mathfrak{M}_2}) \in \mathcal{B}\}$ turns out to be a branching ST-bisimulation between N_1 and N_2 :

1. $\mathfrak{M}_{\circ 1} \mathcal{B}_{\text{ST}} \mathfrak{M}_{\circ 2}$ follows from Observation 3.10(1), since $\overline{\mathfrak{M}_{\circ 1}} \mathcal{B} \overline{\mathfrak{M}_{\circ 2}}$ and $\ell(\mathfrak{M}_{\circ 1}) = \ell(\mathfrak{M}_{\circ 2}) = \epsilon$.
2. Suppose $\mathfrak{M}_1 \mathcal{B}_{\text{ST}} \mathfrak{M}_2$ and $\mathfrak{M}_1 \xrightarrow{\eta} \mathfrak{M}'_1$. Then $\overline{\mathfrak{M}_1} \mathcal{B} \overline{\mathfrak{M}_2}$ and $\overline{\mathfrak{M}_1} \xrightarrow{\overline{\eta}} \overline{\mathfrak{M}'_1}$. Hence $\exists \mathfrak{M}^\dagger_2, \mathfrak{M}^\ddagger_2$ such that $\overline{\mathfrak{M}_2} \implies \mathfrak{M}^\dagger_2 \xrightarrow{(\overline{\eta})} \mathfrak{M}^\ddagger_2$, $\overline{\mathfrak{M}_1} \mathcal{B} \mathfrak{M}^\dagger_2$ and $\overline{\mathfrak{M}'_1} \mathcal{B} \mathfrak{M}^\ddagger_2$. As N_2 is plain, $\mathfrak{M}^\dagger_2 = \overline{\mathfrak{M}_2}$. By Observation 3.10(5), using that $\overline{\mathfrak{M}_2} \xrightarrow{(\overline{\eta})} \mathfrak{M}^\ddagger_2$, $\exists \mathfrak{M}''_2, \eta'$ such that $\mathfrak{M}_2 \xrightarrow{(\eta')} \mathfrak{M}''_2$, $\overline{\eta'} = \overline{\eta}$ and $\overline{\mathfrak{M}''_2} = \mathfrak{M}^\ddagger_2$. By Lemma 3.11, there is an ST-marking \mathfrak{M}''_2 such that $\mathfrak{M}_2 \xrightarrow{(\eta)} \mathfrak{M}''_2$, $\ell(\mathfrak{M}''_2) = \ell(\mathfrak{M}'_1)$, and $\overline{\mathfrak{M}''_2} = \overline{\mathfrak{M}'_2} = \mathfrak{M}^\ddagger_2$. It follows that $\mathfrak{M}'_1 \mathcal{B}_{\text{ST}} \mathfrak{M}''_2$.
3. Suppose $\mathfrak{M}_1 \mathcal{B}_{\text{ST}} \mathfrak{M}_2$ and $\mathfrak{M}_2 \xrightarrow{\eta} \mathfrak{M}'_2$. Then $\overline{\mathfrak{M}_1} \mathcal{B} \overline{\mathfrak{M}_2}$ and $\overline{\mathfrak{M}_2} \xrightarrow{\overline{\eta}} \overline{\mathfrak{M}'_2}$. Hence $\exists \mathfrak{M}^\dagger_1, \mathfrak{M}^\ddagger_1$ such that $\overline{\mathfrak{M}_1} \implies \mathfrak{M}^\dagger_1 \xrightarrow{(\overline{\eta})} \mathfrak{M}^\ddagger_1$, $\mathfrak{M}^\dagger_1 \mathcal{B} \overline{\mathfrak{M}_2}$ and $\mathfrak{M}^\ddagger_1 \mathcal{B} \overline{\mathfrak{M}'_2}$. By Observation 3.10(7), $\exists \mathfrak{M}^*_1$ such that $\mathfrak{M}_1 \implies \mathfrak{M}^*_1$ and $\overline{\mathfrak{M}^*_1} = \mathfrak{M}^\ddagger_1$. By Observation 3.12, $\ell(\mathfrak{M}^*_1) = \ell(\mathfrak{M}_1) = \ell(\mathfrak{M}_2)$, so $\mathfrak{M}^*_1 \mathcal{B}_{\text{ST}} \mathfrak{M}_2$. Since N_2 is plain, $\eta \neq \tau$.
 - Let $\eta = a^+$ for some $a \in \text{Act}$. Using that $\overline{\mathfrak{M}^*_1} \xrightarrow{(\overline{\eta})} \mathfrak{M}^\ddagger_1$, by Observation 3.10(5) $\exists \mathfrak{M}'_1, \eta'$ such that $\mathfrak{M}^*_1 \xrightarrow{(\eta')} \mathfrak{M}'_1$, $\overline{\eta'} = \overline{\eta}$ and $\overline{\mathfrak{M}'_1} = \mathfrak{M}^\ddagger_1$. It must be that $\eta' = \eta = a^+$ and $\ell(\mathfrak{M}'_1) = \ell(\mathfrak{M}^*_1)a = \ell(\mathfrak{M}_2)a = \ell(\mathfrak{M}'_2)$. Hence $\mathfrak{M}'_1 \mathcal{B}_{\text{ST}} \mathfrak{M}'_2$.

- Let $\eta = a^{-n}$ for some $a \in \text{Act}$ and $n > 0$. By Observation 3.13, $\exists \mathfrak{M}'_1$ with $\mathfrak{M}'_1 \xrightarrow{\eta} \mathfrak{M}'_1$. By Part 2. of this proof, $\exists \mathfrak{M}''_2$ such that $\mathfrak{M}_2 \xrightarrow{(\eta)} \mathfrak{M}''_2$ and $\mathfrak{M}'_1 \mathcal{B}_{\text{ST}} \mathfrak{M}''_2$. By Observation 3.14 $\mathfrak{M}''_2 = \mathfrak{M}'_2$.

Since the net N_2 is plain, it has no divergence. In such a case, the requirement “with explicit divergence” requires N_1 to be free of divergence as well, regardless of whether split or ST-semantics is used. \square

In this paper we will not consider causal semantics. The reason is that our distributed implementations will not fully preserve the causal behaviour of nets. We will further comment on this in the conclusion.

4. DISTRIBUTED SYSTEMS

In this section, we stipulate what we understand by a distributed system, and subsequently formalise a model of distributed systems in terms of Petri nets.

- A distributed system consists of components residing on different locations.
- Components work concurrently.
- Interactions between components are only possible by explicit communications.
- Communication between components is time consuming and asynchronous.

Asynchronous communication is the only interaction mechanism in a distributed system for exchanging signals or information.

- The sending of a message happens always strictly before its receipt (there is a causal relation between sending and receiving a message).
- A sending component sends without regarding the state of the receiver; in particular there is no need to synchronise with a receiving component. After sending the sender continues its behaviour independently of receipt of the message.

As explained in the introduction, we will add another requirement to our notion of a distributed system, namely that its components only allow sequential behaviour.

4.1. LSGA nets. Formally, we model distributed systems as nets consisting of component nets with sequential behaviour and interfaces in terms of input and output places.

Definition 4.1. Let $N = (S, T, F, M_0, \ell)$ be a Petri net, $I, O \subseteq S$, $I \cap O = \emptyset$ and $O^\bullet = \emptyset$.

1. (N, I, O) is a *component with interface* (I, O) .
2. (N, I, O) is a *sequential component with interface* (I, O) iff $\exists Q \subseteq S \setminus (I \cup O)$ with $\forall t \in T: |\bullet t \upharpoonright Q| = 1 \wedge |t^\bullet \upharpoonright Q| = 1$ and $|M_0 \upharpoonright Q| = 1$.

An input place $i \in I$ of a component $\mathcal{C} = (N, I, O)$ can be regarded as a mailbox of \mathcal{C} for a specific type of messages. An output place $o \in O$, on the other hand, is an address outside \mathcal{C} to which \mathcal{C} can send messages. Moving a token into o is like posting a letter. The condition $o^\bullet = \emptyset$ says that a message, once posted, cannot be retrieved by the component.³

A set of places like Q above is a special case of an *S-invariant*. The requirements guarantee that the number of tokens in these places remains constant, in this case 1. It

³We could have required that $\bullet I = \emptyset$, thereby disallowing a component to put messages in its own mailbox. This would not lead to a loss of generality in the class of distributed systems that can be obtained as the asynchronous parallel composition of sequential components, defined below. However, this property is not preserved under asynchronous parallel composition (defined below), and we like the composition of a set of (sequential) components to be a component itself (but not a sequential one).

follows that no two transitions can ever fire concurrently (in one step). Conversely, whenever a net is sequential, in the sense that no two transitions can fire in one step, it is easily converted into a behaviourally equivalent net with the required S -invariant, namely by adding a single marked place with a self-loop to all transitions. This modification preserves virtually all semantic equivalences on Petri nets from the literature, including \approx_{bSTb}^Δ .

Next we define an operator for combining components with asynchronous communication by fusing input and output places.

Definition 4.2. Let \mathfrak{K} be an index set.

Let $((S_k, T_k, F_k, M_{0k}, \ell_k), I_k, O_k)$ with $k \in \mathfrak{K}$ be components with interface such that $(S_k \cup T_k) \cap (S_l \cup T_l) = (I_k \cup O_k) \cap (I_l \cup O_l)$ for all $k, l \in \mathfrak{K}$ with $k \neq l$ (components are disjoint except for interface places) and $I_k \cap I_l = \emptyset$ for all $k, l \in \mathfrak{K}$ with $k \neq l$ (mailboxes cannot be shared; any message has a unique recipient).

Then the *asynchronous parallel composition* of these components is defined by

$$\parallel_{i \in \mathfrak{K}} ((S_k, T_k, F_k, M_{0k}, \ell_k), I_k, O_k) = ((S, T, F, M_0, \ell), I, O)$$

with $S = \bigcup_{k \in \mathfrak{K}} S_k$, $T = \bigcup_{k \in \mathfrak{K}} T_k$, $F = \bigcup_{k \in \mathfrak{K}} F_k$, $M_0 = \sum_{k \in \mathfrak{K}} M_{0k}$, $\ell = \bigcup_{k \in \mathfrak{K}} \ell_k$ (componentwise union of all nets), $I = \bigcup_{k \in \mathfrak{K}} I_k$ (we accept additional inputs from outside), and $O = \bigcup_{k \in \mathfrak{K}} O_k \setminus \bigcup_{k \in \mathfrak{K}} I_k$ (once fused with an input, $o \in O_I$ is no longer an output).

Note that the asynchronous parallel composition of components with interfaces is again a component with interface.

Observation 4.3. \parallel is associative.

This follows directly from the associativity of the (multi)set union operator. \square

We are now ready to define the class of nets representing systems of asynchronously communicating sequential components.

Definition 4.4. A Petri net N is an *LSGA net* (a *locally sequential globally asynchronous net*) iff there exists an index set \mathfrak{K} and sequential components with interface \mathcal{C}_k , $k \in \mathfrak{K}$, such that $(N, I, O) = \parallel_{k \in \mathfrak{K}} \mathcal{C}_k$ for some I and O .

Up to \approx_{bSTb}^Δ —or any reasonable equivalence preserving causality and branching time but abstracting from internal activity—the same class of LSGA systems would have been obtained if we had imposed, in Definition 4.1 of sequential components, that I , O and Q form a partition of S and that $\bullet I = \emptyset$.⁴ However, it is essential that our definition allows multiple transitions of a component to read from the same input place.

⁴First of all, any $i \in I$ with $\bullet i \neq \emptyset$ can be split into a pure input place, receiving tokens only from outside the component, and an internal place, which is the target of all arcs that used to go to i . Any transition t with $i \in \bullet t$ now needs to be split into one that takes its input token from the pure input place and one that takes it from the internal incarnation of i . In fact, if $F(i, t) = n$ then t needs to be split into $n+1$ copies. The result of this transformation is that $\bullet I = \emptyset$.

Next, any component $\mathcal{C} = ((S, T, F, M_0, \ell), I, O)$ with $\bullet I = \emptyset$ can be replaced by an equivalent component $((S', T', F', M'_0, \ell'), I, O)$ whose places S' are $I \dot{\cup} O \dot{\cup} Q$, where Q is the set of markings of \mathcal{C} , each restricted to the places outside I and O . For each transition t and markings M, M' of the component such that $M \upharpoonright t M'$, writing $q := M \upharpoonright (S \setminus (I \cup O))$ and $q' := M' \upharpoonright (S \setminus (I \cup O))$, there will be a transition $t_q \in T'$ with $F'(i, t_q) = F(i, t)$ for all $i \in I$, $F'(t_q, o) = F(t, o)$ for all $o \in O$, $F'(q, t) = F'(t, q') = 1$, and $F'(p, t) = F'(t, p) = 0$ otherwise. Moreover, $\ell'(t_q) = \ell(t)$ and M'_0 consists of the single place $M_0 \upharpoonright (S \setminus (I \cup O))$. This component clearly has the required properties.

4.2. Distributed nets. In the remainder of this section we give a more abstract characterisation of Petri nets representing distributed systems, namely as *distributed* Petri nets, which we introduced in [GGS08]. This will be useful in Section 5, where we investigate distributability using this more semantic characterisation. We show below that the concrete characterisation of distributed systems as LSGA nets and this abstract characterisation agree.

Following [BCD02], to arrive at a class of nets representing distributed systems, we associate *localities* to the elements of a net $N = (S, T, F, M_0, \ell)$. We model this by a function $D : S \cup T \rightarrow \text{Loc}$, with Loc a set of possible locations. We refer to such a function as a *distribution* of N . Since the identity of the locations is irrelevant for our purposes, we can just as well abstract from Loc and represent D by the equivalence relation \equiv_D on $S \cup T$ given by $x \equiv_D y$ iff $D(x) = D(y)$.

Following [GGS08], we impose a fundamental restriction on distributions, namely that when two transitions can occur in one step, they cannot be co-located. This reflects our assumption that at a given location actions can only occur sequentially.

In [GGS08] we observed that Petri nets incorporate a notion of synchronous interaction, in that a transition can fire only by synchronously taking the tokens from all of its preplaces. In general the behaviour of a net would change radically if a transition would take its input tokens one by one—in particular deadlocks may be introduced. Therefore we insist that in a distributed Petri net, a transition and all its input places reside on the same location. There is no reason to require the same for the output places of a transition, for the behaviour of a net would not change significantly if transitions were to deposit their output tokens one by one [GGS08].

This leads to the following definition of a distributed Petri net.

Definition 4.5. [GGS08] A Petri net $N = (S, T, F, M_0, \ell)$ is *distributed* iff there exists a distribution D such that

- (1) $\forall s \in S, t \in T. s \in \bullet t \Rightarrow t \equiv_D s,$
- (2) $\forall t, u \in T. t \smile u \Rightarrow t \not\equiv_D u.$

A typical example of a net which is not distributed is shown in Figure 5 on Page 19. Transitions t and v are concurrently executable and hence should be placed on different locations. However, both have preplaces in common with u which would enforce putting all three transitions on the same location. In fact, distributed nets can be characterised in the following semi-structural way.

Observation 4.6. A Petri net is distributed iff there is no sequence t_0, \dots, t_n of transitions with $t_0 \smile t_n$ and $\bullet t_{i-1} \cap \bullet t_i \neq \emptyset$ for $i = 1, \dots, n$. \square

Since a structural conflict net is defined as a net without such a sequence with $n = 1$ (cf. Definition 2.5), we obtain:

Observation 4.7. Every distributed Petri net is a structural conflict net. \square

Further on, we use a more liberal definition of a distributed net, called *essentially distributed*. We will show that up to \approx_{bSTb}^Δ any essentially distributed net can be converted into a distributed net. In [GGS08] we employed an even more liberal definition of a distributed net, which we call here *externally distributed*. Although we showed that up to step failures equivalence any externally distributed net can be converted into a distributed net, this does not hold for \approx_{bSTb}^Δ .

Definition 4.8. A net $N = (S, T, F, M_0, \ell)$ is *essentially distributed* iff there exists a distribution D satisfying (1) of Definition 4.5 and

$$(2') \quad \forall t, u \in T. t \smile u \wedge \ell(t) \neq \tau \Rightarrow t \not\equiv_D u.$$

It is *externally distributed* iff there exists a distribution D satisfying (1) and

$$(2'') \quad \forall t, u \in T. t \smile u \wedge \ell(t), \ell(u) \neq \tau \Rightarrow t \not\equiv_D u.$$

Instead of ruling out co-location of concurrent transitions in general, essentially distributed nets permit concurrency of internal transitions—labelled τ —at the same location. Externally distributed nets even allow concurrency between visible and silent transitions at the same location. If the transitions t and v in the net of Figure 5 would both be labelled τ , the net would be essentially distributed, although not distributed; in case only v would be labelled τ the net would be externally distributed but not essentially distributed. Essentially distributed nets need not be structural conflict nets; in fact, *any* net without visible transitions is essentially distributed.

Definition 4.9. Given any Petri net N , the *canonical co-location relation* \equiv_C on N is the equivalence relation on the places and transitions of N generated by Condition (1) of Definition 4.5, i.e. the smallest equivalence relation \equiv_D satisfying (1). The *canonical distribution* of N is the distribution C that maps each place or transition to its \equiv_C -equivalence class.

Observation 4.10. A Petri net that is distributed (resp. essentially or externally distributed) w.r.t. any distribution D , is distributed (resp. essentially or externally distributed) w.r.t. its canonical distribution.

This follows because whenever a co-location relation \equiv_D satisfies Condition (2) of Definition 4.5 (resp. Condition (2') or (2'') of Definition 4.8), then so does any smaller co-location relation. Hence a net is distributed (resp. essentially or externally distributed) iff its canonical distribution D satisfies (2) (resp. (2') or (2'')).

4.3. Correspondence between LSGA nets and distributed nets. We proceed to show that the classes of LSGA nets, distributable nets and essentially distributable nets essentially coincide.

That every LSGA net is distributed follows because we can place each sequential component on a separate location. The following two lemmas constitute a formal argument. Here we call a component with interface (N, I, O) distributed iff N is distributed.

Lemma 4.11. *Any sequential component with interface is distributed.*

Proof. As a sequential component displays no concurrency, it suffices to co-locate all places and transitions. \square

Lemma 4.12 states that the class of distributed nets is closed under asynchronous parallel composition.

Lemma 4.12. *Let $\mathcal{C}_k = (N_k, I_k, O_k)$, $k \in \mathfrak{K}$, be components with interface, satisfying the requirements of Definition 4.2, which are all distributed. Then $\parallel_{k \in \mathfrak{K}} \mathcal{C}_k$ is distributed.*

Proof. We need to find a distribution D satisfying the requirements of Definition 4.5.

Every component \mathcal{C}_k is distributed and hence comes with a distribution D_k . Without loss of generality the codomains of all D_k can be assumed disjoint.

Considering each D_k as a function from net elements onto locations, a partial function D'_k can be defined which does not map any places in O_k , denoting that the element may be located arbitrarily, and behaves as D_k for all other elements. As an output place has no posttransitions within a component, any total function larger than (i.e. a superset of) D'_k is still a valid distribution for N_k .

Now $D' = \bigcup_{k \in \bar{K}} D'_k$ is a (partial) function, as every place shared between components is an input place of at most one. The required distribution D can be chosen as any total function extending D' ; it satisfies the requirements of Definition 4.5 since the D_k 's do. \square

Corollary 4.13. *Every LSGA net is distributed.* \square

Corollary 4.14. *Every LSGA net is a structural conflict net.* \square

Conversely, any distributed net N , and even any essentially distributed net N , can be transformed in an LSGA net by choosing co-located transitions with their pre- and postplaces as sequential components and declaring any place that belongs to multiple components to be an input place of component N_k if it is a preplace of a transition in N_k , and an output place of component N_l if it is a postplace of a transition in N_l and not an input place of N_l . As transitions sharing a preplace are co-located, a place will be an input place of at most one component. Furthermore, in order to guarantee that the components are sequential in the sense of Definition 4.1, an explicit control place is added to each component—without changing behaviour—as explained below Definition 4.1. It is straightforward to check that the asynchronous parallel composition of all so-obtained components is an LSGA net, and that it is equivalent to N (using $\approx_{\mathcal{F}}$, \approx_{bSTb}^{Δ} , or any other reasonable equivalence).

Theorem 4.15. *For any essentially distributed net N there is an LSGA net N' with $N' \approx_{bSTb}^{\Delta} N$.*

Proof. Let $N = (S, T, F, M_0, \ell)$ be an essentially distributed net with a distribution D . Then an equivalent LSGA net N' can be constructed by composing sequential components with interfaces as follows.

For each equivalence class $[x]$ of net elements according to D a sequential component $(N_{[x]}, I_{[x]}, O_{[x]})$ is created. Each such component contains one new and initially marked place $p_{[x]}$ which is connected via self-loops to all transitions in $[x]$. The interface of the component is formed by $I_{[x]} := (S \cap [x])^5$ and $O_{[x]} := ([x] \cap T)^{\bullet} \setminus [x]$. Formally, $N_{[x]} := (S_{[x]}, T_{[x]}, F_{[x]}, M_{0[x]}, \ell_{[x]})$ with

- $S_{[x]} = ((S \cap [x]) \cup O_{[x]} \cup \{p_{[x]}\})$,
- $T_{[x]} = T \cap [x]$,
- $F_{[x]} = F \upharpoonright (S_{[x]} \cup T_{[x]})^2 \cup \{(p_{[x]}, t), (t, p_{[x]}) \mid t \in T_{[x]}\}$,
- $M_{0[x]} = (M_0 \upharpoonright [x]) \cup \{p_{[x]}\}$, and
- $\ell_{[x]} = \ell \upharpoonright [x]$.

All components overlap at interfaces only, as the sole places not in an interface are the newly created $p_{[x]}$. The $I_{[x]}$ are disjoint as the equivalence classes $[x]$ are, so $(N', I', O') := \bigsqcup_{[x] \in (S \cup T)/D} (N_{[x]}, O_{[x]}, I_{[x]})$ is well-defined. It remains to be shown that $N' \approx_{bSTb}^{\Delta} N$. The elements of N' are exactly those of N plus the new places $p_{[x]}$, which stay marked continuously except when a transition from $[x]$ is firing, and never connect two concurrently enabled transitions.

⁵Alternatively, we could take $I_{[x]} := (T \setminus [x])^{\bullet} \cap [x]$.




Figure 1: The LSGA net obtained from converting the essentially distributed net of Figure 4.

As we cannot have concurrently firing visible transitions on a single location, $|\overline{U} \cap [x]| \leq 1$ for any reachable ST-marking (M, U) of N and any $x \in S \cup T$, i.e. for any location $[x]$. Here \overline{U} is the multiset representation of the sequence U , defined in Section 3. The relation

$$\{((M, U), (M \cup P_U, U)) \mid (M, U) \text{ is a reachable ST-marking of } N, P_U = \{p_{[x]} \mid \overline{U} \cap [x] = \emptyset\}\}$$

is a bijection between the reachable ST-markings of N' and N that preserves the ST-transition relations between them. In particular, if $(M, U) \xrightarrow{\tau} (M', U')$, using a silent transition that belongs to the equivalence class $[x]$, then $U' = U$ and $\overline{U} \cap [x] = \emptyset$, i.e. no transition at location $[x]$ is currently firing, using that N is essentially distributed. Hence $p_{[x]} \in P_U$ and thus $(M \cup P_U, U) \xrightarrow{\tau} (M' \cup P_U, U)$. (This argument does not extend to externally distributed nets N .) From this it follows that $N' \approx_{bSTb}^{\Delta} N$. \square

Example 4.16. In Figure 4 appears an example of an essentially distributed net; the location borders are indicated. This net is not distributed, and thus not an LSGA net, because the two topmost τ -transitions are co-located but can be fired concurrently. Applying the construction in the proof of Theorem 4.15 turns this net into the distributed net of Figure 1.

Likewise, up to $\approx_{\mathcal{F}}$ any externally distributed net can be converted into a distributed net.

Proposition 4.17. [GGS08] *For any externally distributed net N there is a distributed net N' with $N' \approx_{\mathcal{F}} N$.*

Proof. The same construction applies. The relation

$$\{(M, M \cup P) \mid M \text{ is a reachable marking of } N, P = \{p_{[x]} \mid [x] \text{ is a location}\}\}$$

is a bijection between the reachable markings of N' and N that preserves the step transition relations between them. Here we use that the transitions in the associated LTS involve either a multiset of concurrently firing *visible* transitions (that all reside on different locations and thus do not share a preplace $p_{[x]}$), or a single internal one. It follows that $N' \approx_{\mathcal{F}} N$. \square

Example 4.18. Figure 2 shows an externally distributed net; the (canonical) location borders are dotted. It is not essentially distributed, because the transitions t and v are co-located but can be fired concurrently, while $\ell(t) \neq \tau$. Applying the construction in the proof of Proposition 4.17 turns this net into the step failures equivalent LSGA net of Figure 3.

The counterexample in Figure 2 shows that up to \approx_{bSTb}^{Δ} not all externally distributed nets can be converted into distributed nets. Sequentialising the component with actions a , b and τ (as happens in Figure 3) would disable the execution $\xrightarrow{a^+} \xRightarrow{\tau} \xrightarrow{c^+}$.




Figure 2: Externally distributed, but not convertible into a distributed net up to \approx_{bSTb}^Δ .




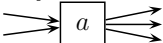

Figure 3: The LSGA net obtained from converting the externally distributed net of Figure 2.

5. DISTRIBUTABLE SYSTEMS

We now consider Petri nets as specifications of concurrent systems and ask the question which of those specifications can be implemented as distributed systems. This question can be formalised as

Which Petri nets are semantically equivalent to distributed nets?

Of course the answer depends on the choice of a suitable semantic equivalence. Here we will answer this question using the two equivalences discussed in the introduction. We will give a precise characterisation of those nets for which we can find semantically equivalent distributed nets. For the negative part of this characterisation, stating that certain nets are not distributable, we will use step failures equivalence, which is one of the simplest and least discriminating equivalences imaginable that abstracts from internal actions, but preserves branching time, concurrency and divergence to some small degree.⁶ Giving up on any of these latter three properties would make any Petri net distributable, but in a rather trivial and unsatisfactory way:

- Every net can be converted into an essentially distributed net by refining every transition  into the net segment . This construction appears in [BD12] where it is criticised for putting “all relevant choice resolutions” on one location. The construction does not introduce or remove concurrency or divergence. So it preserves even causality respecting linear time equivalences like pomset trace equivalence [GG01]. It does not preserve branching time equivalences, because a choice between two visible transitions a and b in the original net is implemented by a choice between two internal transitions preceding a and b . The resulting net is essentially distributed because all new τ -transitions can be placed on the same location, whereas all other transitions get allocated a location of their own. Hence, using Theorem 4.15, it can be converted into an equivalent distributed net.
- When working in interleaving semantics, any net can be converted into an equivalent distributed net by removing all concurrency between transitions. This can be accomplished

⁶In [GGS12] we used *step readiness equivalence*, a slightly more discriminating equivalence with roughly the same properties. By moving to step failures equivalence we strengthen our result.




Figure 4: A busy-wait implementation of the net in Figure 5, location borders dotted.

by adding a new, initially marked place, with an arc to and from every transition in the net.

- When fully abstracting from divergence, even when respecting causality and branching time, the net of Figure 5 is equivalent to the essentially distributed net of Figure 4, and in fact it is not hard to see that this type of implementation is possible for any given net. Yet, the implementation may diverge, as the nondeterministic choices might consistently be decided in an unhelpful way. This argument is elaborated in Section 5.1 below. The clause $M \xrightarrow{\tau}$ in Definition 3.6 is strong enough to rule out this type of implementation, even though our step failures semantics abstracts from other forms of divergence.

For the positive part, namely that all other nets are indeed distributable, we will use the most discriminating equivalence for which our implementation works, namely branching ST-bisimilarity with explicit divergence, which is finer than step failures equivalence. Hence we will obtain the strongest possible results for both directions and it turns out that the concept of distributability is fairly robust w.r.t. the choice of a suitable equivalence: any equivalence notion between step failures equivalence and branching ST-bisimilarity with explicit divergence will yield the same characterisation.

Definition 5.1. A Petri net N' is *distributable* up to an equivalence \approx iff there exists a distributed net N with $N \approx N'$.

Formally we give our characterisation of distributability by classifying which finitary plain structural conflict nets can be implemented as distributed nets, and hence as LSGA nets. In such implementations, we use invisible transitions. We study the concept “distributable” for plain nets only, but in order to get the largest class possible we allow non-plain implementations, where a given transition may be split into multiple transitions carrying the same label.

5.1. Characterising Distributability. It is well known that sometimes a global protocol is necessary to implement synchronous interactions present in system specifications. In particular, this may be needed for deciding choices in a coherent way, when these choices require agreement of multiple components. The simple net in Figure 5 shows a typical situation of this kind. Independent decisions of the two choices might lead to incorrect system behaviour. If p and q both decide to send their respective tokens leftwards, a can fire, yet the token from q gets stuck as b never receives a second token. Compared to the correct semantics, a firing of c after a is missing. It can be argued that for this particular net there exists no satisfactory distributed implementation that fully respects the reactive behaviour of the original system: Transitions t and v are supposed to be concurrently

executable (if we do not want to restrict performance of the system), and hence reside on different locations. Thus at least one of them, say t , cannot be co-located with transition u . However, both transitions are in conflict with u .

As we use nets as models of reactive systems, we allow the environment of a net to influence decisions at runtime by blocking some of the possibilities. Equivalently we can say it is the environment that fires transitions, and this can only happen for transitions that are currently enabled in the net. If the net decides between t and u before the actual execution of the chosen transition, the environment might change its mind in between, leading to a state of deadlock. Therefore we work in a branching time semantics, in which the option to perform t stays open until either t or u occurs. Hence the decision to fire u can only be taken at the location of u , namely by firing u , and similarly for t . Assuming that it takes time to propagate any message from one location to another, in no distributed implementation of this net can t and u be simultaneously enabled, because in that case we cannot exclude that both of them happen. Thus, the only possible implementation of the choice between t and u is to alternate the right to fire between t and u , by sending messages between them (cf. Figure 4). But if the environment only sporadically tries to fire t or u it may repeatedly miss the opportunity to do so, leading to an infinite loop of control messages sent back and forth, without either transition ever firing.




Figure 5: A fully reachable pure \mathbf{M} .

Indeed such \mathbf{M} -structures, representing interference between concurrency and choice, turn out to play a crucial rôle for characterising distributability. To be specific, it is only those \mathbf{M} s that are *pure*, i.e. don't have extra arcs from their places to their transitions besides those in Figure 5, and are *fully reachable*, i.e. for which there exists a reachable marking enabling all three transitions at the same time.

Definition 5.2. Let $N = (S, T, F, M_0, \ell)$ be a Petri net. N has a *fully reachable pure \mathbf{M}* iff $\exists t, u, v \in T. \bullet t \cap \bullet u \neq \emptyset \wedge \bullet u \cap \bullet v \neq \emptyset \wedge \bullet t \cap \bullet v = \emptyset \wedge \exists M \in [M_0]. \bullet t \cup \bullet u \cup \bullet v \leq M$.

Note that Definition 5.2 implies that $t \neq u$, $u \neq v$ and $t \neq v$.

Observation 5.3. A net with a fully reachable pure \mathbf{M} is not distributed. □

We now give an upper bound on the class of distributable nets by adapting a result from [GGS08]: We show that fully reachable pure \mathbf{M} 's that are present in a plain structural conflict net are preserved under step failures equivalence. In [GGS08] we showed this for step readiness equivalence.

Lemma 5.4. Let $N = (S, T, F, M_0, \ell)$ be a plain structural conflict net. If N has a fully reachable pure \mathbf{M} , then there are $\sigma \in \text{Act}^*$ and $a, b, c \in \text{Act}$ with $a \neq c$, such that $\langle \sigma, \{\{a, c\}\} \rangle, \langle \sigma, \{\{b\}\} \rangle \notin \mathcal{F}(N)$ and $\langle \sigma, \{\{a, b\}, \{b, c\}\} \rangle \in \mathcal{F}(N)$. (It is implied that $a \neq b \neq c$.)

Proof. N has a fully reachable pure \mathbf{M} , so there exist $t, u, v \in T$ and $M \in [M_0]$ such that $\bullet t \cap \bullet u \neq \emptyset \wedge \bullet u \cap \bullet v \neq \emptyset \wedge \bullet t \cap \bullet v = \emptyset \wedge \bullet t \cup \bullet u \cup \bullet v \leq M$. Let $\sigma \in \text{Act}^*$ such that

$M_0 \xrightarrow{\sigma} M$. Let $a := \ell(t)$, $b := \ell(u)$ and $c := \ell(v)$. Then $M \xrightarrow{\{a,c\}}$ and $M \xrightarrow{\{b\}}$. Moreover, using that N is a structural conflict net, $M \xrightarrow{\{a,b\}}$ and $M \xrightarrow{\{b,c\}}$. Since N is a plain net, $M \not\xrightarrow{\tau}$, and there is no $M' \neq M$ with $M_0 \xrightarrow{\sigma} M'$. Hence $\langle \sigma, \{\{a, c\}\} \rangle, \langle \sigma, \{\{b\}\} \rangle \notin \mathcal{F}(N)$ and $\langle \sigma, \{\{a, b\}, \{b, c\}\} \rangle \in \mathcal{F}(N)$. \square

Lemma 5.5. *Let $N = (S, T, F, M_0, \ell)$ be a structural conflict net. If there are $\sigma \in \text{Act}^*$ and $a, b, c \in \text{Act}$ with $a \neq c$, such that $\langle \sigma, \{\{a, c\}\} \rangle, \langle \sigma, \{\{b\}\} \rangle \notin \mathcal{F}(N)$ and $\langle \sigma, \{\{a, b\}, \{b, c\}\} \rangle \in \mathcal{F}(N)$, then N has a fully reachable pure \mathbf{M} .*

Proof. Let $M \in \mathbb{N}^S$ be the marking that gives rise to the step failure pair $\langle \sigma, \{\{a, b\}, \{b, c\}\} \rangle$, i.e. $M_0 \xrightarrow{\sigma} M$, $M \xrightarrow{\{a,b\}}$ and $M \xrightarrow{\{b,c\}}$. Since $\langle \sigma, \{a, c\} \rangle \notin \mathcal{F}(N)$, it must be that $M \xrightarrow{\{a,c\}}$. Likewise, $M \xrightarrow{\{b\}}$.

As $a \neq b \neq c \neq a$ there must exist three transitions $t, u, v \in T$ with $\ell(t) = a \wedge \ell(u) = b \wedge \ell(v) = c$ and $M[\{t, v\}] \wedge M[\{u\}] \wedge \neg(M[\{t, u\}]) \wedge \neg(M[\{u, v\}])$. From $M[\{t, v\}] \wedge M[\{u\}]$ it follows that $\bullet t \cup \bullet u \cup \bullet v \leq M$ and $\bullet t \cap \bullet v = \emptyset$, using that N is a structural conflict net. From $\neg(M[\{t, u\}])$ then follows $\bullet t \cap \bullet u \neq \emptyset$ and analogously for u and v . Hence N has a fully reachable pure \mathbf{M} . \square

Note that the lemmas above give a behavioural property that for plain structural conflict nets is equivalent to having a fully reachable pure \mathbf{M} .

Theorem 5.6. *Let N be a plain structural conflict Petri net. If N has a fully reachable pure \mathbf{M} , then N is not distributable up to step failures equivalence.*

Proof. Let N be a plain structural conflict net which has a fully reachable pure \mathbf{M} . Let N' be a net which is step failures equivalent to N . By Lemma 5.4 and Lemma 5.5, also N' has a fully reachable pure \mathbf{M} . By Observation 5.3, N' is not distributed. Thus N is not distributable up to step failures equivalence. \square

Since \approx_{bSTb}^{Δ} is finer than $\approx_{\mathcal{F}}$, this result holds also for distributability up to \approx_{bSTb}^{Δ} (and any equivalence between $\approx_{\mathcal{F}}$ and \approx_{bSTb}^{Δ}).

In the following, we establish that this upper bound is tight, and hence a finitary plain structural conflict net is distributable iff it has no fully reachable pure \mathbf{M} . For this, it is helpful to first introduce a more compact graphical notation for Petri nets as well as macros for reversibility of transitions.

5.2. A compressed Petri net notation. To compress the graphical notation, we allow universal quantifiers of the form $\forall x. \phi(x)$ to appear in the drawing (cf. Figures 6 and 7). A quantifier replaces occurrences of x in place and transition identities with all concrete values for which $\phi(x)$ holds, possibly creating a set of places, respectively transitions, instead of the depicted single one. Accordingly, an arc of which only one end is replicated by a given




Figure 6: A net with quantifiers.




Figure 7: The same net expanded.

quantifier results in a fan of arcs, one for each replicated element. If both ends of an arc are affected by the same quantifier, an arc is created between pairs of elements corresponding to the same x , but not between elements created due to differing values of x .

5.3. Petri nets with reversible transitions. A *Petri net with reversible transitions* generalises the notion of a Petri net; its semantics is given by a translation to an ordinary Petri net, thereby interpreting the reversible transitions as syntactic sugar for certain net fragments. It is defined as a tuple $(S, T, \Omega, \iota, F, M_0, \ell)$ with S a set of places, T a set of (reversible) transitions, labelled by $\ell : T \rightarrow \text{Act} \cup \{\tau\}$, Ω a set of *undo interfaces* with the relation $\iota \subseteq \Omega \times T$ linking interfaces to transitions, $M_0 \in \mathbb{N}^S$ an initial marking, and

$$F : (S \times T \times \{in, early, late, out, far\}) \rightarrow \mathbb{N}$$

the flow relation. When $F(s, t, type) > 0$ for $type \in \{in, early, late, out, far\}$, this is depicted by drawing an arc from s to t , labelled with its arc weight $F(s, t, type)$, of the form \longrightarrow , $\longrightarrow\leftarrow$, \longleftrightarrow , \longleftarrow , $\longleftarrow\leftarrow$, respectively. For $t \in T$ and $type \in \{in, early, late, out, far\}$, the multiset of places $t^{type} \in \mathbb{N}^S$ is given by $t^{type}(s) = F(s, t, type)$. When $s \in t^{type}$ for $type \in \{in, early, late\}$, the place s is called a *preplace* of t of type $type$; when $s \in t^{type}$ for $type \in \{out, far\}$, s is called a *postplace* of t of type $type$. For each undo interface $\omega \in \Omega$ and transition t with $\iota(\omega, t)$ there must be places $\text{undo}_\omega(t)$, $\text{reset}_\omega(t)$ and $\text{ack}_\omega(t)$ in S . A transition with a nonempty set of interfaces is called *reversible*; the other (*standard*) transitions may have pre- and postplaces of types *in* and *out* only—for these transitions $t^{in} = \bullet t$ and $t^{out} = t \bullet$. In case $\Omega = \emptyset$, the net is just a normal Petri net.

A global state of a Petri net with reversible transitions is given by a marking $M \in \mathbb{N}^S$, together with the state of each reversible transition “currently in progress”. Each transition in the net can fire as usual. A reversible transition can moreover take back (some of) its output tokens, and be *undone* and *reset*. (The use in our implementation will be that every reversible transition that fires is undone and reset later.) When a transition t fires, it consumes $\sum_{type \in \{in, early, late\}} F(s, t, type)$ tokens from each of its preplaces s and produces $\sum_{type \in \{out, far\}} F(s, t, type)$ tokens in each of its postplaces s . A reversible transition t that has fired can start its reversal by consuming a token from $\text{undo}_\omega(t)$ for one of its interfaces ω . Subsequently, it can take back the tokens from its postplaces of type *far*. After it has retrieved all its output of type *far*, the transition is undone, thereby returning $F(s, t, early)$ tokens in each of its preplaces s of type *early*. Afterwards, by consuming a token from $\text{reset}_\omega(t)$, for the same interface ω that started the undo-process, the transition terminates its chain of activities by returning $F(s, t, late)$ tokens in each of its *late* preplaces s . At that occasion it also produces a token in $\text{ack}_\omega(t)$. Alternatively, two tokens in $\text{undo}_\omega(t)$ and $\text{reset}_\omega(t)$ can annihilate each other without involving the transition t ; this also produces a token in $\text{ack}_\omega(t)$. The latter mechanism comes in action when trying to undo a transition that has not yet fired.

Figure 8 shows the translation of a reversible transition t with $\ell(t) = a$ into an ordinary net fragment. The arc weights on the green (or grey) arcs are inherited from the untranslated net; the other arcs have weight 1. Formally, a net $(S, T, \Omega, \iota, F, M_0, \ell)$ with reversible transitions translates into the Petri net containing all places S , all standard transitions in T , labelled according to ℓ , along with their pre- and postplaces, and furthermore all net elements mentioned in Table 1, T^{\leftarrow} denoting the set of reversible transitions in T . The initial marking is exactly M_0 .




Figure 8: A reversible transition and its macro expansion.

Transition	at label	Preplaces	Postplaces	for all	
$t \cdot \text{fire}$	t	$\ell(t)$	$t^{in}, t^{early}, t^{late}$	$\text{fired}(t), t^{out}, t^{far}$	$t \in T^{\leftarrow}$
$t \cdot \text{undo}_\omega$	$t\text{-undo}$	τ	$\text{undo}_\omega(t), \text{fired}(t)$	$\rho_\omega(t), \text{take}(f, t)$	$t \in T^{\leftarrow}, \iota(\omega, t), f \in t^{far}$
$t \cdot \text{undo}(f)$	f	τ	$\text{take}(f, t), f$	$\text{took}(f, t)$	$t \in T^{\leftarrow}, f \in t^{far}$
$t \cdot \text{undone}$	$t\text{-undo}$	τ	$\text{took}(f, t)$	$\rho(t), t^{early}$	$t \in T^{\leftarrow}, f \in t^{far}$
$t \cdot \text{reset}_\omega$	$t\text{-undo}$	τ	$\text{reset}_\omega(t), \rho_\omega(t), \rho(t)$	$t^{late}, \text{ack}_\omega(t)$	$t \in T^{\leftarrow}, \iota(\omega, t)$
$t \cdot \text{elide}_\omega$	$t\text{-undo}$	τ	$\text{undo}_\omega(t), \text{reset}_\omega(t)$	$\text{ack}_\omega(t)$	$t \in T^{\leftarrow}, \iota(\omega, t)$

Table 1: Expansion of a Petri net with reversible transitions into a place/transition system.

A distribution of a Petri net with reversible transitions can be given as a function $D : S \cup T \rightarrow \text{Loc}$. As in Condition (1) of Definition 4.5 we require that a transition and its preplaces (of types *in*, *early* or *late*) reside on the same location. Additionally, for any given transition t , all its undo-interface places $\text{undo}_\omega(t)$ and $\text{reset}_\omega(t)$ for all $\omega \in \Omega$ must reside on the same location—we refer to this location as $t\text{-undo}$. The second column of Table 1 indicates how such a distribution is translated under expansion of reversible transitions into ordinary net fragments: The location of a reversible transition t is really the location of $t \cdot \text{fire}$; it should be the same as all preplaces of t . Furthermore, the transition $t \cdot \text{undo}(f)$ and its preplace $\text{take}(f, t)$ reside on the same location as the place $f \in t^{far}$. All other net elements that are part of the macro expansion of t , except for $\text{ack}_\omega(t)$, reside at the location $t\text{-undo}$. The resulting distribution of the expanded net is now guaranteed to satisfy (1). Whether a

Petri net with reversible translations is (essentially) distributed requires checking Condition (2) of Definition 4.5 (or Condition (2') of Definition 4.8) on its expansion.

5.4. The conflict replicating implementation. Now we establish that a finitary plain structural conflict net that has no fully reachable pure \mathbf{M} is distributable. We do this by proposing the *conflict replicating implementation* of any such net, and show that this implementation is always (a) essentially distributed, and (b) equivalent to the original net. In order to get the strongest possible result, for (b) we use branching ST-bisimilarity with explicit divergence.

To define the conflict replicating implementation of a net $N' = (S', T', F', M'_0, \ell')$ we fix an arbitrary well-ordering $<$ on its transitions. We let $b, c, g, h, i, j, k, l, u$ range over these ordered transitions, and write

- $i \# j$ iff $i \neq j \wedge \bullet i \cap \bullet j \neq \emptyset$ (transitions i and j are *in conflict*), and $i \stackrel{\#}{=} j$ iff $i \# j \vee i = j$,
- $i <^{\#} j$ iff $i < j \wedge i \# j$, and $i \leq^{\#} j$ iff $i <^{\#} j \vee i = j$.

Figure 9 shows the conflict replicating implementation of N' . It is presented as a Petri net

$$\mathcal{I}(N') = (S, T, F, \Omega, \iota, M_0, \ell)$$

with reversible transitions. The set Ω of undo interfaces is T' , and for $i \in \Omega$ we have $\iota(i, t)$ iff $t \in \Omega_i$, where the sets of transitions $\Omega_i \subseteq T$ are specified in Figure 9. The implementation $\mathcal{I}(N')$ inherits the places of N' (i.e. $S \supseteq S'$), and we define $M_0 \upharpoonright S'$ to be M'_0 . Given this, Figure 9 is not merely an illustration of $\mathcal{I}(N')$ —it provides a complete and accurate description of it, thereby defining the conflict replicating implementation of any net. In interpreting this figure it is important to realise that net elements are completely determined by their name (identity), and exist only once, even if they show up multiple times in the figure. For instance, the place $\pi_{h\#j}$ with $h=2$ and $j=5$ (when using natural numbers for the transitions in T') is the same as the place $\pi_{j\#l}$ with $j=2$ and $l=5$; it is a standard preplace of execute_2^i (for all $i \leq^{\#} 2$), a standard postplace of fetched_2^i , as well as a late preplace of transfer_5^j . Figure 10 depicts the same net after expanding the macros for reversible transitions. An alternative description of the latter net appears in Table 2 on Page 39.

The rôle of the transitions distribute_p for $p \in S'$ is to distribute a token in p to copies p_j of p in the localities of all transitions $j \in T'$ with $p \in \bullet j$. In case j is enabled in N' , the transition initialise_j will become enabled in $\mathcal{I}(N')$. These transitions put tokens in the places pre_k^j , which are preconditions for all transitions execute_k^j , which model the execution of j at the location of k . When two conflicting transitions h and j are both enabled in N' , the first steps initialise_h and initialise_j towards their execution in $\mathcal{I}(N')$ can happen in parallel. To prevent them from executing both, execute_j^j (of j at its own location) is only possible after transfer_j^h , which disables execute_h^h . This happens because transfer_j^h takes the initially present token from the place $\pi_{h\#j}$, which is needed to fire execute_h^h .

The main idea behind the conflict replicating implementation is that a transition $h \in T'$ is primarily executed by a sequential component of its own, but when a conflicting transition j gets enabled, the sequential component implementing j may “steal” the possibility to execute h from the home component of h , by putting a token in trans_j^h -in and getting transfer_j^h to fire, and then keep the options to do h and j open on the home component of j until one of them occurs. To prevent h and j from stealing each other’s initiative, which would result in deadlock, a global asymmetry is built in by ordering the transitions. Transition j can steal the initiative from h only when $h < j$.




Figure 9: The entire conflict replicating implementation, drawn with emphasis on the structure of the component of j ; location borders dotted.

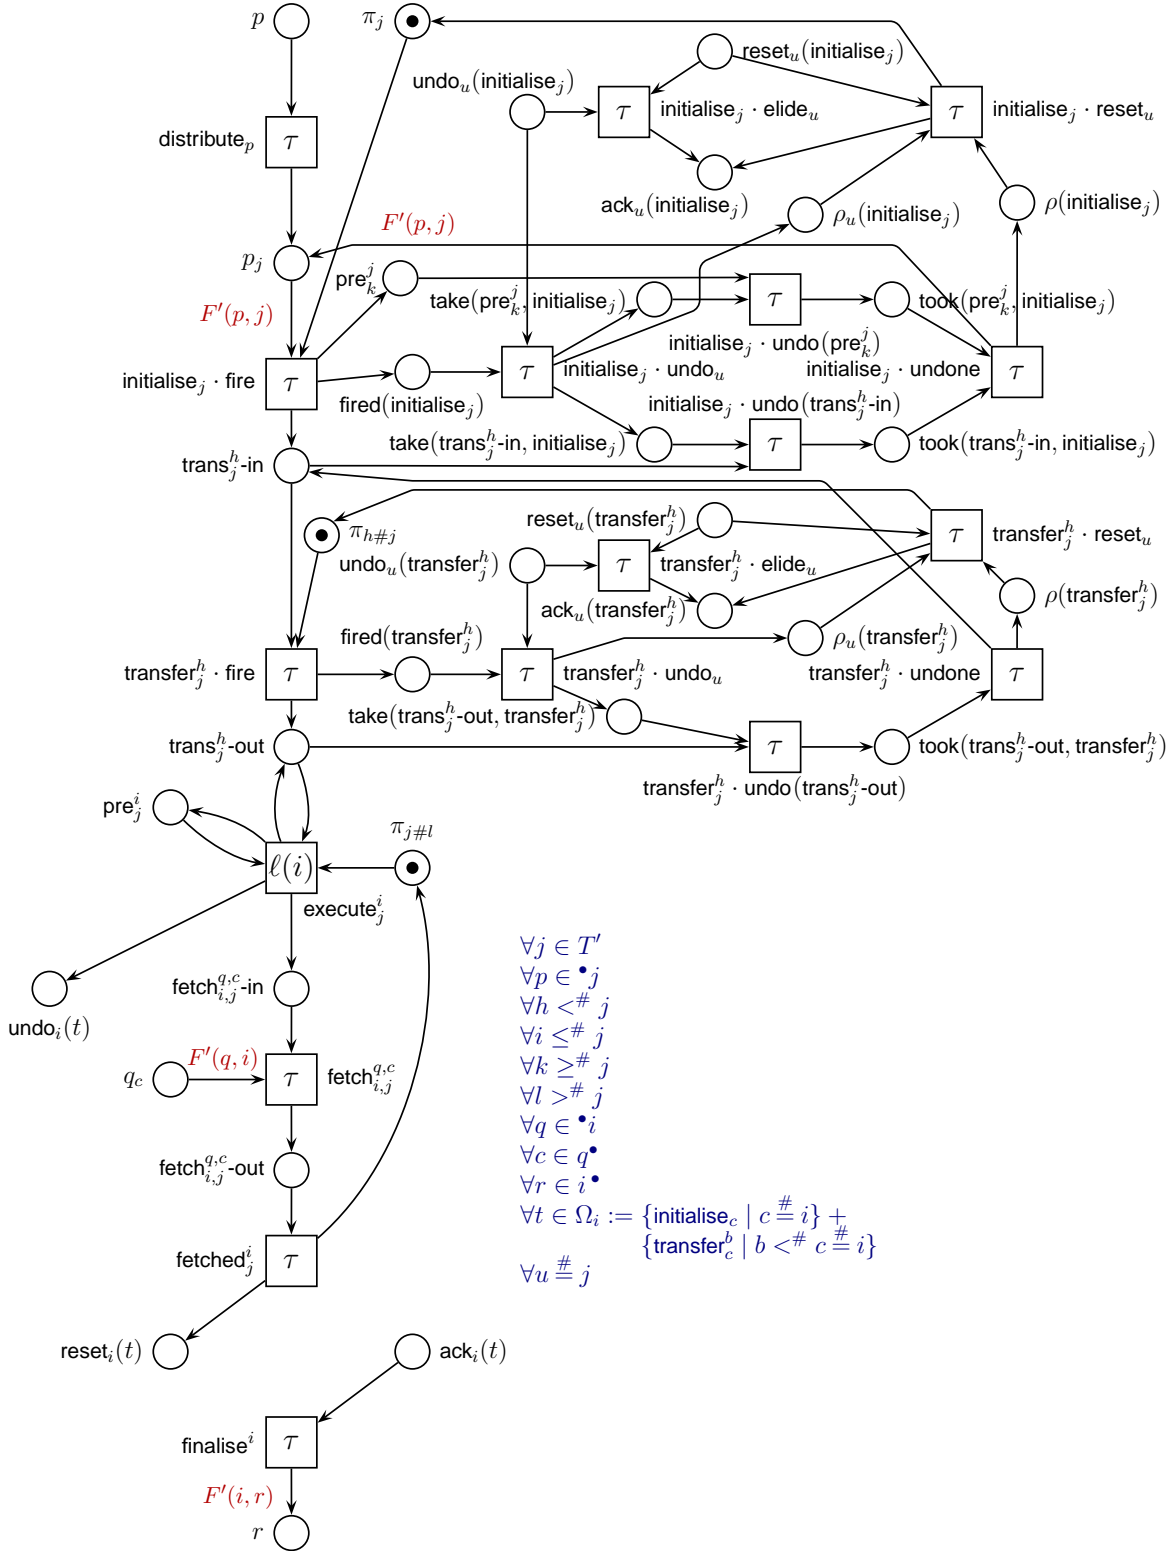


Figure 10: The entire conflict replicating implementation (with macros expanded).

In case j is also in conflict with a transition l , with $j < l$, the initiative to perform j may subsequently be stolen by l . In that case either h and l are in conflict too—then l takes responsibility for the execution of h as well—or h and l are concurrent—in that case h will not be enabled, due to the absence of fully reachable pure \mathbf{M} s in N' . The absence of fully reachable pure \mathbf{M} s also guarantees that it cannot happen that two concurrent transitions j and k both steal the initiative from an enabled transition h .

After the firing of execute_j^i , all tokens that were left behind in the process of carefully orchestrating this firing will have to be cleaned up, in order to prepare the net for the next activity in the same neighbourhood. This is the reason for the reversibility of the transitions preparing the firing of execute_j^i . Hence there is an undo interface for each transition $i \in T'$, cleaning up the mess made in preparation of firing execute_j^i for some $j \geq^\# i$. Ω_i is the set of all transitions t that could possibly have contributed to this. For each of them the undo interface i is activated, by execute_j^i depositing a token in $\text{undo}_i(t)$. After all preparatory transitions that have fired are undone, tokens appear in the places p_c for all $p \in \bullet i$ and $c \in p^\bullet$. These are collected by $\text{fetch}_{i,j}^{p,c}$, after which all transitions in Ω_i get a reset signal. Those that have fired and were undone are reset, and those that never fired perform $\text{elide}_i(t)$. In either case a token appears in $\text{ack}_i(t)$. These are collected by finalise^i , which finishes the process of executing i by depositing tokens in its postplaces.

We allow multiple tokens to reside on the same place in the specification. To ensure that this does never lead to the component implementing a transition j starting the firing protocol again, even though it has not yet completed an earlier round, we introduce a place π_j which only holds a token while the component is idle.

By means of location boundaries, Figure 9 also displays a distribution of $\mathcal{I}(N')$. It has

- a location p for every place $p \in S'$, containing distribute_p and p ;
- locations initialise_j and execute_j for every $j \in T'$ —collectively referred to as “the location of j ”—the latter containing all transitions execute_j^i for $i \leq^\# j \in T'$;
- locations fetched_j^i for every $i \leq^\# j \in T'$;
- locations $\text{initialise}_j\text{-undo}$ for every $j \in T'$;
- locations $\text{transfer}_j^h\text{-undo}$ for every $h <^\# j \in T'$;
- and locations finalise^i for every $i \in T'$.

A transition transfer_j^h resides at location execute_h , due to its common preplace $\pi_{h\#j}$ with execute_h^g . Likewise, $\text{fetch}_{i,j}^{p,c}$ resides at location initialise_c . Provided N' is a finitary plain structural conflict net without a fully reachable pure \mathbf{M} , the proof of Theorem 7.11 will show that this distribution makes $\mathcal{I}(N')$ an essentially distributed net.

The conflict replicating implementation is illustrated by means of the finitary plain structural conflict net N' of Figure 11. The places and transitions a - q - b - s - c - x - d in this net constitute a *Long M*: for each pair a - b , b - c and c - d of neighbouring transitions, as well as for the pair a - d of extremal transitions, there exists a reachable marking enabling them both. Moreover, neighbouring transitions in the long \mathbf{M} are in conflict: $a \# b$, $b \# c$ and $c \# d$, whereas the extremal transitions are concurrent: $a \smile d$. However, N' has no fully reachable pure \mathbf{M} : no \mathbf{M} -shaped triple of transitions a - b - c , b - c - d or b - c - e is ever simultaneously enabled.

In [GGS08] we gave a simpler implementation, the *transition-controlled choice implementation*, that works for all finitary plain 1-safe Petri nets without such a long \mathbf{M} . Hence N' constitutes an example where that implementation does not apply, yet the conflict replicating implementation does. In fact, when leaving out the z - e -branch it may be the simplest




Figure 11: An example net.

example with these properties. We have added this branch to illustrate the situation where three transitions are pairwise in conflict.

Figure 12 presents relevant parts of the conflict replicating implementation $\mathcal{I}(N')$ of N' . What corresponds to the ten places of N' can easily be discerned in $\mathcal{I}(N')$, but the transitions of N' are replaced by more complicated net fragments. In Figure 12 we have simplified the rendering of $\mathcal{I}(N')$ by simply just copying the five topmost transitions of N' , instead of displaying the net fragments replacing them. This simplification is possible since the top half of N' is already distributed. To remind the reader of this, we left those transitions unlabelled.⁷

In order to fix a well-ordering $<$ on the remaining transitions, we named them after the first five positive natural numbers. The ordered conflicts between those transitions now are $1 \leq \#2$, $2 \leq \#3$, $3 \leq \#4$, $3 \leq \#5$ and $4 \leq \#5$. In Figure 12 we have skipped all places, transitions and arcs involved in the cleanup of tokens after firing of a transition. In this example the cleanup is not necessary, as no place of N' is visited twice. Thus, we displayed only the non-reversible part of the transitions initialise_j and transfer_j^h —i.e. $\text{initialise}_j \cdot \text{fire}$ and $\text{transfer}_j^h \cdot \text{fire}$ —as well as the transitions distribute_p and execute_j^i . Likewise, we omitted the outgoing arcs of execute_j^i , the places π_j , and those places that have arcs only to omitted transitions. We leave it to the reader to check this net against the definition in Figure 9, and to play the token game on this net, to see that it correctly implements N' .

In Section 7 we will show, for any finitary plain structural conflict net N' without a fully reachable pure \mathbf{M} , that $\mathcal{I}(N') \approx_{bSTb}^\Delta N'$, and that $\mathcal{I}(N')$ is essentially distributed. Hence $\mathcal{I}(N')$ is an essentially distributed implementation of N' . By Theorem 4.15 this implies that N' is distributable up to \approx_{bSTb}^Δ . Together with Theorem 5.6 it follows that, for any equivalence between $\approx_{\mathcal{F}}$ and \approx_{bSTb}^Δ , a finitary plain structural conflict net is distributable iff it has no fully reachable pure \mathbf{M} .

Given the complexity of our construction, no techniques known to us were adequate for performing the equivalence proof. We therefore had to develop an entirely new method for rigorously proving the equivalence of two Petri nets up to \approx_{bSTb}^Δ , one of which known to be plain. This method is presented in Section 6.

⁷While it is highly desirable in practical applications to use such simplifications to reduce the implementation size, we refrained from doing so in the formal definition of our implementation. It would have become less regular and the proofs correspondingly longer.




Figure 12: The (relevant parts of the) conflict replicating implementation of the net in Figure 11, location borders dotted.

6. PROVING IMPLEMENTATIONS CORRECT

This section presents a method for establishing the equivalence of two Petri nets, one of which known to be plain, up to branching ST-bisimilarity with explicit divergence. It appears as Theorem 6.8. First approximations of this method are presented in Lemmas 6.3 and 6.4. The progression from Lemma 6.3 to Lemma 6.4 and to Theorem 6.8 makes the method more specific (so less general) and more powerful. By means of a simplification a similar method can be obtained, also in three steps, for establishing the equivalence of two Petri nets up to interleaving branching bisimilarity with explicit divergence. This is elaborated at the end of this section.

We sometimes illustrate the results of this section in terms of the conflict replicating implementation of a net defined in Section 5.4. However, the actual application of these results to show the correctness of that implementation is presented in Section 7.

Definition 6.1. A labelled transition system $(\mathfrak{S}, \mathfrak{T}, \mathfrak{M}_o)$ is called *deterministic* if for all reachable states $\mathfrak{M} \in [\mathfrak{M}_o)$ we have $\mathfrak{M} \not\rightarrow^\tau$ and if $\mathfrak{M} \xrightarrow{a} \mathfrak{M}'$ and $\mathfrak{M} \xrightarrow{a} \mathfrak{M}''$ for some $a \in \mathfrak{Act}$ then $\mathfrak{M}' = \mathfrak{M}''$.

Deterministic systems may not have reachable τ -transitions at all; this way, if $\mathfrak{M} \xRightarrow{\sigma} \mathfrak{M}'$ and $\mathfrak{M} \xRightarrow{\sigma} \mathfrak{M}''$ for some $\sigma \in \mathfrak{Act}^*$ then $\mathfrak{M}' = \mathfrak{M}''$. Note that the labelled transition system associated to a plain Petri net is deterministic; the same applies to the ST-LTS, the split LTS or the step LTS associated to such a net.

Lemma 6.2. Let $(\mathfrak{S}_1, \mathfrak{T}_1, \mathfrak{M}_{o1})$ and $(\mathfrak{S}_2, \mathfrak{T}_2, \mathfrak{M}_{o2})$ be two labelled transition systems, the latter being deterministic. Suppose there is a relation $\mathcal{B} \subseteq \mathfrak{S}_1 \times \mathfrak{S}_2$ such that

- (a) $\mathfrak{M}_{o1} \mathcal{B} \mathfrak{M}_{o2}$,
- (b) if $\mathfrak{M}_1 \mathcal{B} \mathfrak{M}_2$ and $\mathfrak{M}_1 \xrightarrow{\tau} \mathfrak{M}'_1$ then $\mathfrak{M}'_1 \mathcal{B} \mathfrak{M}_2$,
- (c) if $\mathfrak{M}_1 \mathcal{B} \mathfrak{M}_2$ and $\mathfrak{M}_1 \xrightarrow{a} \mathfrak{M}'_1$ for some $a \in \mathfrak{Act}$ then $\exists \mathfrak{M}'_2. \mathfrak{M}_2 \xrightarrow{a} \mathfrak{M}'_2 \wedge \mathfrak{M}'_1 \mathcal{B} \mathfrak{M}'_2$,
- (d) if $\mathfrak{M}_1 \mathcal{B} \mathfrak{M}_2$ and $\mathfrak{M}_2 \xrightarrow{a}$ for some $a \in \mathfrak{Act}$ then either $\mathfrak{M}_1 \xrightarrow{a}$ or $\mathfrak{M}_1 \xrightarrow{\tau}$
- (e) and there is no infinite sequence $\mathfrak{M}_1 \xrightarrow{\tau} \mathfrak{M}'_1 \xrightarrow{\tau} \mathfrak{M}''_1 \xrightarrow{\tau} \dots$ with $\mathfrak{M}_1 \mathcal{B} \mathfrak{M}_2$ for some \mathfrak{M}_2 .

Then \mathcal{B} is a branching bisimulation with explicit divergence, and the two LTSs are branching bisimilar with explicit divergence.

Proof: It suffices to show that \mathcal{B} satisfies Conditions 1–3 of Definition 3.2; the condition on explicit divergence follows immediately from (e), using that a deterministic LTS admits no divergence at all.

- (1) By (a).
- (2) In case $\alpha = \tau$ this follows directly from (b), and otherwise from (c). In both cases $\mathfrak{M}_2^\dagger := \mathfrak{M}_2$ and when $\alpha = \tau$ also $\mathfrak{M}'_2 := \mathfrak{M}_2$.
- (3) Suppose $\mathfrak{M}_1 \mathcal{B} \mathfrak{M}_2$ and $\mathfrak{M}_2 \xrightarrow{\alpha} \mathfrak{M}'_2$. Since $(\mathfrak{S}_2, \mathfrak{T}_2, \mathfrak{M}_{o2})$ is deterministic, $\alpha = a \in \mathfrak{Act}$. By (d) we have either $\mathfrak{M}_1 \xrightarrow{a} \mathfrak{M}_1^1$ or $\mathfrak{M}_1 \xrightarrow{\tau} \mathfrak{M}_1^1$ for some $\mathfrak{M}_1^1 \in \mathfrak{S}_1$. In the latter case (b) yields $\mathfrak{M}_1^1 \mathcal{B} \mathfrak{M}_2$, and using (d) again, either $\mathfrak{M}_1^1 \xrightarrow{a} \mathfrak{M}_1^2$ or $\mathfrak{M}_1^1 \xrightarrow{\tau} \mathfrak{M}_1^2$ for some $\mathfrak{M}_1^2 \in \mathfrak{S}_1$. Repeating this argument, if the choice between a and τ is made k times in favour of τ (with $k \geq 0$), we obtain $\mathfrak{M}_1^k \mathcal{B} \mathfrak{M}_2$ (where $\mathfrak{M}_1^0 := \mathfrak{M}_1$) and either $\mathfrak{M}_1^k \xrightarrow{a} \mathfrak{M}_1^{k+1}$ or $\mathfrak{M}_1^k \xrightarrow{\tau} \mathfrak{M}_1^{k+1}$. By (e), at some point the choice must be made in favour of a , say at \mathfrak{M}_1^k . Thus $\mathfrak{M}_1 \xRightarrow{\alpha} \mathfrak{M}_1^k \xrightarrow{a} \mathfrak{M}_1^{k+1}$, with $\mathfrak{M}_1^k \mathcal{B} \mathfrak{M}_2$. We take \mathfrak{M}_1^\dagger and \mathfrak{M}'_1 from Definition 3.2 to be \mathfrak{M}_1^k and \mathfrak{M}_1^{k+1} . It remains to show that $\mathfrak{M}_1^{k+1} \mathcal{B} \mathfrak{M}'_2$.

By (c) there is an $\mathfrak{M}_2'' \in \mathfrak{S}_2$ with $\mathfrak{M}_2 \xrightarrow{a} \mathfrak{M}_2''$ and $\mathfrak{M}_1^{k+1} \mathcal{B} \mathfrak{M}_2''$. Since $(\mathfrak{S}_2, \mathfrak{T}_2, \mathfrak{M}_{o2})$ is deterministic, $\mathfrak{M}_2'' = \mathfrak{M}_2'$. \square

Lemma 6.3. *Let $N = (S, T, F, M_0, \ell)$ and $N' = (S', T', F', M'_0, \ell')$ be two nets, N' being plain. Suppose there is a relation $\mathcal{B} \subseteq (\mathbb{N}^S \times \mathbb{N}^T) \times (\mathbb{N}^{S'} \times \mathbb{N}^{T'})$ such that*

- (a) $(M_0, \emptyset) \mathcal{B} (M'_0, \emptyset)$,
- (b) if $(M_1, U_1) \mathcal{B} (M'_1, U'_1)$ and $(M_1, U_1) \xrightarrow{\tau} (M_2, U_2)$ then $(M_2, U_2) \mathcal{B} (M'_1, U'_1)$,
- (c) if $(M_1, U_1) \mathcal{B} (M'_1, U'_1)$ and $(M_1, U_1) \xrightarrow{\eta} (M_2, U_2)$ for some $\eta \in \text{Act}^\pm$ then $\exists (M'_2, U'_2)$. $(M'_1, U'_1) \xrightarrow{\eta} (M'_2, U'_2) \wedge (M_2, U_2) \mathcal{B} (M'_2, U'_2)$,
- (d) if $(M_1, U_1) \mathcal{B} (M'_1, U'_1)$ and $(M'_1, U'_1) \xrightarrow{\eta}$ with $\eta \in \text{Act}^\pm$ then either $(M_1, U_1) \xrightarrow{\eta}$ or $(M_1, U_1) \xrightarrow{\tau}$
- (e) and there is no infinite sequence $(M, U) \xrightarrow{\tau} (M_1, U_1) \xrightarrow{\tau} (M_2, U_2) \xrightarrow{\tau} \dots$ with $(M, U) \mathcal{B} (M', U')$ for some (M', U') .

Then \mathcal{B} is a branching split bisimulation with explicit divergence, and $N \approx_{bSTb}^\Delta N'$.

Proof. That N and N' are branching split bisimilar with explicit divergence follows directly from Lemma 6.2 by taking $(\mathfrak{S}_1, \mathfrak{T}_1, \mathfrak{M}_{o1})$ and $(\mathfrak{S}_2, \mathfrak{T}_2, \mathfrak{M}_{o2})$ to be the split LTSs associated to N and N' respectively. Here we use that the split LTS associated to a plain net is deterministic. The final conclusion follows by Proposition 3.15. \square

Lemma 6.3 provides a method for proving $N \approx_{bSTb}^\Delta N'$ that can be more efficient than directly checking the definition. In particular, the intermediate states \mathfrak{M}^\dagger and the sequence of τ -transitions \implies from Definition 3.2 do not occur in Lemma 6.2, and hence not in Lemma 6.3. Moreover, in Condition (d) one no longer has to match the targets of corresponding transitions. Lemma 6.4 below, when applicable, provides an even more efficient method: it is no longer necessary to specify the branching split bisimulation \mathcal{B} , and the targets have disappeared from the transitions in Condition 2c as well. Instead, we have acquired Condition 1, but this is a structural property, which is relatively easy to check.

Lemma 6.4. *Let $N = (S, T, F, M_0, \ell)$ be a net and $N' = (S', T', F', M'_0, \ell')$ be a plain net with $S' \subseteq S$ and $M'_0 = M_0 \upharpoonright S'$. Suppose:*

- (1) $\forall t \in T, \ell(t) \neq \tau. \exists t' \in T', \ell(t') = \ell(t). \exists G \in_F \mathbb{N}^T, \ell(G) \equiv \emptyset. \llbracket t' \rrbracket = \llbracket t + G \rrbracket$.
- (2) For any $G \in_F \mathbb{Z}^T$ with $\ell(G) \equiv \emptyset, M' \in \mathbb{N}^{S'}, U' \in \mathbb{N}^{T'}$ and $U \in \mathbb{N}^T$ with $\ell'(U') = \ell(U), M' + \bullet U' \in [M'_0]_{N'}$ and $M := M' + \bullet U' + (M_0 - M'_0) + \llbracket G \rrbracket - \bullet U \in \mathbb{N}^S$ with $M + \bullet U \in [M_0]_N$, it holds that:
 - (a) there is no infinite sequence $M \xrightarrow{\tau} M_1 \xrightarrow{\tau} M_2 \xrightarrow{\tau} \dots$
 - (b) if $M' \xrightarrow{a}$ with $a \in \text{Act}$ then $M \xrightarrow{a}$ or $M \xrightarrow{\tau}$
 - (c) and if $M \xrightarrow{a}$ with $a \in \text{Act}$ then $M' \xrightarrow{a}$.

Then $N \approx_{bSTb}^\Delta N'$.

Proof:⁸ Define $\mathcal{B} \subseteq (\mathbb{N}^S \times \mathbb{N}^T) \times (\mathbb{N}^{S'} \times \mathbb{N}^{T'})$ by $(M, U) \mathcal{B} (M', U') :\Leftrightarrow \ell'(U') = \ell(U) \wedge M' + \bullet U' \in [M'_0]_{N'} \wedge \exists G \in_F \mathbb{Z}^T. \ell(G) \equiv \emptyset \wedge M + \bullet U = M' + \bullet U' + (M_0 - M'_0) + \llbracket G \rrbracket \in [M_0]_N$. It suffices to show that \mathcal{B} satisfies Conditions (a)–(e) of Lemma 6.3.

- (a) Take $G = \emptyset$.
- (b) Suppose $(M_1, U_1) \mathcal{B} (M'_1, U'_1)$ and $(M_1, U_1) \xrightarrow{\tau} (M_2, U_2)$. Then $\ell'(U'_1) = \ell(U_1) \wedge M'_1 + \bullet U'_1 \in [M'_0]_{N'} \wedge \exists G \in_F \mathbb{Z}^T. \ell(G) \equiv \emptyset \wedge M_1 = M'_1 + \bullet U'_1 + (M_0 - M'_0) + \llbracket G \rrbracket - \bullet U_1 \wedge M_1 + \bullet U \in [M_0]_N$ and moreover $M_1 \xrightarrow{\tau} M_2 \wedge U_2 = U_1$. So $M_1[t]M_2$ for some $t \in T$ with $\ell(t) = \tau$. Hence

⁸For didactic reason it may be preferable to skip ahead and read the (simpler) proof of Lemma 6.10 first.

$M_2 = M_1 + \llbracket t \rrbracket = M'_1 + \bullet U'_1 + (M_0 - M'_0) + \llbracket G + t \rrbracket - \bullet U_1$. Since $(M_1 + \bullet U_1)[t](M_2 + \bullet U_1)$, we have $M_2 + \bullet U_1 \in [M_0]_N$. Since also $\ell(G + t) \equiv \emptyset$ it follows that $(M_2, U_1)\mathcal{B}(M'_1, U'_1)$.

(c) Suppose $(M_1, U_1)\mathcal{B}(M'_1, U'_1)$ and $(M_1, U_1) \xrightarrow{\eta} (M_2, U_2)$, with $\eta \in \text{Act}^\pm$. Then $\ell'(U'_1) = \ell(U_1)$, $M'_1 + \bullet U'_1 \in [M'_0]_{N'}$ and

$$\exists G \in_F \mathbb{Z}^T. \ell(G) \equiv \emptyset \wedge M_1 + \bullet U_1 = M'_1 + \bullet U'_1 + (M_0 - M'_0) + \llbracket G \rrbracket \in [M_0]_N. \quad (6.1)$$

First suppose $\eta = a^+$. Then $\exists t \in T$. $\ell(t) = a \wedge M_1[t] \wedge M_2 = M_1 - \bullet t \wedge U_2 = U_1 + \{t\}$. Using that $M_1 \xrightarrow{a}$ with $a \in \text{Act}$, by Condition 2c we have $M'_1 \xrightarrow{a}$, i.e. $M'_1[t']$ for some $t' \in T$ with $\ell'(t') = a$. Let $M'_2 := M'_1 - \bullet t$ and $U'_2 := U'_1 + \{t'\}$. Then $(M'_1, U'_1) \xrightarrow{a^+} (M'_2, U'_2)$. Moreover, $\ell(U_2) = \ell(U'_2)$, $M'_2 + \bullet U'_2 = M'_1 + \bullet U'_1 \in [M'_0]_{N'}$ and $M_2 + \bullet U_2 = M_1 + \bullet U_1$. In combination with (6.1) this yields

$$M_2 + \bullet U_2 = M_1 + \bullet U_1 = M'_1 + \bullet U'_1 + (M_0 - M'_0) + \llbracket G \rrbracket = M'_2 + \bullet U'_2 + (M_0 - M'_0) + \llbracket G \rrbracket,$$

so $(M_2, U_2)\mathcal{B}(M'_2, U'_2)$.

Now suppose $\eta = a^-$. Then $\exists t \in U_1$. $\ell(t) = a \wedge U_2 = U_1 - \{t\} \wedge M_2 = M_1 + t^\bullet$. Since $\ell'(U'_1) = \ell(U_1)$ there is a $t' \in U'_1$ with $\ell'(t') = a$. Let $M'_2 := M'_1 + t'^\bullet$ and $U'_2 := U'_1 - \{t'\}$. Then $(M'_1, U'_1) \xrightarrow{a^-} (M'_2, U'_2)$. By construction, $\ell(U_2) = \ell(U'_2)$. Moreover, $M_2 + \bullet U_2 = M_1 + t^\bullet + \bullet U_1 - \bullet t = (M_1 + \bullet U_1) + \llbracket t \rrbracket$, and likewise

$$M'_2 + \bullet U'_2 = (M'_1 + \bullet U'_1) + \llbracket t' \rrbracket \quad (6.2)$$

so $(M'_1 + \bullet U'_1)[t'](M'_2 + \bullet U'_2)$. Since $M'_1 + \bullet U'_1 \in [M'_0]_{N'}$, this yields $M'_2 + \bullet U'_2 \in [M'_0]_{N'}$. Moreover, $M_2 + \bullet U_2 = M_1 + t^\bullet + \bullet U_1 - \bullet t = M_1 + \bullet U_1 + \llbracket t \rrbracket \in [M_0]_N$. Furthermore, combining (6.1) and (6.2) gives

$$\exists G \in_F \mathbb{Z}^T. \ell(G) \equiv \emptyset \wedge M_2 + \bullet U_2 - \llbracket t \rrbracket = M'_2 + \bullet U'_2 - \llbracket t' \rrbracket + (M_0 - M'_0) + \llbracket G \rrbracket. \quad (6.3)$$

By Condition 1 of Lemma 6.4, $\exists t'' \in T'$, $\ell(t'') = \ell(t)$. $\exists G_t \in_F \mathbb{N}^T$, $\ell(G_t) \equiv \emptyset$. $\llbracket t \rrbracket = \llbracket t'' - G_t \rrbracket$. Since N' is a plain net, it has only one transition t^\dagger with $\ell(t^\dagger) = a$, so $t'' = t'$. Substitution of $\llbracket t' - G_t \rrbracket$ for $\llbracket t \rrbracket$ in (6.3) yields

$$\exists G \in_F \mathbb{Z}^T. \ell(G) \equiv \emptyset \wedge M_2 + \bullet U_2 = M'_2 + \bullet U'_2 + (M_0 - M'_0) + \llbracket G - G_t \rrbracket.$$

Since $\ell(G - G_t) \equiv \emptyset$ we obtain $(M_2, U_2)\mathcal{B}(M'_2, U'_2)$.

(d) Follows directly from Condition 2b and Definition 3.9.

(e) Follows directly from Condition 2a and Definition 3.9. \square

To illustrate the use of Lemmas 6.10 and 6.4, let N' be a plain net and N be its conflict replicating implementation, depicted in Figure 10. Condition (1) says that for any visible transition t in the implementation—this must be **execute** $_j^i$ for some i and j —there must be a transition t' in N' with the same label—this must be i —such that the same token replacement $\llbracket t' \rrbracket$ that results from firing t' in the net N' can also be achieved by t in N together with a multiset G of internal transitions of N . For this to even make sense it is necessary that $S' \subseteq S$, so that $\llbracket t' \rrbracket$ can just as well be seen as a token replacement of N . This condition can be fulfilled by taking G to contain **distribute** $_p$ for every preplace p of i , **fetch** $_{i,j}^{p,c}$ for every preplace p of i and every $c \in p^\bullet$, **fetch** $_j^i$, $u \cdot \text{elide}_i$ for $u \in \Omega_i$, and **finalise** $_i$.

In the proof of Lemma 6.10/6.4, a branching bisimulation is constructed between the markings of N' and N , by relating any reachable marking M' of N' with the corresponding marking $M' + (M_0 - M'_0)$ of N ; the latter is the marking M' seen as a marking of N , together with those places in $S \setminus S'$ that are marked initially (or by default). In addition, M' is also related to markings obtained from $M' + (M_0 - M'_0)$ by adding or subtracting the token replacement due to firing some internal transitions of N . For instance, compared to the

state of N given by the marking $M' + (M_0 - M'_0)$ it could be that **finalise** ^{i} has not yet fired—so that **ack** _{i} (t) is marked for all $t \in \Omega_i$ instead of the postplaces r of i —and that **distribute** _{p} has already fired for some place p . This gives rise to the marking $M' + (M_0 - M'_0) + \llbracket G \rrbracket$ being related to M' , with $G = -\{\text{finalise}^i\} + \{\text{distribute}_p\}$. To show that the relation really is a branching bisimulation with explicit divergence it suffices to check the conditions (a)–(c). That these are enough to obtain the stronger conditions (a)–(e) of Lemma 6.9/6.3 follows with help of the new condition (1).

In the proof of Lemma 6.4 the bisimulation constructed in the proof of Lemma 6.10 is strengthened to a split bisimulation by taking account of the sets U' and U of transitions currently firing in N' and N . Here we need to require that U' and U carry the same multiset of labels. Moreover, the preplaces of U' and U need to be added to M' and M when determining that they are reachable markings, and in relating these markings to each other; for these purposes we thus use the markings we would have had before starting the transitions that are currently firing. On the other hand, M' and M themselves need to be markings (i.e. put a nonnegative number of tokens in each place), and in conditions (a)–(c) only those transitions matter that can be fired from M' and M themselves—without the preplaces of U' and U .

In Lemma 6.4 a relation is explored between markings \bar{M} and $\bar{M} + \llbracket H \rrbracket$ (where \bar{M} is $M' + \bullet U' + (M_0 - M'_0)$ of Lemma 6.4, $H := G$, and $\bar{M} + \llbracket H \rrbracket$ is $M + \bullet U$ of Lemma 6.4). In such a case, we can think of \bar{M} as an “original marking”, and of $\bar{M} + \llbracket H \rrbracket$ as a modification of this marking by the token replacement $\llbracket H \rrbracket$. The next lemma provides a method to trace certain places s marked by $\bar{M} + \llbracket H \rrbracket$ (or transitions t that are enabled under $\bar{M} + \llbracket H \rrbracket$) back to places that must have been marked by \bar{M} before taking into account the token replacement $\llbracket H \rrbracket$. Such places are called *faithful origins* of s (or t). In tracking the faithful origins of places and transitions, we assume that the places marked by \bar{M} are taken from a set S_+ and the transitions in H from a set T_+ . In Lemma 6.7 we furthermore assume that the flow relation restricted to $S \cup T_+$ is acyclic. We will need this lemma in proving the correctness of our final method of proving $N \approx_{bSTb}^{\Delta} N'$.

Definition 6.5. Let $N = (S, T, F, M_0, \ell)$ be a Petri net, $T_+ \subseteq T$ a set of transitions and $S_+ \subseteq S$ a set of places.

- A *path* in N is an alternating sequence $\pi = x_0 x_1 x_2 \cdots x_n \in (S \cup T)^*$ of places and transitions, such that $F(x_i, x_{i+1}) > 0$ for $0 \leq i < n$. The *arc weight* $F(\pi)$ of such a path is the product $\prod_0^{n-1} F(x_i, x_{i+1})$.
- A place $s \in S$ is called *faithful* w.r.t. T_+ and S_+ iff $|\{s\} \cap S_+| + \sum_{t \in T_+} F(t, s) = 1$.
- A path $x_0 x_1 x_2 \cdots x_n \in (S \cup T)^*$ from x_0 to x_n is *faithful* w.r.t. T_+ and S_+ iff all intermediate nodes x_i for $0 \leq i < n$ are either transitions in T_+ or faithful places w.r.t. T_+ and S_+ .
- For $x \in S \cup T$, the *infinitary multiset* $*x \in (\mathbb{N} \cup \{\infty\})^{S_+}$ of *faithful origins* of x is given by $*x(s) = \sup\{F(\pi) \mid \pi \text{ is a faithful path from } s \in S_+ \text{ to } x\}$. (So $*x(s) = 0$ if no such path exists.)

Suppose a marking M is reachable from a marking $\bar{M} \in \mathbb{N}^{S_+}$ by firing transitions from T_+ only. So $M = \bar{M} + \llbracket H \rrbracket$ for some $H \in_F \mathbb{N}^{T_+}$. Then, if a faithful place s bears a token under M —i.e. $M(s) > 0$ —this token has a unique source: if $s \in S_+$ it must stem from \bar{M} and otherwise it must be produced by the unique transition $t \in T_+$ with $F(t, s) = 1$.

Now consider a period in the evolution of the net N that starts with the marking \bar{M} , and during which only transitions from T_+ fire. Suppose $\pi = x_0 x_1 x_2 \cdots x_n$ is a faithful

path from a place $x_0 \in S_+$ to a either a faithful place x_n that gets marked at some point during this period or a transition x_n that fires during (or right after) this period. In that case a token, left on x_0 by the marking \bar{M} , must have travelled along that path from x_0 to x_n —where a token is understood to visit a transition when that transition fires. Namely, if x_{i+1} is a transition that fired at some point, then its (faithful) preplace x_i must have been marked right beforehand; and if a faithful place x_{j+i} was marked at some point, then $x_{j+i} \notin S_+$ and the token in x_{j+i} must have been produced by the transition $x_i \in T_+$.

Note that $F(\pi)$ is the product of all arc weights in the path on arcs from places to transitions; for all the weights on arcs from transitions in T_+ to faithful places are 1. Taking arc weights into account, for every token in x_n as many as $F(\pi)$ tokens must have started in x_0 . Namely, for a transition x_{i+1} to fire once, $F(x_i, x_{i+1})$ tokens must have come from place x_i , and for each token in a faithful place x_{j+1} , the transition x_j must have fired once.

In a net without arc weights, $*x$ is always a set, namely the set of places s in S_+ from which the flow relation of the net admits a path to x that passes only through faithful places and transitions from T_+ (with the possible exception of x itself). For nets with arc weights, the underlying set of $*x$ is the same, and the multiplicity of $s \in *x$ is obtained by multiplying all arc weights on the qualifying path from s to x ; in case of multiple such paths, we take the upper bound over all such paths (which could yield the value ∞). It follows from the analysis above that if a faithful place x gets marked, or a transition x enabled, during a period as described above, then at least $*x(s)$ tokens must have been present in s at the beginning of this period. Lemma 6.7 formalises this analysis by comparing a marking $\bar{M} + \llbracket H \rrbracket$ that marks or enables x (possibly multiple times) with the marking \bar{M} that marks the faithful origins $*x$ of x . Here $H \in_F \mathbb{N}^{T_+}$ is the multiset of transitions whose firing converts \bar{M} into $\bar{M} + \llbracket H \rrbracket$. However, Lemma 6.7 does not require that this multiset actually can be fired in any particular order. To enable that generalisation, it must assume that $F \upharpoonright (S \cup T_+)$ is acyclic.

For $k \neq 0$, we have $k \cdot *x(s) = \sup\{k \cdot F(\pi) \mid \pi \text{ is a faithful path from } s \in S_+ \text{ to } x\}$. In order to also have this equality for $k = 0$ and $*x(s) = \infty$ we define $0 \cdot \infty := 0$ in this context.

Observation 6.6. Let (S, T, F, M_0, ℓ) be a Petri net, $T_+ \subseteq T$ a set of transitions and $S_+ \subseteq S$ a set of places. For faithful places s and transitions $t \in T$ we have

$$*s = \begin{cases} \{s\} & \text{if } s \in S_+ \\ *t & \text{if } t \in T_+ \wedge F(t, s) = 1 \end{cases} \quad *t = \bigcup \{F(s, t) \cdot *s \mid s \in \bullet t \wedge s \text{ faithful}\}.$$

Lemma 6.7. Let (S, T, F, M_0, ℓ) be a Petri net, $T_+ \subseteq T$ a set of transitions such that $F \upharpoonright (S \cup T_+)$ is acyclic, and $S_+ \subseteq S$ a set of places. Let $\bar{M} \in \mathbb{N}^{S_+}$ and $H \in_F \mathbb{N}^{T_+}$, such that $\bar{M} + \llbracket H \rrbracket \in \mathbb{N}^S$ (i.e. places occur only non-negatively in $\bar{M} + \llbracket H \rrbracket$). Then

- (a) for any faithful place s w.r.t. T_+ and S_+ we have $(\bar{M} + \llbracket H \rrbracket)(s) \cdot *s \leq \bar{M}$;
- (b) for any $k \in \mathbb{N}$, and any transition t with $(\bar{M} + \llbracket H \rrbracket)[k \cdot \{t\}]$, we have $k \cdot *t \leq \bar{M}$.

Proof. We apply induction on $|H|$. In the base case, $H = \emptyset$, which formally is included in the induction step, (a) follows directly from the assumption that $\bar{M} \in \mathbb{N}^{S_+}$ and the observation that $*s = \{s\}$.

(a). When $(\bar{M} + \llbracket H \rrbracket)(s) = 0$ it trivially follows that $(\bar{M} + \llbracket H \rrbracket)(s) \cdot *s \leq \bar{M}$. So suppose $(\bar{M} + \llbracket H \rrbracket)(s) > 0$. Then either $s \in S_+$ or there is a unique $t \in T_+$ with $H(t) > 0$ and $F(t, s) = 1$. In the first case, using that $s \in u^\bullet$ for no $u \in T_+$, we have $(\bar{M} + \llbracket H \rrbracket)(s) \leq \bar{M}(s)$, so $(\bar{M} + \llbracket H \rrbracket)(s) \cdot *s \leq \bar{M}(s) \cdot \{s\} \leq \bar{M}$. In the latter case, we have $(\bar{M} + \llbracket H \rrbracket)(s) \leq$

$\bar{M}(s) + \sum_{u \in T_+} H(u) \cdot F(u, s) = \bar{M}(s) + H(t) = H(t)$ and $*s = *t$. Thus:

$$(\bar{M} + \llbracket H \rrbracket)(s) \cdot *s \leq H(t) \cdot *t. \quad (6.4)$$

Let $U := \{u \in T_+ \mid H(u) > 0 \wedge uF^+t\}$ be the set of transitions occurring in H from which the flow relation of the net offers a non-empty path to t . As $F \upharpoonright (S \cup T_+)$ is acyclic, $t \notin U$, so $H \upharpoonright U < H$. Let s' be any place with $s' \in \bullet u$ for some transition $u \in U$. Then, by construction of U , it cannot happen that $s' \in v^\bullet$ for some transition $v \notin U$ with $H(v) > 0$. Hence $(\bar{M} + \llbracket H \upharpoonright U \rrbracket)(s') \geq (\bar{M} + \llbracket H \rrbracket)(s') \geq 0$. Moreover, for any other place s'' we have $\bullet(H \upharpoonright U)(s'') = 0$ and thus $(\bar{M} + \llbracket H \upharpoonright U \rrbracket)(s'') \geq \bar{M}(s'') \geq 0$. It follows that $\bar{M} + \llbracket H \upharpoonright U \rrbracket \in \mathbb{N}^S$.

For each $s''' \in \bullet t$ we have $(H - H \upharpoonright U)^\bullet(s''') = 0$ and $\bullet(H - H \upharpoonright U)(s''') \geq H(t) \cdot \bullet t(s''')$ and therefore $0 \leq (\bar{M} + \llbracket H \rrbracket)(s''') \leq (\bar{M} + \llbracket H \upharpoonright U \rrbracket)(s''') - H(t) \cdot \bullet t(s''')$. For this reason, $H(t) \cdot \bullet t \leq \bar{M} + \llbracket H \upharpoonright U \rrbracket$. It follows that $(\bar{M} + \llbracket H \upharpoonright U \rrbracket)[H(t) \cdot \{t\}]$. Thus, by (6.4) and induction, $(\bar{M} + \llbracket H \rrbracket)(s) \cdot *s \leq H(t) \cdot *t \leq \bar{M}$.

(b). Let $(\bar{M} + \llbracket H \rrbracket)[k \cdot \{t\}]$. For any faithful $s \in \bullet t$ we have $(\bar{M} + \llbracket H \rrbracket)(s) \geq k \cdot F(s, t)$, and thus, using (a),

$$k \cdot F(s, t) \cdot *s \leq (\bar{M} + \llbracket H \rrbracket)(s) \cdot *s \leq \bar{M}.$$

Therefore, by Observation 6.6, $k \cdot *t = \bigcup \{k \cdot F(s, t) \cdot *s \mid s \in \bullet t \wedge s \text{ faithful}\} \leq \bar{M}$. \square

As a (forthcoming) application of Lemma 6.7—in fact the only one we'll need in this paper—consider the branching split bisimulation with explicit divergence between a net N' and its conflict replicating implementation N that is constructed according to the proof of Lemma 6.4. When a split marking (M', U') is related to (M, U) , then $M + \bullet U = M' + \bullet U' + (M_0 - M'_0) + \llbracket G \rrbracket$ for a signed multiset G of internal transitions of N . Furthermore suppose that G is a true multiset over the set of transitions T_+ , consisting of **distribute_p**, **initialise_j** · **fire** and **transfer_j^h** · **fire** only (for arbitrary p, j and h). Take $\bar{M} := M' + \bullet U' + (M_0 - M'_0)$, $H := G$ and thus $M + \bullet U = \bar{M} + \llbracket H \rrbracket$. Let $S_+ := S' \cup \{s \in S \mid (M_0 - M'_0)(s) > 0\}$. Then p **distribute_p** p_i **initialise_i** · **fire** **pre_jⁱ** **execute_jⁱ** is a faithful path from p to **execute_jⁱ**. The arc weight of this path is $F'(p, i)$. So $*\text{execute}_j^i \geq F'(p, i)$. Thus if **execute_jⁱ** is enabled under $M + \bullet U$ then \bar{M} must place at least $F'(p, i)$ tokens in the place p . As this reasoning applies to every preplace p of i , it follows that i is enabled under $M' + \bullet U'$.

The following theorem is the main result of this section. It presents a method for proving $N \approx_{bSTb}^{\Delta} N'$ for N a net and N' a plain net. Its main advantage w.r.t. directly using the definition, or w.r.t. application of Lemma 6.3 or 6.4, is the replacement of requirements on the dynamic behaviour of nets by structural requirements. Such requirements are typically easier to check. Replacing the requirement “ $M + \bullet U \in [M_0]_N$ ” in Condition 5 by “ $M + \bullet U \in \mathbb{N}^S$ ” would have yielded an even more structural version of this theorem; however, that version turned out not to be strong enough for the verification task performed in Section 7.

Theorem 6.8. *Let $N = (S, T, F, M_0, \ell)$ be a net and $N' = (S', T', F', M'_0, \ell')$ be a plain net with $S' \subseteq S$ and $M'_0 = M_0 \upharpoonright S'$. Suppose there exist sets $T_+ \subseteq T$ and $T_- \subseteq T$ and a class $NF \subseteq \mathbb{Z}^T$, such that*

- (1) $F \upharpoonright (S \cup T_+)$ is acyclic.
- (2) $F \upharpoonright (S \cup T_-)$ is acyclic.
- (3) $\forall t \in T, \ell(t) \neq \tau. \exists t' \in T', \ell(t') = \ell(t). (\bullet t' \leq *t \wedge \exists G \in_F \mathbb{N}^T, \ell(G) \equiv \emptyset. \llbracket t' \rrbracket = \llbracket t + G \rrbracket)$.

Here $*t$ is the multiset of faithful origins of t w.r.t. T_+ and $S' \cup \{s \in S \mid M_0(s) > 0\}$.

- (4) There exists a function $f : T \rightarrow \mathbb{N}$ with $f(t) > 0$ for all $t \in T$, extended to \mathbb{Z}^T as in Definition 2.1, such that for each $G \in_F \mathbb{Z}^T$ with $\ell(G) \equiv \emptyset$ there is an $H \in_F NF$ with $\ell(H) \equiv \emptyset$, $\llbracket H \rrbracket = \llbracket G \rrbracket$ and $f(H) = f(G)$.
- (5) For every $M' \in \mathbb{N}^{S'}$, $U' \in \mathbb{N}^{T'}$ and $U \in \mathbb{N}^T$ with $\ell(U) = \ell'(U')$ and $M' + \bullet U' \in [M'_0]_{N'}$, there is an $H_{M',U} \in_F \mathbb{N}^{T+}$ with $\ell(H_{M',U}) \equiv \emptyset$, such that for each $H \in_F NF$ with $M := M' + \bullet U' + (M_0 - M'_0) + \llbracket H \rrbracket - \bullet U \in \mathbb{N}^S$ and $M + \bullet U \in [M_0]_N$:
- (a) $M_{M',U} := M' + \bullet U' + (M_0 - M'_0) + \llbracket H_{M',U} \rrbracket - \bullet U \in \mathbb{N}^S$,
 - (b) if $M' \xrightarrow{a}$ with $a \in \text{Act}$ then $M_{M',U} \xrightarrow{a}$,
 - (c) $H \leq H_{M',U}$.
 - (d) if $H(u) < 0$ then $u \in T_-$,
 - (e) if $H(u) < 0$ and $H(t) > 0$ then $\bullet u \cap \bullet t = \emptyset$,
 - (f) if $H(u) < 0$ and $(M + \bullet U)[t]$ with $\ell(t) \neq \tau$ then $\bullet u \cap \bullet t = \emptyset$,
 - (g) if $(M + \bullet U)[\{t\} + \{u\}]$ and $t', u' \in T'$ with $\ell'(t') = \ell(t)$ and $\ell'(u') = \ell(u)$, then $\bullet t' \cap \bullet u' = \emptyset$.

Then $N \approx_{bSTb}^{\Delta} N'$.

Proof: It suffices to show that Condition 2 of Lemma 6.4 holds (as Condition 1 of Lemma 6.4 is part of Condition 3 above). So let $G \in_F \mathbb{Z}^T$ with $\ell(G) \equiv \emptyset$, $M' \in \mathbb{N}^{S'}$, $U' \in \mathbb{N}^{T'}$ and $U \in \mathbb{N}^T$ with $\ell'(U') = \ell(U)$, $M' + \bullet U' \in [M'_0]_{N'}$, $M := M' + \bullet U' + (M_0 - M'_0) + \llbracket G \rrbracket - \bullet U \in \mathbb{N}^S$ and $M + \bullet U \in [M_0]_N$.

- (a) Suppose $M \xrightarrow{\tau} M_1 \xrightarrow{\tau} M_2 \xrightarrow{\tau} \dots$. Then there are transitions $t_i \in T$ with $\ell(t_i) = \tau$, for all $i \geq 1$, such that $M[t_1]M_1[t_2]M_2[t_3]\dots$. As also $(M + \bullet U)[t_1](M_1 + \bullet U)[t_2](M_2 + \bullet U)[t_3]\dots$, it follows that $(M_i + \bullet U) \in [M_0]_N$ for all $i \geq 1$. Let $G_0 := G$ and for all $i \geq 1$ let $G_i := G_{i-1} + \{t_i\}$. Then $\ell(G_i) \equiv \emptyset$ and $M_i = M' + \bullet U' + (M_0 - M'_0) + \llbracket G_i \rrbracket - \bullet U$. Moreover, $f(G_i) = f(G_{i-1}) + f(t_i) > f(G_{i-1})$. For all $i \geq 0$, using Condition 4, let $H_i \in_F NF$ be so that $\llbracket H_i \rrbracket = \llbracket G_i \rrbracket$ and $f(H_i) = f(G_i)$. Then $M_i = M' + \bullet U' + (M_0 - M'_0) + \llbracket H_i \rrbracket - \bullet U$ and $f(H_0) < f(H_1) < f(H_2) < \dots$. However, from Condition 5c we get $f(H_i) \leq f(H_{M'})$ for all $i \geq 0$. The sequence $M \xrightarrow{\tau} M_1 \xrightarrow{\tau} M_2 \xrightarrow{\tau} \dots$ therefore must be finite.
- (b) Now suppose $M' \xrightarrow{a}$ with $a \in \text{Act}$. By Condition 4 above there exists an $H \in_F NF$ such that $\ell(H) \equiv \emptyset$ and $\llbracket H \rrbracket = \llbracket G \rrbracket$, and hence $M = M' + \bullet U' + (M_0 - M'_0) + \llbracket H \rrbracket - \bullet U$. Let $H_- := \{u \in T \mid H(u) < 0\}$.

- First suppose $H_- \neq \emptyset$. By Condition 5d, $H_- \subseteq T_-$. By Condition 2, the relation $<_- := (F \upharpoonright (S \cup T_-))^+$ is a partial order on $S \cup T_-$, and hence on H_- . Let u be a minimal transition in H_- w.r.t. $<_-$. By definition, for all $s \in S$,

$$M(s) = M'(s) + \bullet U'(s) + (M_0 - M'_0)(s) + \sum_{t \in T} H(t) \cdot F(t, s) + \sum_{t \in T} -H(t) \cdot F(s, t) + \sum_{t \in U} -U(t) \cdot F(t, s). \quad (6.5)$$

As $M'_0 = M_0 \upharpoonright S'$, we have $M'_0 \leq M_0$. Hence the first three summands in this equation are always nonnegative. Now assume $s \in \bullet u$. Since u is minimal w.r.t. $<_-$, there is no $t \in T$ with $H(t) < 0$ and $F(t, s) \neq 0$. Hence also all summands $H(t) \cdot F(t, s)$ are nonnegative. By Condition 5e, there is no $t \in T$ with $H(t) > 0$ and $F(s, t) \neq 0$, so all summands $-H(t) \cdot F(s, t)$ are nonnegative as well. By Condition 5f, there is no $t \in T$ with $U(t) > 0$ and $F(s, t) \neq 0$, for this would imply that $\ell(t) \neq \tau$ and $(M + \bullet U)[t]$, so no summands in (6.5) are negative. Thus $0 \leq -H(u) \cdot F(s, u) \leq M(s)$. Since

$H(u) \leq -1$, this implies $M(s) \geq F(s, u)$. Hence u is enabled in M . As $\ell(u) = \tau$, we have $M \xrightarrow{\tau}$.

- Next suppose $H_- = \emptyset$ but $H \neq H_{M',U}$. Let $H^\smile := \{u \in T \mid H_{M',U}(u) - H(u) > 0\}$. Then $H^\smile \neq \emptyset$ by Condition 5c. Since $H_{M',U} \in_F \mathbb{N}^{T_+}$, $H^\smile \subseteq T_+$. By Condition 1, $<_+ := (F \upharpoonright (S \cup T_+))^+$ is a partial order on $S \cup T_+$, and hence on H^\smile . Let u be a minimal transition in H^\smile w.r.t. $<_+$. We have $M = M' + \bullet U' + (M_0 - M'_0) + \llbracket H_{M',U} + (H - H_{M',U}) \rrbracket - \bullet U = M_{M',U} + \llbracket H - H_{M',U} \rrbracket$. Hence, for all $s \in S$,

$$M(s) = M_{M',U}(s) + \sum_{t \in T} (H - H_{M',U})(t) \cdot F(t, s) + \sum_{t \in T} -(H - H_{M',U})(t) \cdot F(s, t). \quad (6.6)$$

By Condition 5a, $M_{M',U} \in \mathbb{N}^S$. By Condition 5c, $H - H_{M',U} \leq 0$. For $s \in \bullet u$ there is moreover no $t \in H^\smile$ with $s \in t^\bullet$, so no $t \in T$ with $(H - H_{M',U})(t) < 0$ and $F(t, s) \neq 0$. Hence no summands in (6.6) are negative. It thereby follows that $0 \leq -(H - H_{M',U})(u) \cdot F(s, t) \leq M(s)$. Since $(H - H_{M',U})(u) \leq -1$, this implies $M(s) \geq F(s, u)$. Hence u is enabled in M . As $\ell(u) = \tau$, we have $M \xrightarrow{\tau}$.

- Finally suppose $H = H_{M',U}$. Then $M = M_{M',U}$ and $M \xrightarrow{a}$ follows by Condition 5b.
- (c) Next suppose $M \xrightarrow{a}$ with $a \in \text{Act}$. Then there is a $t \in T$ with $\ell(t) = a \neq \tau$ and $M[t]$. So $(M + \bullet U)[t]$. We will first show that $(M' + \bullet U') \xrightarrow{a}$. By Condition 4 there exists an $H_0 \in_F NF \subseteq \mathbb{Z}^T$ such that $\ell(H_0) \equiv \emptyset$ and $\llbracket H_0 \rrbracket = \llbracket G \rrbracket$, and hence $M + \bullet U = M' + \bullet U' + (M_0 - M'_0) + \llbracket H_0 \rrbracket \in [M_0]_N$. For our first step, it suffices to show that whenever $H \in_F NF$ with $M_H := M' + \bullet U' + (M_0 - M'_0) + \llbracket H \rrbracket \in [M_0]$ and $M_H[t]$, then $(M' + \bullet U') \xrightarrow{a}$. We show this by induction on $f(H_{M',U} - H)$, observing that $f(H_{M',U} - H) \in \mathbb{N}$ by Conditions 5c (with empty U) and 4.

We consider two cases, depending on the emptiness of $H_- := \{u \in T \mid H(u) < 0\}$.

First assume $H_- = \emptyset$. Then $H \in_F \mathbb{N}^T$. By Condition 5c (with empty U) we even have $H \in_F \mathbb{N}^{T_+}$. Let $*t$ denote the multiset of faithful origins of t w.r.t. T_+ and $S_+ := S' \cup \{s \in S \mid M_0(s) > 0\}$. By Lemma 6.7(b), taking $k=1$ and $\bar{M} := M' + \bullet U' + (M_0 - M'_0)$, and using Condition 1 of Theorem 6.8, $*t \leq M' + \bullet U' + (M_0 - M'_0)$. So by Condition 3 of Theorem 6.8 there is a $t' \in T'$ with $\ell(t') = \ell(t)$ and $\bullet t' \leq M' + \bullet U' + (M_0 - M'_0)$. Since $\bullet t' \in \mathbb{N}^{S'}$ and $M'_0 = M_0 \upharpoonright S'$, this implies $\bullet t' \leq M' + \bullet U'$. It follows that $(M' + \bullet U')[t']_{N'}$ and hence $(M' + \bullet U') \xrightarrow{a}$.

Now assume $H_- \neq \emptyset$. By the same proof as for (b) above, case $H_- \neq \emptyset$, there is a transition $u \in H_-$ that is enabled in M_H . So $M_H[u]M_1$ for some $M_1 \in [M_0]_N$, and $M_1 = M' + \bullet U' + (M_0 - M'_0) + \llbracket H + u \rrbracket$. By Condition 5f of Theorem 6.8 (still with empty U), $\bullet u \cap \bullet t = \emptyset$, and thus $M_1[t]$. By Condition 4 of Theorem 6.8 there exists an $H_1 \in_F NF$ such that $\ell(H_1) \equiv \emptyset$, $\llbracket H_1 \rrbracket = \llbracket H + u \rrbracket$, and $f(H_1) = f(H + u) > f(H)$. Thus $M_1 = M_{H_1}$ and $f(H_{M',U} - H_1) < f(H_{M',U} - H)$. By induction we obtain $(M' + \bullet U') \xrightarrow{a}$.

By the above reasoning, there is a $t' \in T'$ such that $\ell'(t') = \ell(t)$ and $(M' + \bullet U')[t']$. Now take any $u' \in U'$. Then there must be an $u \in U$ with $\ell'(u') = \ell(u)$. Since $M[t]$, we have $(M + \bullet U)[\{t\} + \{u\}]$ and by Condition 5g we obtain $\bullet t' \cap \bullet u' = \emptyset$. It follows that $M'[t']$, and hence $M' \xrightarrow{a}$. \square

Theorem 6.8 will be applied in Section 7 to show the correctness of our conflict replicating implementation N of a given net N' . A crucial observation about N is that its internal transitions can be partitioned into a set T_+ of transitions (3 boxes in Figure 10) that have to occur before firing execute_j^i (for some i and j) and a set T_- of transitions (14 boxes) that can only occur afterwards. In the construction of our bisimulation we consider markings of

the form $M' + \bullet U' + (M_0 - M'_0) + \llbracket H \rrbracket$, where H is a signed multiset of internal transitions that tells how much the marking deviates from the marking $M' + \bullet U' + (M_0 - M'_0)$ of N . The bisimulation relates both markings of N to the marking $M' + \bullet U'$ of N' . When an internal transition of N fires, the related marking of N' remains the same. However, when N fires a visible transition execute_j^i then the related marking of N' becomes $M' + \bullet U' + \llbracket i \rrbracket$, so in view of the structural property in Lemma 6.4(1), a new set H' can be calculated as $H' := H - G$, where G is the signed multiset for which $\llbracket i \rrbracket = \llbracket \text{execute}_j^i + G \rrbracket$. A consequence of this is that elements of T_+ only occur with positive multiplicities in H , whereas elements of T_- occur only with negative multiplicities.

To be precise, it may be that two different sets H_1 and H_2 yield the same token replacement, i.e. $\llbracket H_1 \rrbracket = \llbracket H_2 \rrbracket$. As a result of this, there may be multiple ways to write a marking as $M' + \bullet U' + (M_0 - M'_0) + \llbracket H \rrbracket$ for given M' and U' . The above applies only when converting the signed multisets H to a normal form NF that eliminates this ambiguity.

For given M' and U' , the multiset $H_{M',U'}$ is an upper bound of the possible choices of H for which $M' + \bullet U' + (M_0 - M'_0) + \llbracket H \rrbracket$ can be a reachable marking. This is expressed by Condition 5c. If all internal transitions in $H_{M',U'}$ have fired, the next transition must be an external one. Now the conditions of Theorem 6.8 guarantee that as long as this upper bound is not reached, the net N can perform internal actions, and when it is reached (and possibly also beforehand) it can perform the same actions as the net N' under marking M' . Condition 4 moreover guarantees that this upper bound will be reached in finitely many steps. Due to the need to renormalise the signed multisets H after adding elements to them, this is not straightforward.

These considerations imply that transitions fired by N' can be simulated by N . The other direction involves similar arguments, together with an application of Lemma 6.7.

Digression: Interleaving semantics. Above, a method is presented for establishing the equivalence of two Petri nets, one of which known to be plain, up to branching ST-bisimilarity with explicit divergence. Here, we simplify this result into a method for establishing the equivalence of the two nets up interleaving branching bisimilarity with explicit divergence. This result is not applied in the current paper.

Lemma 6.9. *Let $N = (S, T, F, M_0, \ell)$ and $N' = (S', T', F', M'_0, \ell')$ be two nets, N' being plain. Suppose there is a relation $\mathcal{B} \subseteq \mathbb{N}^S \times \mathbb{N}^{S'}$ such that*

- (a) $M_0 \mathcal{B} M'_0$,
- (b) if $M_1 \mathcal{B} M'_1$ and $M_1 \xrightarrow{\tau} M_2$ then $M_2 \mathcal{B} M'_1$,
- (c) if $M_1 \mathcal{B} M'_1$ and $M_1 \xrightarrow{a} M_2$ for some $a \in \text{Act}$ then $\exists M'_2. M'_1 \xrightarrow{a} M'_2 \wedge M_2 \mathcal{B} M'_2$,
- (d) if $M_1 \mathcal{B} M'_1$ and $M'_1 \xrightarrow{a}$ for some $a \in \text{Act}$ then either $M_1 \xrightarrow{a}$ or $M_1 \xrightarrow{\tau}$
- (e) and there is no infinite sequence $M \xrightarrow{\tau} M_1 \xrightarrow{\tau} M_2 \xrightarrow{\tau} \dots$ with $M \mathcal{B} M'$ for some M' .

Then N and N' are interleaving branching bisimilar with explicit divergence.

Proof. This follows directly from Lemma 6.2 by taking $(\mathfrak{S}_1, \mathfrak{T}_1, \mathfrak{M}_{o1})$ and $(\mathfrak{S}_2, \mathfrak{T}_2, \mathfrak{M}_{o2})$ to be the interleaving LTSs associated to N and N' respectively, using the fact that the LTS associated to a plain net is deterministic. \square

Lemma 6.10. *Let $N = (S, T, F, M_0, \ell)$ be a net and $N' = (S', T', F', M'_0, \ell')$ be a plain net with $S' \subseteq S$ and $M'_0 = M_0 \upharpoonright S'$. Suppose:*

- (1) $\forall t \in T, \ell(t) \neq \tau. \exists t' \in T', \ell(t') = \ell(t). \exists G \in_F \mathbb{N}^T, \ell(G) \equiv \emptyset. \llbracket t' \rrbracket = \llbracket t + G \rrbracket$.

- (2) For any $G \in_F \mathbb{Z}^T$ with $\ell(G) \equiv \emptyset$, $M' \in [M'_0]_{N'}$ and $M := M' + (M_0 - M'_0) + \llbracket G \rrbracket \in [M_0]_N$, it holds that:
- (a) there is no infinite sequence $M \xrightarrow{\tau} M_1 \xrightarrow{\tau} M_2 \xrightarrow{\tau} \dots$,
 - (b) if $M' \xrightarrow{a}$ with $a \in \text{Act}$ then $M \xrightarrow{a}$ or $M \xrightarrow{\tau}$
 - (c) and if $M \xrightarrow{a}$ with $a \in \text{Act}$ then $M' \xrightarrow{a}$.

Then N and N' are interleaving branching bisimilar with explicit divergence.

Proof: Define $\mathcal{B} \subseteq \mathbb{N}^S \times \mathbb{N}^{S'}$ by

$$M \mathcal{B} M' :\Leftrightarrow M' \in [M'_0]_{N'} \wedge \exists G \in_F \mathbb{Z}^T. M = M' + (M_0 - M'_0) + \llbracket G \rrbracket \in [M_0]_N \wedge \ell(G) \equiv \emptyset.$$

It suffices to show that \mathcal{B} satisfies Conditions (a)–(e) of Lemma 6.9.

- (a) Take $G = \emptyset$.
- (b) Suppose $M_1 \mathcal{B} M'_1$ and $M_1 \xrightarrow{\tau} M_2$. Then $\exists G \in_F \mathbb{Z}^T. M_1 = M'_1 + (M_0 - M'_0) + \llbracket G \rrbracket \wedge \ell(G) \equiv \emptyset$ and $\exists t \in T. \ell(t) = \tau \wedge M_2 = M_1 + \llbracket t \rrbracket = M'_1 + (M_0 - M'_0) + \llbracket G + t \rrbracket$. Moreover, $M_1 \in [M_0]_N$ and hence $M_2 \in [M_0]_N$. Furthermore, $M'_1 \in [M'_0]_{N'}$ and $\ell(G + t) \equiv \emptyset$, so $M_2 \mathcal{B} M'_1$.
- (c) Suppose $M_1 \mathcal{B} M'_1$ and $M_1 \xrightarrow{a} M_2$. Then $\exists G \in_F \mathbb{Z}^T. M_1 = M'_1 + (M_0 - M'_0) + \llbracket G \rrbracket \wedge \ell(G) \equiv \emptyset$ and $\exists t \in T. \ell(t) = a \neq \tau \wedge M_2 = M_1 + \llbracket t \rrbracket = M'_1 + (M_0 - M'_0) + \llbracket G + t \rrbracket$. Moreover, $M_1 \in [M_0]_N$ and hence $M_2 \in [M_0]_N$. Furthermore, $M'_1 \in [M'_0]_{N'}$. By Condition 1 of Lemma 6.10, $\exists t' \in T', \ell(t') = \ell(t)$. $\exists G_t \in_F \mathbb{N}^T, \ell(G_t) \equiv \emptyset. \llbracket t \rrbracket = \llbracket t' - G_t \rrbracket$. Substitution of $\llbracket t' - G_t \rrbracket$ for t yields $M_2 = M'_1 + \llbracket t' \rrbracket + (M_0 - M'_0) + \llbracket G - G_t \rrbracket$. By Condition 2c, $M'_1 \xrightarrow{a}$, so $M'_1 \xrightarrow{a} M'_2$ for some $M'_2 \in [M'_0]_{N'}$. As t' is the only transition in T' with $\ell'(t') = a$, we must have $M'_1 \llbracket t' \rrbracket M'_2$. So $M'_1 + \llbracket t' \rrbracket = M'_2$. Since $\ell(G - G_t) \equiv \emptyset$ it follows that $M_2 \mathcal{B} M'_2$.
- (d) Follows directly from Condition 2b.
- (e) Follows directly from Condition 2a. □

The above is a variant of Lemma 6.4 that requires Condition 2 only for $U = U' = \emptyset$, and allows to conclude that N and N' are interleaving branching bisimilar (instead of branching ST-bisimilar) with explicit divergence. Likewise, the below is a variant of Theorem 6.8 that requires Condition 5 only for $U = U' = \emptyset$, and misses Condition 5g.

Theorem 6.11. *Let $N = (S, T, F, M_0, \ell)$ be a net and $N' = (S', T', F', M'_0, \ell')$ be a plain net with $S' \subseteq S$ and $M'_0 = M_0 \upharpoonright S'$. Suppose there exist sets $T_+ \subseteq T$ and $T_- \subseteq T$ and a class $NF \subseteq \mathbb{Z}^T$, such that*

- (1)–(4) *Conditions (1)–(4) from Theorem 6.8 hold, and*
- (5) *For every reachable marking $M' \in [M'_0]_{N'}$ there is an $H_{M'} \in_F \mathbb{N}^{T_+}$ with $\ell(H_{M'}) \equiv \emptyset$, such that for each $H \in_F NF$ with $M := M' + (M_0 - M'_0) + \llbracket H \rrbracket \in [M_0]_N$ one has:*
 - (a) $M_{M'} := M' + (M_0 - M'_0) + \llbracket H_{M'} \rrbracket \in \mathbb{N}^S$,
 - (b) if $M' \xrightarrow{a}$ with $a \in \text{Act}$ then $M_{M'} \xrightarrow{a}$,
 - (c) $H \leq H_{M'}$,
 - (d) if $H(u) < 0$ then $u \in T_-$,
 - (e) if $H(u) < 0$ and $H(t) > 0$ then $\bullet u \cap \bullet t = \emptyset$,
 - (f) if $H(u) < 0$ and $M \llbracket t \rrbracket$ with $\ell(t) \neq \tau$ then $\bullet u \cap \bullet t = \emptyset$.

Then N and N' are interleaving branching bisimilar with explicit divergence.

Proof. A straightforward simplification of the proof of Theorem 6.8. □

7. THE CORRECTNESS PROOF

We now apply the preceding theory to prove the correctness of the conflict replicating implementation.

Theorem 7.1. *Let N' be a finitary plain structural conflict net without a fully reachable pure \mathbf{M} . Then $\mathcal{I}(N') \approx_{bSTb}^{\Delta} N'$.*

Proof: Let $N' = (S', T', F', M'_0, \ell')$ be the given finitary plain structural conflict net without a fully reachable pure \mathbf{M} , and $N = (S, T, F, M_0, \ell)$ be its conflict replicated implementation $\mathcal{I}(N')$. This convention (at the expense of primes in the statement of the theorem) pays off in terms of a significant reduction in the number of primes in this paper.

For future reference, Table 2 provides a place-oriented representation of the conflict replicating implementation of a given net $N' = (S', T', F', M'_0, \ell')$, with the macros for reversible transitions expanded. Here $T^{\leftarrow} = \{\text{initialise}_j \mid j \in T'\} \cup \{\text{transfer}_j^h \mid h <^{\#} j \in T'\}$, $(\text{transfer}_j^h)^{far} = \{\text{trans}_j^h\text{-out}\}$ and $(\text{initialise}_j)^{far} = \{\text{pre}_k^j \mid k \geq^{\#} j\} \cup \{\text{trans}_j^h\text{-in} \mid h <^{\#} j\}$.

Place	Pretransitions arc weights	Posttransitions arc weights	for all
p	finalise_p^i $F'(i, p)$	distribute_p (if $p^{\bullet} \neq \emptyset$)	$p \in S', i \in \bullet p$
p_c	$\begin{cases} \text{distribute}_p \\ \text{initialise}_c \cdot \text{undone} \end{cases}$ $F'(p, c)$	$\text{initialise}_c \cdot \text{fire}$ $F'(p, c)$	$p \in S', c \in p^{\bullet}$
π_c (marked)	$\text{initialise}_c \cdot \text{reset}_i$	$\text{fetch}_{i,j}^{p,c}$ $F'(p, i)$	$j \geq^{\#} i \in p^{\bullet}$
pre_j^i	$\text{initialise}_c \cdot \text{fire}$	$\text{initialise}_c \cdot \text{fire}$	$i \stackrel{\#}{=} c \in T'$
$\text{trans}_j^h\text{-in}$	$\begin{cases} \text{initialise}_i \cdot \text{fire} \\ \text{execute}_j^i \end{cases}$	$\text{initialise}_i \cdot \text{undo}(\text{pre}_j^i)$	$j \geq^{\#} i \in T'$
$\text{trans}_j^h\text{-out}$	$\begin{cases} \text{execute}_j^i \\ \text{transfer}_j^h \cdot \text{undone} \end{cases}$	$\text{transfer}_j^h \cdot \text{fire}$	$h <^{\#} j \in T'$
$\pi_{j\#l}$ (marked)	$\begin{cases} \text{transfer}_j^h \cdot \text{fire} \\ \text{execute}_j^i \end{cases}$	$\text{initialise}_j \cdot \text{undo}(\text{trans}_j^h\text{-in})$	$h <^{\#} j \in T'$
$\text{fetch}_{i,j}^{p,c}\text{-in}$	$\begin{cases} \text{transfer}_j^h \cdot \text{fire} \\ \text{execute}_j^i \end{cases}$	execute_j^i	$h <^{\#} j \in T', i \leq^{\#} j$
$\text{fetch}_{i,j}^{p,c}\text{-out}$	$\begin{cases} \text{fetched}_j^i \\ \text{transfer}_l^j \cdot \text{reset}_c \end{cases}$	$\text{transfer}_j^h \cdot \text{undo}(\text{trans}_j^h\text{-out})$	$h <^{\#} j \in T', i \leq^{\#} j$
$\text{undo}_i(t)$	$\text{execute}_j^i \cdot \text{fire}$	execute_j^i	$i \leq^{\#} j <^{\#} l \in T', c \stackrel{\#}{=} l$
$\text{reset}_i(t)$	execute_j^i	$\text{transfer}_l^j \cdot \text{fire}$	$j \geq^{\#} i \in T', p \in \bullet i, c \in p^{\bullet}$
$\text{ack}_i(t)$	$\text{fetch}_{i,j}^{p,c}$	$\text{fetch}_{i,j}^{p,c}$	$j \geq^{\#} i \in T', p \in \bullet i, c \in p^{\bullet}$
$\text{fired}(t)$	$\text{fetch}_{i,j}^{p,c}\text{-in}$	fetched_j^i	$j \geq^{\#} i \in T', p \in \bullet i, c \in p^{\bullet}$
$\rho_i(t)$	$\text{fetch}_{i,j}^{p,c}\text{-out}$	$\text{execute}_j^i \cdot \text{fire}$	$j \geq^{\#} i \in T', t \in \Omega_i$
$\text{take}(f, t)$	$\text{execute}_j^i \cdot \text{fire}$	$t \cdot \text{undo}_i, t \cdot \text{elide}_i$	$j \geq^{\#} i \in T', t \in \Omega_i$
$\text{took}(f, t)$	fetched_j^i	$t \cdot \text{reset}_i, t \cdot \text{elide}_i$	$j \geq^{\#} i \in T', t \in \Omega_i$
$\rho(t)$	$t \cdot \text{reset}_i, t \cdot \text{elide}_i$	finalise_i	$i \in T', t \in \Omega_i$
$\text{take}(f, t)$	$t \cdot \text{fire}$	initialise_i	$i \in T', t \in \Omega_i$
$\text{took}(f, t)$	$t \cdot \text{fire}$	$t \cdot \text{undo}_i$	$t \in T^{\leftarrow}, \Omega_i \ni t$
$\rho(t)$	$t \cdot \text{undo}_i$	$t \cdot \text{reset}_i$	$t \in T^{\leftarrow}, \Omega_i \ni t$
$\text{take}(f, t)$	$t \cdot \text{undo}_i$	$t \cdot \text{undo}(f)$	$t \in T^{\leftarrow}, \Omega_i \ni t, f \in t^{far}$
$\text{took}(f, t)$	$t \cdot \text{undo}(f)$	$t \cdot \text{undone}$	$t \in T^{\leftarrow}, f \in t^{far}$
$\rho(t)$	$t \cdot \text{undone}$	$t \cdot \text{reset}_i$	$t \in T^{\leftarrow}, \Omega_i \ni t$

Table 2: The conflict replicating implementation.

We will obtain Theorem 7.1 as an application of Theorem 6.8. Following the construction of N described in Section 5.4, we indeed have $S' \subseteq S$ and $M'_0 = M_0 \upharpoonright S'$. Let $T_+ \subseteq T$

be the set of transitions

$$\text{distribute}_p \quad \text{initialise}_j \cdot \text{fire} \quad \text{transfer}_j^h \cdot \text{fire} \quad (7.1)$$

for any applicable values of $p \in S'$ and $h, j \in T'$. Furthermore, $T_- := (T \setminus (T_+ \cup \{\text{execute}_j^i \mid i \leq^\# j \in T'\}))$. We start with checking Conditions 1, 2 and 3 of Theorem 6.8.

1. Let $<_+$ be the partial order on T_+ given by the order of listing in (7.1)—so $\text{initialise}_i \cdot \text{fire} <_+ \text{transfer}_j^h \cdot \text{fire}$, for any $i \in T'$ and $h <^\# j \in T'$, but the transitions $\text{transfer}_j^h \cdot \text{fire}$ and $\text{transfer}_l^k \cdot \text{fire}$ for $(i, j) \neq (k, l)$ are unordered. By examining Table 2 we see that for any place with a pretransition t in T_+ , all its posttransitions u in T_+ appear higher in the $<_+$ -ordering: $t <_+ u$. From this it follows that $F \upharpoonright (S \cup T_+)$ is acyclic.
2. Let $<_-$ be the partial order on T_- given by the column-wise order of the following enumeration of T_- :

$t \cdot \text{undo}_i$	$\text{fetch}_{i,j}^{p,c}$
$\text{transfer}_j^h \cdot \text{undo}(f)$	fetched_j^i
$\text{transfer}_j^h \cdot \text{undone}$	$t \cdot \text{reset}_i$
$\text{initialise}_j \cdot \text{undo}(f)$	$t \cdot \text{elide}_i$
$\text{initialise}_j \cdot \text{undone}$	finalise^i

for any $t \in \{\text{initialise}_j, \text{transfer}_j^h\}$ and any applicable values of $f \in S, p \in S'$, and $h, i, j, c \in T'$. By examining Table 2 we see that for any place with a pretransition t in T_- , all its posttransitions u in T_- appear higher in the $<_-$ -ordering: $t <_- u$. From this it follows that $F \upharpoonright (S \cup T_-)$ is acyclic.

3. The only transitions $t \in T$ with $\ell(t) \neq \tau$ are execute_j^i , with $i \leq^\# j \in T'$. So take $i \leq^\# j \in T'$. Then the only transition $t' \in T'$ with $\ell'(t') = \ell(\text{execute}_j^i)$ is i . Now two statements regarding i and execute_j^i need to be proven. For the first, note that, for any $p \in \bullet i$, the places p, p_i and pre_j^i are faithful w.r.t. T_+ and $S' \cup \{s \in S \mid M_0(s) > 0\}$. Hence $p \text{ distribute}_p p_i \text{ initialise}_i \cdot \text{fire} \text{ pre}_j^i \text{ execute}_j^i$ is a faithful path from p to execute_j^i . The arc weight of this path is $F'(p, i)$. Thus $\bullet i \leq^* \text{execute}_j^i$.

The second statement holds because, for all $i \leq^\# j \in T'$,

$$\llbracket i \rrbracket = \llbracket \text{execute}_j^i \rrbracket + \sum_{p \in \bullet i} (F'(p, i) \cdot \text{distribute}_p + \sum_{c \in p^\bullet} \text{fetch}_{i,j}^{p,c}) + \text{fetched}_j^i + \text{finalise}^i + \sum_{t \in \Omega_i} t \cdot \text{elide}_i. \quad (7.2)$$

To check that these equations hold, note that

$$\begin{aligned} \llbracket \text{distribute}_p \rrbracket &= -\{p\} + \{p_c \mid c \in p^\bullet\}, \\ \llbracket \text{execute}_j^i \rrbracket &= -\{\pi_{j\#l} \mid l \geq^\# j\} + \{\text{fetch}_{i,j}^{p,c}\text{-in} \mid p \in \bullet i, c \in p^\bullet\} + \{\text{undo}_i(t) \mid t \in \Omega_i\}, \\ \llbracket \text{fetch}_{i,j}^{p,c} \rrbracket &= -\{\text{fetch}_{i,j}^{p,c}\text{-in}\} - F'(p, i) \cdot \{p_c\} + \{\text{fetch}_{i,j}^{p,c}\text{-out}\}, \\ \llbracket \text{fetched}_j^i \rrbracket &= -\{\text{fetch}_{i,j}^{p,c}\text{-out} \mid p \in \bullet i, c \in p^\bullet\} + \{\pi_{j\#l} \mid l \geq^\# j\} + \{\text{reset}_i(t) \mid t \in \Omega_i\}, \\ \llbracket t \cdot \text{elide}_i \rrbracket &= -\{\text{undo}_i(t), \text{reset}_i(t) \mid t \in \Omega_i\} + \{\text{ack}_i(t) \mid t \in \Omega_i\}, \\ \llbracket \text{finalise}^i \rrbracket &= -\{\text{ack}_i(t) \mid t \in \Omega_i\} + \sum_{r \in i^\bullet} F'(i, r) \cdot \{r\}. \end{aligned}$$

Before we define the class $NF \subseteq \mathbb{Z}^T$ of signed multisets of transitions in normal form, and verify conditions 4 and 5, we derive some properties of the conflict replicating implementation $N = \mathcal{I}(N')$.

Claim 7.2. For any $M' \in \mathbb{Z}^{S'}$ and $G \in_F \mathbb{Z}^T$ such that $M := M' + (M_0 - M'_0) + \llbracket G \rrbracket \in \mathbb{N}^S$ and for each $i \in T'$ and $t \in \Omega_i$ we have

$$G(t \cdot \text{elide}_i) + G(t \cdot \text{undo}_i) \leq \sum_{j \geq \#i} G(\text{execute}_j^i) \quad (7.3)$$

$$G(\text{finalise}_i^i) \leq G(t \cdot \text{elide}_i) + G(t \cdot \text{reset}_i) \leq \sum_{j \geq \#i} G(\text{fetched}_j^i) \quad (7.4)$$

$$G(t \cdot \text{reset}_i) \leq G(t \cdot \text{undo}_i). \quad (7.5)$$

Moreover, for each $t \in T^{\leftarrow}$ and $f \in t^{\text{far}}$,

$$\sum_{\{\omega | t \in \Omega_\omega\}} G(t \cdot \text{reset}_\omega) \leq G(t \cdot \text{undone}) \leq G(t \cdot \text{undo}(f)) \leq \sum_{\{\omega | t \in \Omega_\omega\}} G(t \cdot \text{undo}_\omega) \leq G(t \cdot \text{fire}) \quad (7.6)$$

and for each appropriate $c, h, i, j, l \in T'$ and $p \in S'$:

$$G(\text{fetched}_j^i) \leq G(\text{fetch}_{i,j}^{p,c}) \leq G(\text{execute}_j^i) \quad (7.7)$$

$$G(\text{initialise}_j \cdot \text{fire}) \leq 1 + \sum_{\omega} G(\text{initialise}_j \cdot \text{reset}_\omega) \quad (7.8)$$

$$G(\text{transfer}_j^h \cdot \text{fire}) - G(\text{transfer}_j^h \cdot \text{undone}) \leq G(\text{initialise}_j \cdot \text{fire}) - G(\text{initialise}_j \cdot \text{undo}(\text{trans}_j^h \text{-in})) \quad (7.9)$$

$$G(\text{transfer}_l^j \cdot \text{fire}) + \sum_{i \leq \#j} G(\text{execute}_j^i) \leq 1 + \sum_{\omega} G(\text{transfer}_l^j \cdot \text{reset}_\omega) + \sum_{i \leq \#j} G(\text{fetched}_j^i) \quad (7.10)$$

$$\text{if } M[\text{execute}_j^i] \text{ then } 1 \leq G(\text{initialise}_i \cdot \text{fire}) - G(\text{initialise}_i \cdot \text{undo}(\text{pre}_j^i)) \quad (7.11)$$

$$\text{if } \exists i. M[\text{execute}_j^i] \text{ then } 1 \leq G(\text{transfer}_j^h \cdot \text{fire}) - G(\text{transfer}_j^h \cdot \text{undo}(\text{trans}_j^h \text{-out})) \quad (7.12)$$

$$F'(p, c) \cdot (G(\text{initialise}_c \cdot \text{fire}) - G(\text{initialise}_c \cdot \text{undone})) + \sum_{j \geq \#i \in p^\bullet} F'(p, i) \cdot G(\text{fetch}_{i,j}^{p,c}) \leq G(\text{distribute}_p) \quad (7.13)$$

$$G(\text{distribute}_p) \leq M'(p) + \sum_{\{i \in T' | p \in i^\bullet\}} G(\text{finalise}_i^i). \quad (7.14)$$

Proof: For any $i \in T'$ and $t \in \Omega_i$, we have

$$M(\text{undo}_i(t)) = \left(\sum_{j \geq \#i} G(\text{execute}_j^i) \right) - G(t \cdot \text{elide}_i) - G(t \cdot \text{undo}_i) \geq 0,$$

given that $M'(\text{undo}_i(t)) = (M_0 - M'_0)(\text{undo}_i(t)) = \emptyset$. In this way, the place $\text{undo}_i(t)$ gives rise to the inequation (7.3) about G . Likewise, the places $\text{ack}_i(t)$, $\text{reset}_i(t)$ and $\rho_i(t)$, respectively, contribute (7.4) and (7.5), whereas $\rho(t)$, $\text{took}(t)$, $\text{take}(t)$ and $\text{fired}(t)$ yield (7.6). The remaining inequations arise from $\text{fetch}_{i,j}^{p,c}$ -out, $\text{fetch}_{i,j}^{p,c}$ -in, π_j , trans_j^h -in, $\pi_{j \neq l}$, pre_j^i , trans_j^h -out, p_c and p , respectively. ■

(7.10) can be rewritten as $T_l^j + \sum_{i \leq \#j} E_j^i \leq 1$, where $T_l^j := G(\text{transfer}_l^j \cdot \text{fire}) - \sum_{\omega} G(\text{transfer}_l^j \cdot \text{reset}_\omega)$ and $E_j^i := G(\text{execute}_j^i) - G(\text{fetched}_j^i)$. By (7.6) $\sum_{\omega} G(\text{transfer}_l^j \cdot \text{reset}_\omega) \leq G(\text{transfer}_l^j \cdot \text{fire})$, so $T_l^j \geq 0$, and likewise, by (7.7), $E_j^i \geq 0$ for all $i \leq \#j$. Hence, for all $i \leq \#j < \#l \in T'$,

$$0 \leq T_l^j \leq 1 \quad 0 \leq E_j^i \leq 1 \quad T_l^j + \sum_{i \leq \#j} E_j^i \leq 1. \quad (7.15)$$

In our next claim we study triples (M, M', G) with

(A) $M \in [M_0]_N$, $M' \in [M'_0]_{N'}$ and $G \in_F \mathbb{Z}^T$,

(B) $M = M' + (M_0 - M'_0) + \llbracket G \rrbracket$,

- (C) $G(\text{finalise}^i) = 0$ for all $i \in T'$,
- (D) $G(\text{distribute}_p) \leq M'(p)$ for all $p \in S'$,
- (E) $G(\text{fetched}_l^k) \geq 0$ for all $k \leq^\# l \in T'$,
- (F) $G(\text{distribute}_p) \geq F'(p, i) \cdot G(\text{execute}_j^i)$ for all $i \leq^\# j \in T'$ and $p \in \bullet i$,
- (G) $0 \leq G(\text{execute}_j^i) \leq 1$ for all $i \leq^\# j \in T'$,
- (H) $G(\text{distribute}_p) \geq F'(p, j) \cdot G(\text{execute}_j^i)$ for all $i \leq^\# j \in T'$ and $p \in \bullet j$,
- (I) (in the notation of (7.15)) if $E_j^i = 1$ with $i \leq^\# j \in T'$ then $T_j^h = 1$ for all $h <^\# j$,
- (J) there are no $j \geq^\# i \stackrel{\#}{=} k \leq^\# l \in T'$ with $(i, j) \neq (k, l)$, $G(\text{execute}_j^i) > 0$ and $G(\text{execute}_l^k) > 0$,
- (K) there are no $i \leq^\# j \stackrel{\#}{=} k \leq^\# l \in T'$ with $(i, j) \neq (k, l)$, $G(\text{execute}_j^i) > 0$ and $G(\text{execute}_l^k) > 0$.

Given such a triple (M_1, M'_1, G_1) and a transition $t \in T$, we define $\text{next}(M_1, M'_1, G_1, t) =: (M, M', G)$ as follows: Let $G_2 := G_1 + \{t\}$. Take $M := M_1 + \llbracket t \rrbracket = M'_1 + (M_0 - M'_0) + \llbracket G_2 \rrbracket$. In case t is not of the form finalise^i we take $M' := M'_1 \in [M'_0]_{N'}$ and $G := G_2 \in_F \mathbb{Z}^T$. In case $t = \text{finalise}^i$ for some $i \in T'$ then $1 = G_2(\text{finalise}^i) \leq \sum_{j \geq^\# i} G_2(\text{execute}_j^i) = \sum_{j \geq^\# i} G_1(\text{execute}_j^i)$ by (C), (7.4) and (7.7), so by (G) and (J) there is a unique $j \geq^\# i$ with $G_1(\text{execute}_j^i) = 1$. We take $M' := M'_1 + \llbracket i \rrbracket$ and $G := G_2 - G_j^i$, where G_j^i is the right-hand side of (7.2).

Claim 7.3.

- (1) If $M_1[t]$ and (M_1, M'_1, G_1) satisfies (A)-(K), then so does $\text{next}(M_1, M'_1, G_1, t)$.
- (2) For any $M \in [M_0]_N$ there exist M' and G such that (A)-(K) hold.

Proof: (2) follows from (1) via induction on the reachability of M . In case $M = M_0$ we take $M' := M'_0$ and $G := \emptyset$. Clearly, (A)–(K) are satisfied.

Hence we now show (1). Let $(M, M', G) := \text{next}(M_1, M'_1, G_1, t)$. We check that (M, M', G) satisfies the requirements (A)–(K).

- (A) By construction, $M \in [M_0]_N$ and $G \in_F \mathbb{Z}^T$. If t is not of the form finalise^i we have $M' = M'_1 \in [M'_0]_{N'}$. Otherwise, by (D) and (F) we have $M'_1(p) \geq G_1(\text{distribute}_p) \geq F'(p, i)$ for all $p \in \bullet i$, and hence $M'_1 \llbracket i \rrbracket$. This in turn implies that $M' = M'_1 + \llbracket i \rrbracket \in [M'_0]_{N'}$.
- (B) In case t is not of the form finalise^i we have

$$M = M_1 + \llbracket t \rrbracket = M'_1 + (M_0 - M'_0) + \llbracket G_1 + t \rrbracket = M' + (M_0 - M'_0) + \llbracket G \rrbracket.$$

In case $t = \text{finalise}^i$ we have $M = M'_1 + (M_0 - M'_0) + \llbracket G_2 \rrbracket = M' + (M_0 - M'_0) + \llbracket G \rrbracket$, using that $\llbracket i \rrbracket = \llbracket G_j^i \rrbracket$.

- (C) In case $t = \text{finalise}^i$ we have $G(\text{finalise}^i) = G_1(\text{finalise}^i) + 1 - G_j^i(\text{finalise}^i) = 0 + 1 - 1 = 0$. Otherwise $G(\text{finalise}^i) = G_1(\text{finalise}^i) + 0 = 0 + 0 = 0$.
- (D) This follows immediately from (C) and (7.14).
- (E) The only time that this invariant is in danger is when $t = \text{finalise}^i$. Then $G = G_1 + \{\text{finalise}^i\} - G_j^i$ for a certain $j \geq^\# i$ with $G_1(\text{execute}_j^i) = 1$. By (J)⁹ $G_1(\text{execute}_l^i) \leq 0$ for all $l \geq^\# i$ with $l \neq j$. Hence by (7.7) $G_1(\text{fetched}_l^i) \leq 0$ for all such l . By (C) $G_2(\text{finalise}^i) = G_1(\text{finalise}^i) + 1 = 1$, so by (7.4) $\sum_{l \geq^\# i} G_1(\text{fetched}_l^i) = \sum_{l \geq^\# i} G_2(\text{fetched}_l^i) > 0$; hence it must be that $G_1(\text{fetched}_j^i) > 0$. By (E)⁹ $G_1(\text{fetched}_l^k) \geq 0$ for all $k \leq^\# l \in T'$. Given that $G_j^i(\text{fetched}_j^i) = 1$ and $G_j^i(\text{fetched}_l^k) = 0$ for all $(k, l) \neq (i, j)$, we obtain $G(\text{fetched}_l^k) \geq 0$ for all $k \leq^\# l \in T'$.

⁹We use (J) and (E) for G_1 only, making use of the induction hypothesis.

- (F) Take $i \leq^{\#} j \in T'$ and $p \in \bullet i$. There are two occasions where the invariant is in danger: when $t = \text{execute}_j^i$ and when $t = \text{finalise}^k$ with $k \in T'$. First let $t = \text{execute}_j^i$. Then $M_1[\text{execute}_j^i]$. Thus,

$$\begin{aligned}
& G(\text{distribute}_p) \\
& \geq F'(p, i) \cdot (G(\text{initialise}_i \cdot \text{fire}) - G(\text{initialise}_i \cdot \text{undone})) + \sum_{h \geq^{\#} g \in p^{\bullet}} F'(p, g) \cdot G(\text{fetch}_{g,h}^{p,i}) \\
& \geq F'(p, i) \cdot (G(\text{initialise}_i \cdot \text{fire}) - G(\text{initialise}_i \cdot \text{undone})) + \sum_{h \geq^{\#} g \in p^{\bullet}} F'(p, g) \cdot G(\text{fetched}_h^g) \\
& \geq F'(p, i) \cdot (G(\text{initialise}_i \cdot \text{fire}) - G(\text{initialise}_i \cdot \text{undone})) + F'(p, i) \cdot G(\text{fetched}_j^i) \\
& \geq F'(p, i) \cdot \left((G(\text{initialise}_i \cdot \text{fire}) - G(\text{initialise}_i \cdot \text{undo}(\text{pre}_j^i))) + G(\text{fetched}_j^i) \right) \\
& \geq F'(p, i) \cdot (1 + G(\text{fetched}_j^i)) \\
& \geq F'(p, i) \cdot G(\text{execute}_j^i)
\end{aligned}$$

by (7.13), (7.7), (E), (7.6), (7.11) and (7.15), respectively. By (7.6) $G(\text{initialise}_i \cdot \text{fire}) - G(\text{initialise}_i \cdot \text{undone}) \geq 0$. So by (7.13), (E), and (7.7) $G(\text{distribute}_p) \geq 0$. For this reason we may assume, w.l.o.g., that $G(\text{execute}_j^i) \geq 1$.

We have $G = G_1 + \{\text{finalise}^k\} - G_l^k$ for certain $l \geq^{\#} k$ with $G_1(\text{execute}_l^k) = 1$. Since $G_j^i(\text{execute}_j^i) \geq 0$, we also have $G_1(\text{execute}_j^i) \geq 1$. By (J) this implies that $\neg(i \stackrel{\#}{=} k)$ or $(i, j) = (k, l)$. In the latter case $G(\text{execute}_j^i) = G_1(\text{execute}_j^i) - G_j^i(\text{execute}_j^i) = 1 - 1 = 0$, contradicting our assumption. In the former case $p \notin \bullet k$, so $G_l^k(\text{distribute}_p) = 0$ and hence $G(\text{distribute}_p) = G_1(\text{distribute}_p) \geq F'(p, i) \cdot G_1(\text{execute}_j^i) = F'(p, i) \cdot G(\text{execute}_j^i)$.

- (G) That $G(\text{execute}_j^i) \geq 0$ follows from (E) and (7.7). If $G(\text{execute}_j^i) \geq 2$ for some $i \leq^{\#} j \in T'$ then $M'(p) \geq G(\text{distribute}_p) \geq 2 \cdot F'(p, i)$ for all $p \in \bullet i$, using (D) and (F), so $M'[2 \cdot \{i\}]_{N'}$. Since N' is a finitary structural conflict net, it has no self-concurrency, so this is impossible.
- (H) Take $i \leq^{\#} j \in T'$ and $p \in \bullet j$. The case $i = j$ follows from (F), so assume $i <^{\#} j$. By (7.6) we have $G(\text{initialise}_i \cdot \text{fire}) - G(\text{initialise}_i \cdot \text{undone}) \geq 0$. So by (7.13), (E), and (7.7) $G(\text{distribute}_p) \geq 0$. Hence, using (G), we may assume, w.l.o.g., that $G(\text{execute}_j^i) = 1$. We need to investigate the same two cases as in the proof of (F) above. First let $t = \text{execute}_j^i$. Then $M_1[\text{execute}_j^i]$. Thus,

$$\begin{aligned}
& G(\text{distribute}_p) && \text{(by (7.13))} \\
& \geq F'(p, j) \cdot (G(\text{initialise}_j \cdot \text{fire}) - G(\text{initialise}_j \cdot \text{undone})) + \sum_{h \geq^{\#} g \in p^{\bullet}} F'(p, g) \cdot G(\text{fetch}_{g,h}^{p,j}) \\
& \geq F'(p, j) \cdot (G(\text{initialise}_j \cdot \text{fire}) - G(\text{initialise}_j \cdot \text{undone})) && \text{(by (E) and (7.7))} \\
& \geq F'(p, j) \cdot (G(\text{initialise}_j \cdot \text{fire}) - G(\text{initialise}_j \cdot \text{undo}(\text{trans}_j^i\text{-in}))) && \text{(by (7.6))} \\
& \geq F'(p, j) \cdot (G(\text{transfer}_j^i \cdot \text{fire}) - G(\text{transfer}_j^i \cdot \text{undone})) && \text{(by (7.9))} \\
& \geq F'(p, j) \cdot (G(\text{transfer}_j^i \cdot \text{fire}) - G(\text{transfer}_j^i \cdot \text{undo}(\text{trans}_j^i\text{-out}))) && \text{(by (7.6))} \\
& \geq F'(p, j) && \text{(by (7.12)).}
\end{aligned}$$

Now let $t = \text{finalise}^k$ with $k \in T'$. We have $G = G_1 + \{\text{finalise}^k\} - G_l^k$ for certain $l \geq^{\#} k$ with $G_1(\text{execute}_l^k) = 1$. Since $G_j^i(\text{execute}_j^i) \geq 0$, we also have $G_1(\text{execute}_j^i) \geq 1$. By (K) this implies that $\neg(j \stackrel{\#}{=} k)$ or $(i, j) = (k, l)$. In the latter case $G(\text{execute}_j^i) = G_1(\text{execute}_j^i) - G_j^i(\text{execute}_j^i) = 1 - 1 = 0$, contradicting our assumption. In the former case $p \notin \bullet k$, so $G_l^k(\text{distribute}_p) = 0$ and hence $G(\text{distribute}_p) = G_1(\text{distribute}_p) \geq F'(p, j) \cdot G_1(\text{execute}_j^i) = F'(p, j) \cdot G(\text{execute}_j^i)$.

- (I) Let $i \leq \#j \in T'$ and $h < \#j$. Since, for all $k \leq \#l \in T'$, $G_l^k(\text{transfer}_j^h \cdot \text{fire}) = \sum_{\omega} G_l^k(\text{transfer}_j^h \cdot \text{reset}_{\omega}) = 0$ and $G_l^k(\text{execute}_j^i) = G_l^k(\text{fetched}_j^i)$, the invariant is preserved when t has the form **finalise**^{*b*}. Using (7.15), it is in danger only when $t = \text{execute}_j^i$ or $t = \text{transfer}_j^h \cdot \text{reset}_{\omega}$ for some ω with $\text{transfer}_j^h \in \Omega_{\omega}$.

First assume $M_1[\text{execute}_j^i]$ and $T_j^h = G_1(\text{transfer}_j^h \cdot \text{fire}) - \sum_{\omega} G_1(\text{transfer}_j^h \cdot \text{reset}_{\omega}) = 0$. Then

$$\begin{aligned} 1 &\leq G_1(\text{transfer}_j^h \cdot \text{fire}) - G_1(\text{transfer}_j^h \cdot \text{undo}(\text{trans}_j^h \cdot \text{out})) \quad (\text{by (7.12)}) \\ &\leq G_1(\text{transfer}_j^h \cdot \text{fire}) - \sum_{\omega} G_1(\text{transfer}_j^h \cdot \text{reset}_{\omega}) = 0 \quad (\text{by (7.6)}), \end{aligned}$$

which is a contradiction.

Next assume $t = \text{transfer}_j^h \cdot \text{reset}_k$ with $k \not\equiv j$, and $E_j^i = 1$. By (E) and (G) the latter implies that $G_1(\text{execute}_j^i) = 1$ and $G_1(\text{fetched}_j^i) = 0$. Then

$$\begin{aligned} 0 &= G_1(\text{finalise}^k) \quad (\text{by (C)}) \\ &\leq G_1(\text{transfer}_j^h \cdot \text{elide}_k) + G_1(\text{transfer}_j^h \cdot \text{reset}_k) \quad (\text{by (7.4)}) \\ &< G(\text{transfer}_j^h \cdot \text{elide}_k) + G(\text{transfer}_j^h \cdot \text{reset}_k) \\ &\leq \sum_{l \geq \#k} G(\text{fetched}_l^k) \quad (\text{by (7.4)}). \end{aligned}$$

Hence $G_1(\text{fetched}_l^k) = G(\text{fetched}_l^k) > 0$ for some $l \geq \#k$, and by (7.7) also $G_1(\text{execute}_l^k) > 0$. Using (K) we obtain $(i, j) = (k, l)$, thereby obtaining a contradiction ($0 = G_1(\text{fetched}_j^i) = G_1(\text{fetched}_l^k) > 0$).

- (J) Let $j \geq \#i \not\equiv k \leq \#l \in T'$ with $(i, j) \neq (k, l)$. The invariant is in danger only when $t = \text{execute}_j^i$ or $t = \text{execute}_l^k$. W.l.o.g. let $t = \text{execute}_l^k$, with $G_1(\text{execute}_l^k) = 0$ and $G_1(\text{execute}_j^i) \geq 1$.

Making a case distinction, first assume $G(\text{fetched}_j^i) \geq 1$. Using (D), (F) and that $G(\text{execute}_l^k) = 1$, $M'(p) \geq G(\text{distribute}_p) \geq F'(p, k)$ for all $p \in \bullet k$. Likewise, $M'(p) \geq G(\text{distribute}_p) \geq F'(p, i)$ for all $p \in \bullet i$. Moreover, just as in the proof of (F), we derive, for all $p \in \bullet i \cap \bullet k$,

$$\begin{aligned} M'(p) &\geq G(\text{distribute}_p) \\ &\geq F'(p, k) \cdot (G(\text{initialise}_k \cdot \text{fire}) - G(\text{initialise}_k \cdot \text{undone})) + \sum_{h \geq \#g \in p^{\bullet}} F'(p, g) \cdot G(\text{fetch}_{g,h}^{p,k}) \\ &\geq F'(p, k) \cdot (G(\text{initialise}_k \cdot \text{fire}) - G(\text{initialise}_k \cdot \text{undone})) + \sum_{h \geq \#g \in p^{\bullet}} F'(p, g) \cdot G(\text{fetched}_h^g) \\ &\geq F'(p, k) \cdot (G(\text{initialise}_k \cdot \text{fire}) - G(\text{initialise}_k \cdot \text{undone})) + F'(p, i) \cdot G(\text{fetched}_j^i) \\ &\geq F'(p, k) \cdot (G(\text{initialise}_k \cdot \text{fire}) - G(\text{initialise}_k \cdot \text{undo}(\text{pre}_l^k))) + F'(p, i) \cdot G(\text{fetched}_j^i) \\ &\geq F'(p, k) + F'(p, i) \end{aligned}$$

by (D), (7.13), (7.7), (E), (7.6) and (7.11), respectively. It follows that $M'[\{k\} + \{i\}]$. As $i \not\equiv k$ and N' is a finitary structural conflict net, this is impossible. (Note that this argument holds regardless whether $i = k$.)

Now assume $G(\text{fetched}_j^i) \leq 0$. Then, in the notation of (7.15), $E_j^i = 1$. As $G_1(\text{execute}_l^k) = 0$, (E) and (7.7) yield $G_1(\text{fetched}_l^k) = 0$. Hence $G(\text{execute}_l^k) = 1$ and $G(\text{fetched}_l^k) = 0$, so $E_l^k = 1$. We will conclude the proof by deriving a contradiction from $E_j^i = E_l^k = 1$. In case $j = l$ this contradiction emerges immediately from (7.15). By symmetry it hence suffices to consider the case $j < l$.

By (D) and (H) we have $M'(p) \geq G(\text{distribute}_p) \geq F'(p, j)$ for all $p \in \bullet j$, so $M'[j]$. Likewise $M'[l]$ and, using (F), $M'[i]$ and $M'[k]$. Since $j \not\equiv i \not\equiv k$ and N' has no

fully reachable pure \mathbf{M} , $j \stackrel{\#}{=} k$. Since $j \stackrel{\#}{=} k \stackrel{\#}{=} l$ and N' has no fully reachable pure \mathbf{M} , $j \stackrel{\#}{=} l$. So $j <^{\#} l$. By (7.15), using that $E_j^i = 1$, $T_l^j = 0$. This is in contradiction with $E_l^k = 1$ and (I).

- (K) Suppose that $G(\text{execute}_j^i) > 0$ and $G(\text{execute}_l^k) > 0$, with $i \leq^{\#} j \stackrel{\#}{=} k \leq^{\#} l \in T'$. By (D) and (H) we have $M'(p) \geq G(\text{distribute}_p) \geq F(p, j)$ for all $p \in \bullet j$, so $M'[j]$. Likewise, using (F), $M'[i]$ and $M'[k]$. Since $i \stackrel{\#}{=} j \stackrel{\#}{=} k$ and N' has no fully reachable pure \mathbf{M} , $i \stackrel{\#}{=} k$. Using this, the result follows from (J). \blacksquare

Claim 7.4. For any $M \in [M_0]_N$ there exist $M' \in [M'_0]_{N'}$ and $G \in_F \mathbb{Z}^T$ satisfying (A)–(K) from Claim 7.3, and

- (L) there are no $j \geq^{\#} i \stackrel{\#}{=} k \leq^{\#} l \in T'$ with $M[\text{execute}_j^i]$ and $G(\text{execute}_l^k) > 0$,
(M) there are no $i \leq^{\#} j \stackrel{\#}{=} k \leq^{\#} l \in T'$ with $M[\text{execute}_j^i]$ and $G(\text{execute}_l^k) > 0$,
(N) if $M[\text{execute}_j^i]$ for $i \leq^{\#} j \in T'$ then $M'[j]$.

Proof: Given M , by Claim 7.3(2) there are M' and G so that the triple (M, M', G) satisfies (A)–(K). Assume $M[\text{execute}_j^i]$ for some $i \leq^{\#} j \in T'$. Let $M_1 := M + \llbracket \text{execute}_j^i \rrbracket$ and $G_1 := G + \{\text{execute}_j^i\}$. By (G) $G(\text{execute}_j^i) \geq 0$, so $G_1(\text{execute}_j^i) > 0$. By Claim 7.3(1) the triple (M_1, M', G_1) satisfies (A)–(K).

- (L) Suppose $G(\text{execute}_l^k) > 0$ for certain $l \geq^{\#} k \stackrel{\#}{=} i$. In case $(i, j) = (k, \ell)$, $G_1(\text{execute}_j^i) \geq 2$, contradicting (G). In case $(i, j) \neq (k, \ell)$, G_1 fails (J), also a contradiction.
(M) Suppose $G(\text{execute}_l^k) > 0$ for certain $l \geq^{\#} k \stackrel{\#}{=} j$. Then G_1 fails (G) or (K), a contradiction.
(N) By (D) and (H) $M'(p) \geq G_1(\text{distribute}_p) \geq F(p, j)$ for all $p \in \bullet j$, so $M'[j]$. \blacksquare

Claim 7.5. If $M[\{\text{execute}_j^i\} + \{\text{execute}_l^k\}]$ for some $M \in [M_0]_N$ then $\neg(i \stackrel{\#}{=} k)$.

Proof: Suppose $M[\{\text{execute}_j^i\} + \{\text{execute}_l^k\}]$ for some $M \in [M_0]_N$. By Claim 7.3(2) there exist $M' \in [M'_0]_{N'}$ and $G \in_F \mathbb{Z}^T$ satisfying (A)–(K). Let $M_1 := M + \llbracket \text{execute}_l^k \rrbracket$ and $G_1 := G + \{\text{execute}_l^k\}$. By Claim 7.3(1) the triple (M_1, M', G_1) satisfies (A)–(K). Let $M_2 := M_1 + \llbracket \text{execute}_j^i \rrbracket$ and $G_2 := G_1 + \{\text{execute}_j^i\}$. Again by Claim 7.3(1), the triple (M_2, M', G_2) also satisfies (A)–(K). As (G) implies $G(\text{execute}_j^i) \geq 0$, in case $(i, j) = (k, l)$ we obtain $G_2(\text{execute}_j^i) \geq 2$, contradicting (G). Hence $(i, j) \neq (k, l)$. Moreover, $G_2(\text{execute}_l^k) > 0$ and $G_2(\text{execute}_j^i) > 0$. Now (J) implies $\neg(i \stackrel{\#}{=} k)$. \blacksquare

For any $t \in \{\text{initialise}_j, \text{transfer}_j^h\}$ with $h, j \in T'$, and any $\omega \in \Omega$ with $t \in \Omega_\omega$, we write

$$t(\omega) := t \cdot \text{fire} + t \cdot \text{undo}_\omega + \left(\sum_{f \in t^{\text{far}}} t \cdot \text{undo}(f) \right) + t \cdot \text{undone} + t \cdot \text{reset}_\omega.$$

The transition t has no preplaces of type in , nor postplaces of type out . By checking in Table 1 or Figure 8 that each other place occurs as often in $\bullet u(\omega) + (u \cdot \text{elide}_\omega) \bullet$ as in $u(\omega) \bullet + \bullet (u \cdot \text{elide}_\omega)$, one verifies, for any $\omega \in \Omega$ with $t \in \Omega_\omega$, that

$$\llbracket t(\omega) \rrbracket = \llbracket t \cdot \text{elide}_\omega \rrbracket. \quad (7.16)$$

Let \equiv be the congruence relation on finite signed multisets of transitions generated by

$$t(\omega) \equiv t \cdot \text{elide}_\omega \quad (7.17)$$

for all $t \in \{\text{initialise}_j, \text{transfer}_j^h \mid h, j \in T'\}$ and $\omega \in \Omega$ with $\Omega_\omega \ni t$. Here *congruence* means that $G_1 \equiv G_2$ implies $k \cdot G_1 \equiv k \cdot G_2$ and $G_1 + H \equiv G_2 + H$ for all $k \in \mathbb{Z}$ and $H \in_F \mathbb{Z}^T$. Using (7.16) $G_1 \equiv G_2$ implies $\llbracket G_1 \rrbracket = \llbracket G_2 \rrbracket$.

Claim 7.6. If $M' = \llbracket G \rrbracket$ for $M' \in \mathbb{Z}^{S'}$ and $G \in_F \mathbb{Z}^T$ such that for all $i \in T'$ we have $G(\text{finalise}^i) = 0$ and either $\forall j \geq^\# i. G(\text{execute}_j^i) \geq 0$ or $\forall j \geq^\# i. G(\text{execute}_j^i) \leq 0$, then $G \equiv \emptyset$.

Proof: Let M' and G be as above. W.l.o.g. we assume $G(t \cdot \text{elide}_\omega) = 0$ for all $t \in \{\text{initialise}_j, \text{transfer}_j^h\}$ and all $\omega \in \Omega$ with $t \in \Omega_\omega$, for any G can be brought into that form by applying (7.17). For each $s \in S \setminus S'$ we have $M'(s) = 0$, and using this the inequations (7.3)–(7.7) and (7.13) of Claim 7.2 turn into equations. For each $i \in T'$ we have $G(\sum_{j \geq^\# i} \text{execute}_j^i) = 0$, using (the equational form of) (7.3)–(7.5), and that $G(\text{finalise}^i) = 0$. Since $G(\text{execute}_j^i) \geq 0$ (or ≤ 0) for all $j \geq^\# i$, this implies that $G(\text{execute}_j^i) = 0$ for each $i \leq^\# j \in T'$. With (7.7) we obtain $G(\text{fetched}_j^i) = G(\text{fetch}_{i,j}^{p,c}) = 0$ for each applicable p, c, i, j . Using that $G(t \cdot \text{elide}_\omega) = 0$ for each applicable t and ω , with (7.4)–(7.6) and (7.13) we find $G(t) = 0$ for all $t \in T$. \blacksquare

Claim 7.7. Let $M := M' + (M_0 - M'_0) + \llbracket H \rrbracket \in [M_0]_N$ for $M' \in [M'_0]_{N'}$ and $H \in_F \mathbb{Z}^T$ with $H(\text{execute}_j^i) = 0$ for all $i \leq^\# j \in T'$.

- (a) If $H(\text{finalise}^i) < 0$ and $H(\text{finalise}^k) < 0$ for certain $i, k \in T'$ then $\neg(i \# k)$.
- (b) If $M[\text{execute}_j^i]$ and $H(\text{finalise}^k) < 0$ for certain $i, k \in T'$ then $\neg(i \# k)$ and $\neg(j \# k)$.
- (c) $H(\text{distribute}_p) \geq 0$ for all $p \in S'$ (with $p^\bullet \neq \emptyset$).
- (d) Let $c \# i \in T'$. If $H(\text{distribute}_p) \geq F'(p, c)$ for all $p \in \bullet c$, then $H(\text{finalise}^i) = 0$.
- (e) If $M[\text{execute}_j^i]$ with $i \leq^\# j \in T'$ then $M'[j]$.

Proof: By Claim 7.4 there exist $M'_1 \in [M'_0]_{N'}$ and $G_1 \in_F \mathbb{Z}^T$ satisfying (B)–(N) (with M, M'_1 and G_1 playing the rôles of M, M' and G). In particular, $M = M'_1 + (M_0 - M'_0) + \llbracket G_1 \rrbracket$, $G_1(\text{finalise}^i) = 0$ for all $i \in T'$, and $G_1(\text{execute}_j^i) \geq 0$ for all $i \leq^\# j \in T'$. Using (J), for each $i \in T'$ there is at most one $j \geq^\# i$ with $G_1(\text{execute}_j^i) > 0$; we denote this j by $f(i)$, and let $f(i) := i$ when there is no such j . This makes $f : T' \rightarrow T'$ a function, satisfying $G_1(\text{execute}_j^i) = 0$ for all $j \geq^\# i$ with $j \neq f(i)$.

Given that $H(\text{execute}_j^i) = 0$ for all $i \leq^\# j \in T'$, (7.3)–(7.5) (or (7.4) and (7.7)) imply $H(\text{finalise}^i) \leq 0$ for all $i \in T'$. Let $M'_2 := M' + \sum_{i \in T'} H(\text{finalise}^i) \cdot \llbracket i \rrbracket$ and $G_2 := H - \sum_{i \in T'} H(\text{finalise}^i) \cdot G_{f(i)}^i$, where G_j^i is the right-hand side of (7.2). Then $M = M' + (M_0 - M'_0) + \llbracket H \rrbracket = M'_2 + (M_0 - M'_0) + \llbracket G_2 \rrbracket$, using that $\llbracket i \rrbracket = \llbracket G_{f(i)}^i \rrbracket$. Moreover, $G_2(\text{finalise}^i) = 0$ for all $i \in T'$, using that $G_{f(i)}^i(\text{finalise}^i) = 1$.

It follows that $M'_1 - M'_2 = \llbracket G_2 - G_1 \rrbracket$. Moreover, we have $(G_2 - G_1)(\text{finalise}^i) = 0$ for all $i \in T'$. We proceed to show that $G_2 - G_1$ satisfies the remaining precondition of Claim 7.6. So let $i \in T'$. In case $H(\text{finalise}^i) = 0$, for all $j \geq^\# i$ we have $G_2(\text{execute}_j^i) = 0$, and $G_1(\text{execute}_j^i) \geq 0$ by (G). Hence $(G_2 - G_1)(\text{execute}_j^i) \leq 0$. In case $H(\text{finalise}^i) < 0$, we have $G_2(\text{execute}_{f(i)}^i) \geq 1$, and hence, using (G), $(G_2 - G_1)(\text{execute}_{f(i)}^i) \geq 0$. Furthermore, for all $j \neq f(i)$, $G_2(\text{execute}_j^i) \geq 0$ and $G_1(\text{execute}_j^i) = 0$, so again $(G_2 - G_1)(\text{execute}_j^i) \geq 0$.

Thus we may apply Claim 7.6, which yields $G_2 \equiv G_1$. It follows that $M'_2 = M'_1 \in [M'_0]_{N'}$.

- (a) Suppose that $H(\text{finalise}^i) < 0$ and $H(\text{finalise}^k) < 0$ for certain $i \# k \in T'$. Then $G_2(\text{execute}_{f(i)}^i) > 0$ and $G_2(\text{execute}_{f(k)}^k) > 0$, so $G_1(\text{execute}_{f(i)}^i) > 0$ and $G_1(\text{execute}_{f(k)}^k) > 0$, contradicting (J).
- (b) Suppose that $M[\text{execute}_j^i]$ and $H(\text{finalise}^k) < 0$ for certain $k \# i$ or $k \# j$. Then $G_1(\text{execute}_{f(k)}^k) = G_2(\text{execute}_{f(k)}^k) > 0$, contradicting (L) or (M).
- (c) By (a), for any given $p \in S'$ there is at most one $i \in p^\bullet$ with $H(\text{finalise}^i) < 0$. For all $i \in T'$ with $i \notin p^\bullet$ we have $G_{f(i)}^i(\text{distribute}_p) = 0$. First suppose $k \in p^\bullet$ satisfies

$H(\text{finalise}^k) < 0$. Then

$$\begin{aligned} G_1(\text{execute}_{f(k)}^k) &= G_2(\text{execute}_{f(k)}^k) \\ &= H(\text{execute}_{f(k)}^k) - \sum_{i \in T'} H(\text{finalise}^i) \cdot G_{f(i)}^i(\text{execute}_{f(k)}^k) \\ &= 0 - H(\text{finalise}^k), \end{aligned}$$

so by (F) $G_1(\text{distribute}_p) \geq -F'(p, k) \cdot H(\text{finalise}^k)$. Hence

$$\begin{aligned} H(\text{distribute}_p) &= G_2(\text{distribute}_p) + \sum_{i \in T'} H(\text{finalise}^i) \cdot G_{f(i)}^i(\text{distribute}_p) \\ &= G_1(\text{distribute}_p) + H(\text{finalise}^k) \cdot G_{f(k)}^k(\text{distribute}_p) \\ &\geq -F'(p, k) \cdot H(\text{finalise}^k) + H(\text{finalise}^k) \cdot F'(p, k) = 0. \end{aligned}$$

In case there is no $i \in p^\bullet$ with $H(\text{finalise}^i) < 0$ we have

$$H(\text{distribute}_p) = G_2(\text{distribute}_p) + \sum_{i \in T'} H(\text{finalise}^i) \cdot G_{f(i)}^i(\text{distribute}_p) = G_1(\text{distribute}_p) \geq 0$$

by (F) and (G).

- (d) Since $H(\text{finalise}^i) \leq 0$ and $G_{f(i)}^i(\text{distribute}_p) \geq 0$ for all $i \in T'$, also using (c), all summands in $H(\text{distribute}_p) + \sum_{i \in T'} -H(\text{finalise}^i) \cdot G_{f(i)}^i(\text{distribute}_p)$ are positive. Now suppose $H(\text{finalise}^i) < 0$ for certain $i \in T'$. Then, using (D), for all $p \in \bullet i$,

$$M'_1(p) \geq G_1(\text{distribute}_p) = G_2(\text{distribute}_p) \geq G_{f(i)}^i(\text{distribute}_p) = F'(p, i).$$

Furthermore, let $c \stackrel{\#}{=} i$ and suppose $H(\text{distribute}_p) \geq F'(p, c)$ for all $p \in \bullet c$. Then, using (D),

$$M'_1(p) \geq G_1(\text{distribute}_p) = G_2(\text{distribute}_p) \geq H(\text{distribute}_p) \geq F'(p, c)$$

for all $p \in \bullet c$. Moreover, if $p \in \bullet c \cap \bullet i$ then

$$M'_1(p) \geq G_2(\text{distribute}_p) \geq H(\text{distribute}_p) + G_{f(i)}^i(\text{distribute}_p) \geq F'(p, c) + F'(p, i).$$

Hence $M'_2[\{c\} + \{i\}]$. However, since $c \stackrel{\#}{=} i$ and N' is a structural conflict net, this is impossible.

- (e) Suppose $M[\text{execute}_j^i]$ with $i \leq \# j \in T'$. Then $M'_1[j]$ by (N).

Now $M' = M'_1 + \sum_{k \in T'} -H(\text{finalise}^k) \cdot \llbracket k \rrbracket$, with $-H(\text{finalise}^k) \geq 0$ for all $k \in T'$. Whenever $-H(\text{finalise}^k) > 0$ then $\neg(j \stackrel{\#}{=} k)$ by (b). Hence $M'[j]$. \blacksquare

We now define the class $NF \subseteq \mathbb{Z}^T$ of signed multisets of transitions in *normal form* by $H \in NF$ iff $\ell(H) \equiv \emptyset$ and, for all $t \in \{\text{initialise}_j, \text{transfer}_j^h \mid h, j \in T'\}$:

- (NF-1) $H(t \cdot \text{elide}_\omega) \leq 0$ for each $\omega \in \Omega$,
- (NF-2) $H(t \cdot \text{undo}_\omega) \geq 0$ for each $\omega \in \Omega$, or $H(t \cdot \text{fire}) \geq 0$,
- (NF-3) and if $H(t \cdot \text{elide}_\omega) < 0$ for any $\omega \in \Omega$, then $H(t \cdot \text{undo}_\omega) \leq 0$ and $H(t \cdot \text{fire}) \leq 0$.

We proceed verifying the remaining conditions of Theorem 6.8.

4. By applying (7.17), each signed multiset $G \in_{\mathcal{F}} \mathbb{Z}^T$ with $\ell(G) \equiv \emptyset$ can be converted into a signed multiset $H \in_{\mathcal{F}} NF$ with $\ell(H) \equiv \emptyset$, such that $\llbracket H \rrbracket = \llbracket G \rrbracket$. Namely, for any $t \in \{\text{initialise}_j, \text{transfer}_j^h \mid h, j \in T'\}$, first of all perform the following three transformations, until none is applicable:

- (i) correct a positive count of a transition $t \cdot \text{elide}_\omega$ in G by adding $t(\omega) - t \cdot \text{elide}_\omega$ to G ;
- (ii) if both $H(t \cdot \text{undo}_\omega) < 0$ for some ω and $H(t \cdot \text{fire}) < 0$, correct this in the same way;
- (iii) and if, for some ω , $t \cdot \text{elide}_\omega$ has a negative and $t \cdot \text{undo}_\omega$ a positive count, add $t \cdot \text{elide}_\omega - t(\omega)$.

Note that transformation (iii) will never be applied to the same ω as (i) or (ii), so termination is ensured. Properties (NF-1) and (NF-2) then hold for t . After termination of (i)–(iii), perform

(iv) if, for some ω , $H(t \cdot \text{elide}_\omega) < 0$ and $H(t \cdot \text{fire}) > 0$, add $t \cdot \text{elide}_\omega - t(\omega)$.

This will ensure that also (NF-3) is satisfied, while preserving (NF-1) and (NF-2).

Define the function $f : T \rightarrow \mathbb{N}$ by $f(u) := 1$ for all $u \in T$ not of the form $u = t \cdot \text{elide}_\omega$, and $f(t \cdot \text{elide}_\omega) := f(t(\omega))$ (applying the last item of Definition 2.1). Then surely $f(G) = f(H)$.

5. Let $M' \in \mathbb{N}^{S'}$, $U' \in \mathbb{N}^{T'}$ and $U \in \mathbb{N}^T$ with $\ell(U) = \ell'(U')$ and $M' + \bullet U' \in [M'_0]_{N'}$. Since N' is a finitary structural conflict net, it admits no self-concurrency, so, as $\bullet U' \leq M' + \bullet U' \in [M'_0]_{N'}$, the multiset U' must be a set. As N' is plain, this implies that the multiset $\ell'(U')$ is a set. Since $\ell(U) = \ell'(U')$, also $\ell(U)$, and hence U , must be a set. All its elements have the form execute_j^i for $i \leq^\# j \in T'$, since these are the only transitions in T with visible labels. Note that U' is completely determined by U , namely by $U' = \{i \mid \exists j. \text{execute}_j^i \in U\}$. We take $H_{M',U} :=$

$$\sum_{p \in S'} (M' + \bullet U')(p) \cdot \{\text{distribute}_p\} + \sum_{(M' + \bullet U')[j]} \left(\{\text{initialise}_j \cdot \text{fire}\} + \sum_{h <^\# j, \# \text{execute}_h^g \in U} \{\text{transfer}_j^h \cdot \text{fire}\} \right)$$

Since N' is finitary, $H_{M',U} \in_F \mathbb{N}^{T+}$. Moreover, $\ell(H_{M',U}) \equiv \emptyset$.

Let $H \in_F NF$ with $M := M' + \bullet U' + (M_0 - M'_0) + \llbracket H \rrbracket - \bullet U \in \mathbb{N}^S$ and $M + \bullet U \in [M_0]_N$. Since $H \in NF$, and thus $\ell(H) \equiv \emptyset$, $H(\text{execute}_j^i) = 0$. From here on we apply Claim 7.2 and Claim 7.7 with $M + \bullet U$ and $M' + \bullet U'$ playing the rôles of M and M' . Note that the preconditions of these claims are met.

That $H(\text{execute}_j^i) = 0$ for all $i \leq^\# j \in T'$, together with (7.3) and the requirements (NF-1) and (NF-3) for normal forms, yields $H(t \cdot \text{elide}_i) \leq 0$ as well as $H(t \cdot \text{undo}_i) \leq 0$. Using this, (7.4)–(7.7) imply that

$$H(u) \leq 0 \text{ for each } u \in T_-. \quad (7.18)$$

Claim 7.8. Let $c \in T'$ and $p \in \bullet c$. Then

- if $H(\text{initialise}_c \cdot \text{fire}) > 0$ then $H(\text{fetch}_{i,j}^{p,c}) = 0$ for all $i \in p^\bullet$ and $j \geq^\# i$, and
- if $H(\text{transfer}_c^b \cdot \text{fire}) > 0$ for some $b <^\# c$ then $H(\text{fetch}_{i,j}^{p,c}) = 0$ for all $i \in p^\bullet$ and $j \geq^\# i$.

Proof: Suppose that $H(t \cdot \text{fire}) > 0$, for $t = \text{initialise}_c$ or $t = \text{transfer}_c^b$. Then (7.8) resp. (7.15) together with (7.18) implies that $H(t \cdot \text{reset}_\omega) = 0$ for each ω with $t \in \Omega_\omega$. In other words, $H(t \cdot \text{reset}_i) = 0$ for each $i \stackrel{\#}{=} c$, so in particular for each $i \in p^\bullet$. Furthermore, $H(t \cdot \text{elide}_i) \geq 0$, by requirement (NF-3) of normal forms. With (7.4), this yields $\sum_{j \geq^\# i} H(\text{fetched}_j^i) \geq 0$, and (7.18) implies $H(\text{fetched}_j^i) = 0$ for each $j \geq^\# i$. Now (7.7, 7.18) gives $H(\text{fetch}_{i,j}^{p,c}) = 0$ for each $j \geq^\# i \in p^\bullet$. ■

We proceed to verify the requirements (5a)–(5g) of Theorem 6.8.

- (5a) To show that $M_{M',U} \in \mathbb{N}^S$, it suffices to apply it to the preplaces of transitions in $H_{M',U} + U$:

$$\begin{aligned}
M_{M',U}(p) &= 0 && \text{for all } p \in S'; \\
M_{M',U}(p_j) &= \begin{cases} (M' + \bullet U')(p) - F'(p, j) & \text{if } (M' + \bullet U')[j] \\ (M' + \bullet U')(p) & \text{otherwise} \end{cases} && \text{for } p \in S', j \in p^\bullet; \\
M_{M',U}(\pi_j) &= \begin{cases} 0 & \text{if } (M' + \bullet U')[j] \\ 1 & \text{otherwise} \end{cases} && \text{for } j \in T'; \\
M_{M',U}(\text{pre}_k^j) &= \begin{cases} 1 & \text{if } (M' + \bullet U')[j] \wedge \text{execute}_k^j \notin U \\ -1 & \text{if } \neg(M' + \bullet U')[j] \wedge \text{execute}_k^j \in U \\ 0 & \text{otherwise} \end{cases} && \text{for } j \leq^\# k \in T'; \\
M_{M',U}(\pi_{h\#j}) &= \begin{cases} 0 & \text{if } \exists \text{execute}_h^g \in U \vee (M' + \bullet U')[j] \\ 1 & \text{otherwise} \end{cases} && \text{for } h <^\# j \in T'; \\
M_{M',U}(\text{trans}_j^h\text{-in}) &= \begin{cases} 1 & \text{if } (M' + \bullet U')[j] \wedge \exists \text{execute}_h^g \in U \\ 0 & \text{otherwise} \end{cases} && \text{for } h <^\# j \in T'; \\
M_{M',U}(\text{trans}_j^h\text{-out}) &= \begin{cases} 1 & \text{if } (M' + \bullet U')[j] \wedge \nexists \text{execute}_h^g \in U \wedge \nexists \text{execute}_j^i \in U \\ -1 & \text{if } (\neg(M' + \bullet U')[j] \vee \exists \text{execute}_h^g \in U) \wedge \exists \text{execute}_j^i \in U \\ 0 & \text{otherwise} \end{cases} && \text{for } h <^\# j \in T'.
\end{aligned}$$

For all these places s we indeed have that $M_{M',U}(s) \geq 0$, for the circumstances yielding the two exceptions above cannot occur:

- Suppose $\text{execute}_k^j \in U$ with $j \leq^\# k \in T'$. Then $j \in U'$, so $\bullet j \leq M' + \bullet U'$ and $(M' + \bullet U')[j]$. Consequently, $M_{M',U}(\text{pre}_k^j) \neq -1$ for all $j \leq^\# k \in T'$.
- Suppose $\text{execute}_j^i \in U$ with $i \leq^\# j \in T'$. Then $\bullet \text{execute}_j^i \leq \bullet U$, so $(M' + \bullet U)[\text{execute}_j^i]$. Claim 7.7(e) with $M + \bullet U$ and $M' + \bullet U'$ in the rôles of M and M' yields $(M' + \bullet U')[j]$.

If moreover $\text{execute}_h^g \in U$ with $g \leq^\# h <^\# j$, then $\{g\} + \{i\} \leq U'$, so $\bullet \{g\} + \bullet \{i\} \leq M' + \bullet U'$ and $(M' + \bullet U')[\{g\} + \{i\}]$. In particular, $g \smile i$, and since N' is a structural conflict net, $\bullet g \cap \bullet i = \emptyset$. By Claim 7.7(e)—as above— $(M' + \bullet U')[h]$, so $\bullet g \cup \bullet h \cup \bullet j \cup \bullet i \leq M' + \bullet U' \in [M'_0]_{N'}$. Moreover, since $g \leq^\# h <^\# j \geq^\# i$, we have $\bullet g \cap \bullet h \neq \emptyset$, $\bullet h \cap \bullet i \neq \emptyset$ and $\bullet i \cap \bullet j \neq \emptyset$. Now in case also $\bullet h \cap \bullet i \neq \emptyset$, the transitions g, h and i constitute a fully reachable pure \mathbf{M} ; otherwise $h \smile i$ and h, j and i constitute a fully reachable pure \mathbf{M} . Either way, we obtain a contradiction.

Consequently, $M_{M',U}(\text{trans}_j^h\text{-out}) \neq -1$ for all $h <^\# j \in T'$.

- (5b) Suppose $M' \xrightarrow{a}$; say $M'[i]$ with $\ell'(i) = a$. Let j be the largest transition in T' w.r.t. the well-ordering $<$ on T such that $i \leq^\# j$ and $(M' + \bullet U')[j]$. It suffices to show that $M_{M',U}[\text{execute}_j^i]$, i.e. that $M_{M',U}(\text{pre}_j^i) = 1$, $M_{M',U}(\text{trans}_j^h\text{-out}) = 1$ for all $h <^\# j$, and $M_{M',U}(\pi_{j\#l}) = 1$ for all $l >^\# j$.

If $\text{execute}_j^i \in U$ we would have $i \in U'$ and hence $(M' + \bullet U')[2 \cdot \{i\}]$. Since N' is a finitary structural conflict net, this is impossible. Therefore $\text{execute}_j^i \notin U$ and, using the calculations from (a) above, $M_{M',U}(\text{pre}_j^i) = 1$.

Let $h <^\# j$. To establish that $M_{M',U}(\text{trans}_j^h\text{-out}) = 1$ we need to show that there is no $k \leq^\# j$ with $\text{execute}_k^j \in U$ and no $g \leq^\# h$ with $\text{execute}_h^g \in U$. First suppose $\text{execute}_k^j \in U$ for some $k \leq^\# j$. Then $k \in U'$ and hence $(M' + \bullet U')[\{i\} + \{k\}]$. This implies $i \smile k$, and, as N' is a structural conflict net, $\bullet i \cap \bullet k = \emptyset$. Hence the transitions i, j and k are all different, with $\bullet i \cap \bullet j \neq \emptyset$ and $\bullet j \cap \bullet k \neq \emptyset$ but $\bullet i \cap \bullet k = \emptyset$. Moreover, the reachable marking $M' + \bullet U'$ enables all three of them. Hence N' contains a fully reachable pure \mathbf{M} , which contradicts the assumptions of Theorem 7.1.

Next suppose $\text{execute}_h^g \in U$ for some $g \leq^\# h$. Then $(M + \bullet U)[\text{execute}_h^g]$, so $(M' + \bullet U')[h]$ by Claim 7.7(e). Moreover, $g \in U'$, so $(M' + \bullet U')[\{i\} + \{g\}]$. This implies $g \smile i$, and $\bullet g \cap \bullet i = \emptyset$. Moreover, $\bullet g \cap \bullet h \neq \emptyset$, $\bullet h \cap \bullet j \neq \emptyset$ and $\bullet j \cap \bullet i \neq \emptyset$, while the reachable marking $M' + \bullet U'$ enables all these transitions. Depending on whether $\bullet h \cap \bullet i = \emptyset$, either h, j and i , or g, h and i constitute a fully reachable pure \mathbf{M} , contradicting the assumptions of Theorem 7.1.

Let $l >^\# j$. To establish that $M_{M',U}(\pi_{j\#l}) = 1$ we need to show that there is no $k \leq^\# j$ with $\text{execute}_j^k \in U$ —already done above—and that $\neg(M' + \bullet U')[l]$. Suppose $(M' + \bullet U')[l]$. Considering that j was the largest transition with $i \leq^\# j$ and $(M' + \bullet U')[j]$, we cannot have $i <^\# l$. Hence the transitions i, j and l are all different, with $\bullet i \cap \bullet j \neq \emptyset$ and $\bullet j \cap \bullet l \neq \emptyset$ but $\bullet i \cap \bullet l = \emptyset$. Moreover, the reachable marking $M' + \bullet U'$ enables all three of them. Hence N' contains a fully reachable pure \mathbf{M} , which contradicts the assumptions of Theorem 7.1.

(5c) We have to show that $H(t) \leq H_{M',U}(t)$ for each $t \in T$.

- In case $t \in T_-$ this follows from (7.18) and $H_{M',U} \in \mathbb{N}^{T_+}$.
- In case $t = \text{execute}_j^i$ it follows since $\ell(H) \equiv \emptyset$.
- In case $t = \text{distribute}_p$ it follows from (7.14) and (7.18).
- Next let $t = \text{initialise}_c \cdot \text{fire}$ for some $c \in T'$. In case $H(\text{initialise}_c \cdot \text{fire}) \leq 0$ surely we have $H(\text{initialise}_c \cdot \text{fire}) \leq H_{M',U}(\text{initialise}_c \cdot \text{fire})$. So without limitation of generality we may assume that $H(\text{initialise}_c \cdot \text{fire}) > 0$. By (7.8, 7.18) we have $H(\text{initialise}_c \cdot \text{fire}) = 1$. Using (7.13), Claim 7.8, (7.18) and (7.14) we obtain, for all $p \in \bullet c$,

$$F'(p, c) \cdot H(\text{initialise}_c \cdot \text{fire}) \leq H(\text{distribute}_p) \leq (M' + \bullet U')(p).$$

Hence c is enabled under $M' + \bullet U'$, which implies $H_{M',U}(\text{initialise}_c \cdot \text{fire}) = 1$.

- Let $t = \text{transfer}_c^b \cdot \text{fire}$ for some $b <^\# c \in T'$. As above, we may assume $H(\text{transfer}_c^b \cdot \text{fire}) > 0$. By (7.15, 7.18) we have $H(\text{transfer}_c^b \cdot \text{fire}) = 1$. Using (7.18) and that $H(\text{execute}_b^g) = 0$ for all $g \leq^\# b$, it follows that $(M + \bullet U)(\pi_{b\#c}) = 0$. Hence $\neg(M + \bullet U)[\text{execute}_b^g]$ for all $g \leq^\# b$, and thus $\nexists \text{execute}_b^g \in U$. For all $p \in \bullet c$ we derive

$$\begin{aligned} F'(p, c) \cdot H(\text{transfer}_c^b \cdot \text{fire}) & \\ & \leq F'(p, c) \cdot (H(\text{transfer}_c^b \cdot \text{fire}) - H(\text{transfer}_c^b \cdot \text{undone})) \end{aligned} \quad (7.18)$$

$$\leq F'(p, c) \cdot (H(\text{initialise}_c \cdot \text{fire}) - H(\text{initialise}_c \cdot \text{undo}(\text{trans}_c^b \cdot \text{in}))) \quad (7.9)$$

$$\leq F'(p, c) \cdot (H(\text{initialise}_c \cdot \text{fire}) - H(\text{initialise}_c \cdot \text{undone})) \quad (7.6)$$

$$= [\text{the same as above}] + \sum_{j \geq^\# i \in p^\bullet} F'(p, i) \cdot H(\text{fetch}_{i,j}^{p,c}) \quad (\text{Claim 7.8})$$

$$\leq H(\text{distribute}_p) \quad (7.13)$$

$$\leq (M' + \bullet U')(p) + \sum_{i \in T' \setminus p \in i^\bullet} H(\text{finalise}_i) \quad (7.14)$$

$$\leq (M' + \bullet U')(p) \quad \{i \in T' \setminus p \in i^\bullet\} \quad (7.18).$$

Hence $(M' + \bullet U')[c]$, and thus $H_{M',U}(\text{transfer}_c^b) = 1$.

(5d) If $u \notin T_-$, yet $H(u) \neq 0$, then u is either distribute_p , $\text{initialise}_j \cdot \text{fire}$ or $\text{transfer}_j^h \cdot \text{fire}$ for suitable $p \in S'$ or $h, j \in T'$. For $u = \text{distribute}_p$ the requirement follows from Claim 7.7(c); otherwise Property (NF-2), together with (7.6), guarantees that $H(u) \geq 0$.

(5e) If $H(t) > 0$ and $H(u) < 0$, then $t \in T_+$ and $u \in T_-$. The only candidates for $\bullet t \cap \bullet u \neq \emptyset$ are

- $p_c \in \bullet(\text{initialise}_c \cdot \text{fire}) \cap \bullet(\text{fetch}_{i,j}^{p,c})$ for $p \in S'$, $c, i \in p^\bullet$ and $j \geq^\# i$,
- $\text{trans}_c^b \cdot \text{in} \in \bullet(\text{transfer}_c^b \cdot \text{fire}) \cap \bullet(\text{initialise}_c \cdot \text{undo}(\text{trans}_c^b \cdot \text{in}))$ for $b \leq^\# c \in T'$.

We investigate these possibilities one by one.

- $H(\text{initialise}_c \cdot \text{fire}) > 0 \wedge H(\text{fetch}_{i,j}^{p,c}) < 0$ cannot occur by Claim 7.8.
 - Suppose $H(\text{transfer}_c^b \cdot \text{fire}) > 0$. By (7.15, 7.18) we have $H(\text{transfer}_c^b \cdot \text{fire}) = 1$. Through the derivation above, in the proof of requirement (c), using (7.18, 7.9, 7.6), Claim 7.8 and (7.13), we obtain $H(\text{distribute}_p) \geq F'(p, c)$ for all $p \in \bullet c$. Now Claim 7.7(d) yields $H(\text{finalise}^i) = 0$ for all $i \not\equiv c$. By (7.4) and (7.18) we obtain $H(\text{initialise}_c \cdot \text{reset}_i) = 0$ for each such i . Hence $\sum_{i \not\equiv c} H(\text{initialise}_c \cdot \text{reset}_i) = 0$, and thus $H(\text{initialise}_c \cdot \text{undo}(\text{trans}_c^b\text{-in})) = 0$ by (7.6, 7.18).
- (5f) If $H(u) < 0$ and $(M + \bullet U)[t]$ with $\ell(t) \neq \tau$, then $t = \text{execute}_j^i$ for some $i \leq\# j \in T'$ and $u \in T_-$. The only candidates for $\bullet t \cap \bullet u \neq \emptyset$ are
- $\text{pre}_j^i \in \bullet(\text{execute}_j^i) \cap \bullet(\text{initialise}_j \cdot \text{undo}(\text{pre}_j^i))$ and
 - $\text{trans}_j^h\text{-out} \in \bullet(\text{execute}_j^i) \cap \bullet(\text{transfer}_j^h \cdot \text{undo}(\text{trans}_j^h\text{-out}))$ for $h <\# j$.

We investigate these possibilities one by one.

- Suppose $(M + \bullet U)[\text{execute}_j^i]$. By Claim 7.7(b), $H(\text{finalise}^k) \geq 0$ for each $k \equiv i$. By (7.4) and (7.18) we obtain $H(\text{initialise}_i \cdot \text{reset}_k) = 0$ for each such k . Hence $\sum_{k \equiv i} H(\text{initialise}_i \cdot \text{reset}_k) = 0$, and thus $H(\text{initialise}_i \cdot \text{undo}(\text{pre}_j^i)) = 0$ by (7.6, 7.18).
 - Suppose $(M + \bullet U)[\text{execute}_j^i]$ and $h <\# j$. By Claim 7.7(b), $H(\text{finalise}^k) \geq 0$ for each $k \equiv j$. By (7.4) and (7.18) $H(\text{transfer}_j^h \cdot \text{reset}_k) = 0$ for each such k . So $\sum_{k \equiv j} H(\text{transfer}_j^h \cdot \text{reset}_k) = 0$, and $H(\text{transfer}_j^h \cdot \text{undo}(\text{trans}_j^h\text{-out})) = 0$ by (7.6, 7.18).
- (5g) Suppose $(M + \bullet U)[\{t\} + \{u\}]_N$, and $i, k \in T'$ with $\ell'(i) = \ell(t)$ and $\ell'(k) = \ell(u)$. Since the net N' is plain, t and u must have the form execute_j^i and execute_j^k for some $j >\# i$ and $l >\# k$. Claim 7.5 yields $\neg(i \equiv k)$ and hence $\bullet i \cap \bullet k = \emptyset$. \square

Thus, we have established that the conflict replicating implementation $\mathcal{I}(N')$ of a finitary plain structural conflict net N' without a fully reachable pure \mathbf{M} is branching ST-bisimilar with explicit divergence to N' . It remains to be shown that $\mathcal{I}(N')$ is essentially distributed.

Lemma 7.9. *Let N be the conflict replicating implementation of a finitary net*

$N' = (S', T', F', M'_0, \ell')$; *let $j, l \in T'$, with $l >\# j$. Then no two transitions from the set*

$$\{\text{execute}_j^i \mid i \leq\# j\} \cup \{\text{transfer}_l^j \cdot \text{fire}\} \cup \{\text{transfer}_l^j \cdot \text{undo}(\text{trans}_l^j\text{-out})\} \cup \{\text{execute}_l^k \mid k \leq\# l\}$$

can fire concurrently.

Proof. For each $i \leq\# j$ pick an arbitrary preplace q_i of i . The set

$$\{\text{fetch}_{i,j}^{q_i,i}\text{-in}, \text{fetch}_{i,j}^{q_i,i}\text{-out} \mid i \leq\# j\} \cup \{\pi_{j\#l}, \text{trans}_l^j\text{-out}, \text{took}(\text{trans}_l^j\text{-out}, \text{transfer}_l^j), \rho(\text{transfer}_l^j)\}$$

is an S -invariant: there is always exactly one token in this set. This is the case because there is exactly one token initially (on $\pi_{j\#l}$) and each transition from N has as many (with multiplicities) preplaces as postplaces in this set. The transitions from

$$\{\text{execute}_j^i \mid i \leq\# j\} \cup \{\text{transfer}_l^j \cdot \text{fire}\} \cup \{\text{transfer}_l^j \cdot \text{undo}(\text{trans}_l^j\text{-out})\} \cup \{\text{execute}_l^k \mid k \leq\# l\}$$

each have a preplace in this set. Hence no two of them can fire concurrently. \square

Lemma 7.10. *Let N be the conflict replicating implementation $\mathcal{I}(N')$ of a finitary plain structural conflict net $N' = (S', T', F', M'_0, \ell')$ without a fully reachable pure \mathbf{M} . Then for any $i \leq\# j \equiv c \in T'$ and $f \in (\text{initialise}_c)^{\text{far}}$, the transitions execute_j^i and $\text{initialise}_c \cdot \text{undo}(f)$ cannot fire concurrently.*

Proof. Suppose these transitions can fire concurrently, say from the marking $M \in [M_0]_N$. By Claim 7.4, there are $M' \in [M'_0]_{N'}$ and $G \in_F \mathbb{Z}^T$ such that (B)–(N) hold. Let $t := \text{initialise}_c$, $G_1 := G + \{t \cdot \text{undo}(f)\}$ and $M_1 := M + \llbracket t \cdot \text{undo}(f) \rrbracket$. Then (7.6), applied to the triples (M, M', G) and (M_1, M', G_1) , yields

$$\sum_{\{\omega | t \in \Omega_\omega\}} G(t \cdot \text{reset}_\omega) \leq G(t \cdot \text{undo}(f)) < G_1(t \cdot \text{undo}(f)) \leq \sum_{\{\omega | t \in \Omega_\omega\}} G_1(t \cdot \text{undo}_\omega) = \sum_{\{\omega | t \in \Omega_\omega\}} G(t \cdot \text{undo}_\omega).$$

Hence, there is an ω with $t \in \Omega_\omega$ and $G(t \cdot \text{reset}_\omega) < G(t \cdot \text{undo}_\omega)$. This ω must have the form $k \in T'$ with $k \stackrel{\#}{=} c$. We now obtain

$$\begin{aligned} 0 &= G(\text{finalise}^k) && \text{(by (C))} \\ &\leq G(t \cdot \text{elide}_k) + G(t \cdot \text{reset}_k) && \text{(by (7.4))} \\ &< G(t \cdot \text{elide}_k) + G(t \cdot \text{undo}_k) \\ &\leq \sum_{l \geq \#k} G(\text{execute}_l^k) && \text{(by (7.3)).} \end{aligned}$$

Hence, there is an $l \geq \#k \stackrel{\#}{=} c$ with $G(\text{execute}_l^k) > 0$. By (M) we obtain $\neg(j \stackrel{\#}{=} k)$, so $\bullet j \cap \bullet k = \emptyset$. Additionally, we have $\bullet j \cap \bullet c \neq \emptyset$ and $\bullet c \cap \bullet k \neq \emptyset$. By (N) we obtain $M'[j]$, and by (D) and (F) $M'[k]$. Furthermore, by (7.6), $G(t \cdot \text{undo}(f)) < G_1(t \cdot \text{undo}(f)) \leq G_1(t \cdot \text{fire}) = G(t \cdot \text{fire})$, so, for all $p \in \bullet c$,

$$\begin{aligned} F'(p, c) &\leq F'(p, c) \cdot (G(t \cdot \text{fire}) - G(t \cdot \text{undo}(f))) \\ &\leq F'(p, c) \cdot (G(t \cdot \text{fire}) - G(t \cdot \text{undone})) && \text{(by (7.6))} \\ &\leq G(\text{distribute}_p) - \sum_{j \geq \#i \in p} F'(p, i) \cdot G(\text{fetch}_{i,j}^{p,c}) && \text{(by (7.13))} \\ &\leq G(\text{distribute}_p) && \text{(by (E) and (7.7))} \\ &\leq M'(p) && \text{(by (D)).} \end{aligned}$$

It follows that $M'[c]$. Thus N' contains a fully reachable pure \mathbf{M} , which contradicts the assumptions of Lemma 7.10. \square

Theorem 7.11. *Let N be the conflict replicating implementation $\mathcal{I}(N')$ of a finitary plain structural conflict net N' without a fully reachable pure \mathbf{M} . Then N is essentially distributed.*

Proof. We take the canonical distribution D of N , in which \equiv_D is the equivalence relation on places and transitions generated by Condition (1) of Definition 4.5. We need to show that this distribution satisfies Condition (2') of Definition 4.8. A given transition t with $\ell(t) \neq \tau$ must have the form execute_j^i for some $i \leq \#j \in T'$. By following the flow relation of N one finds the places and transitions that, under the canonical distribution, are co-located with execute_j^i :

$$\begin{array}{c} \pi_{j \neq l} \rightarrow \text{transfer}_l^j \cdot \text{fire} \leftarrow \text{trans}_l^j \text{-in} \rightarrow \text{initialise}_l \cdot \text{undo}(\text{trans}_l^j \text{-in}) \leftarrow \text{take}(\text{trans}_l^j \text{-in}, \text{initialise}_l) \\ \downarrow \\ \text{execute}_j^i \\ \uparrow \\ \text{trans}_j^h \text{-out} \rightarrow \text{transfer}_j^h \cdot \text{undo}(\text{trans}_j^h \text{-out}) \leftarrow \text{take}(\text{trans}_j^h \text{-out}, \text{transfer}_j^h) \\ \downarrow \\ \text{execute}_j^g \\ \uparrow \\ \text{pre}_j^g \rightarrow \text{initialise}_g \cdot \text{undo}(\text{pre}_j^g) \leftarrow \text{take}(\text{pre}_j^g, \text{initialise}_g) \end{array}$$

for all $l > \#j$, $h < \#j$ and $g \leq \#j$. We need to show that none of these transitions can happen concurrently with execute_j^i . For transitions $\text{transfer}_l^j \cdot \text{fire}$ and execute_j^g this follows directly

from Lemma 7.9. For $\text{transfer}_j^h \cdot \text{undo}(\text{trans}_j^h\text{-out})$ this also follows from Lemma 7.9, in which j, k and l play the rôle of the current h, i and j . For the transitions $\text{initialise}_l \cdot \text{undo}(\text{trans}_l^j\text{-in})$ and $\text{initialise}_g \cdot \text{undo}(\text{pre}_j^g)$ this has been established in Lemma 7.10. \square

Our main result follows by combining Theorems 7.1, 7.11 and 4.15:

Theorem 7.12. *Let N be a finitary plain structural conflict net without a fully reachable pure \mathbf{M} . Then N is distributable up to \approx_{bSTb}^Δ .* \square

Corollary 7.13. *Let N be a finitary plain structural conflict net. Then N is distributable iff it has no fully reachable pure \mathbf{M} .* \square

8. CONCLUSION

In this paper, we have given a precise characterisation of distributable Petri nets in terms of a semi-structural property. Moreover, we have shown that our notion of distributability corresponds to an intuitive notion of a distributed system by establishing that any distributable net may be implemented as a network of asynchronously communicating components.

In order to formalise what qualifies as a valid implementation, we needed a suitable equivalence relation. We have chosen step failures equivalence for showing the impossibility part of our characterisation, since it is one of the simplest and least discriminating semantic equivalences imaginable that abstracts from internal actions but preserves branching time, concurrency and divergence to some small degree. For the positive part, stating that all other nets are implementable, we have introduced a combination of several well known rather discriminating equivalences, namely a divergence sensitive version of branching bisimulation adapted to ST-semantics. Hence our characterisation is rather robust against the chosen equivalence; it holds in fact for all equivalences between these two notions. However, ST-equivalence (and our version of it) preserves the causal structure between action occurrences only as far as it can be expressed in terms of the possibility of durational actions to overlap in time. Hence a natural question is whether we could have chosen an even stronger causality sensitive equivalence for our implementability result, respecting e.g. pomset equivalence or history preserving bisimulation. Our conflict replicating implementation does not fully preserve the causal behaviour of nets; we are convinced that we have chosen the strongest possible equivalence for which our implementation works. It is an open problem to find a class of nets that can be implemented distributedly while preserving divergence, branching time and causality in full. Another line of research is to investigate which Petri nets can be implemented as distributed nets when relaxing the requirement of preserving the branching structure. We conjecture that there exists a notion of equivalence that captures some branching time aspects, but not as strongly as step failures equivalence, under which all Petri nets become distributable. However, also in this case it is problematic, in fact even impossible in our setting, to preserve the causal structure, as has been shown in [SPG11]. A similar impossibility result has been obtained in the world of the π -calculus in [PSN11].

In this paper we have sought a characterisation of distributability only for *plain* nets, in which all transitions have a different label and none are internal. Naturally, any distributed implementation that applies to plain nets having a semi-structural property—in particular the one contributed here—also applies to non-plain nets having the same semi-structural property. Namely to implement a non-plain net N , note that N can be written as $\rho(N')$, where N' is a plain net and ρ a relabelling function. A correct implementation of N is

now obtained as $\rho(\mathcal{I}(N'))$, where $\mathcal{I}(N')$ is the distributed implementation of N' . Yet, it appears unlikely that there is a semi-structural characterisation that captures *all* non-plain distributable nets: for any non-trivial semi-structural property there probably are nets that do not have that property, but are semantically equivalent to nets that do. This may happen for instance when some essential transitions that violate the property are labelled τ and can be abstracted away. Thus, we do not expect that a natural characterisation of distributability for non-plain nets exists—where “natural” excludes characterisations that just say, in other words, “being equivalent to a distributed net”.

Our work shows that the main problem in creating distributed implementations of systems arises from the interplay between choice and synchronous communication. This issue has already been investigated in the context of distributed algorithms. Rabin and Lehmann observed in [RL94] that there is no fully symmetric distributed solution to the dining philosophers problem. In [Bou88] Luc Bougé considers the problem of implementing symmetric leader election in the sublanguages of CSP obtained by allowing different forms of communication, combining input and output guards in guarded choice in different ways. He finds that the possibility of implementing leader election depends heavily on the structure of the communication graphs. Truly symmetric schemes are only possible in CSP with arbitrary input and output guards in choices.

Synchronous interaction is a basic concept in many languages for system specification and design, e.g. in statechart-based approaches and in process calculi. For process calculi, language hierarchies have been established which exhibit the expressive power of different forms of synchronous and asynchronous interaction. In [BP91] Frank de Boer and Catuscia Palamidessi consider various dialects of CSP with differing degrees of asynchrony. Similar work is done for the π -calculus in [Pal97] by Catuscia Palamidessi, in [Nes00] by Uwe Nestmann and in [Gor06] by Daniele Gorla. A rich hierarchy of asynchronous π -calculi has been mapped out in these papers. Similar to the findings of Bougé, mixed-choice, i.e. the ability to combine input and output guards in a single choice, plays a central rôle in the implementation of synchronous behaviour.

In [Sel97], 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 focuses on the asynchronous nature of the communications of a system with the outside world.

Dirk Taubner has in [Tau88] given various protocols by which to implement arbitrary Petri nets in the OCCAM programming language. Although this programming language offers synchronous communication he makes no substantial use of that feature in the protocols, thereby effectively providing an asynchronous implementation of Petri nets. He does not indicate a specific equivalence relation, but is effectively using linear-time equivalences to compare implementations to the specification.

Also in hardware design it is an intriguing quest to use interaction mechanisms which do not rely on a global clock, in order to gain performance. Here the simulation of synchrony by asynchrony can be a crucial issue, see for instance [Lam78] and [Lam03].

The idea of modelling asynchronously communicating sequential components by sequential Petri nets interacting through buffer places has already been considered in [Re82]. There Wolfgang Reisig introduces a class of systems, represented as Petri nets, where the




Figure 13: A specification and its Hopkins-implementation which added concurrency.



Figure 14: A distributable net which is not considered distributable in [Ho91], and its implementation.

relative speeds of different components are guaranteed to be irrelevant. His class is a strict subset of our LSGA nets, requiring additionally, amongst others, that all choices in sequential components are free, i.e. do not depend upon the existence of buffer tokens, and that places are output buffers of only one component. Another quite similar approach was taken in [EHH10], where transition labels are classified as being either input or output. There, asynchrony is introduced by adding new buffer places during net composition. This framework does not allow multiple senders for a single receiver.

Other notions of distributed and distributable Petri nets are proposed in [Ho91, BCD02, BD12]. In these works, given a distribution of the transitions of a net, the net is distributable iff it can be implemented by a net that is distributed w.r.t. that distribution. The requirement that concurrent transitions may not be co-located is absent; given the fixed distribution, there is no need for such a requirement. These papers differ from each other, and from ours, in what counts as a valid implementation. Hopkins [Ho91] uses an interleaving equivalence to compare an implementation to the original net, and while allowing a range of implementations, he does require them to inherit some of the structure of the original net. The net classes he describes in his paper are incomparable with our class of distributable nets. One direction of this inequality depends on his choice of interleaving semantics, which allows the implementation in Figure 13. The step failures equivalence we use does not tolerate the added concurrency and the depicted net is not distributable in our sense. The other direction of the inequality stems from the fact that we allow implementations which do not share structure with the specification but only emulate its behaviour. That way, the net in Figure 14 can be implemented in our approach as depicted.

A more abstract approach to the same underlying problem of correctly executing an arbitrary Petri net as a distributed system has been taken in [KP13]. The authors provide a modified net semantics and an algorithm to split the net into agents which can locally decide most choices and resort to a global scheduler in case multiple agents must be coordinated. While such an approach loses branching time equivalence between a net and its implementation, it provides a clear separation of concerns between executing the net and solving the distributed coordination problems.

In [GGS08] we have obtained a characterisation similar to Corollary 7.13, but for a much more restricted notion of distributed implementation (*plain distributability*), disallowing nontrivial transition labellings in distributed implementations. We also proved that fully reachable pure Ms are not implementable in a distributed way, even when using transition labels (Theorem 5.6). However, we were not able to show that this upper bound on the class of distributable systems was tight. Our current work implies the validity of Conjecture 1 of [GGS08]. While in [GGS08] we considered only one-safe place/transition systems, the present paper employs a more general class of place/transition systems, namely structural conflict nets. This enables us to give a concrete characterisation of distributed nets as systems of sequential components interacting via non-safe buffer places.

On the level of applications, we expect our results to be useful for language design. We would like to make a thorough comparison of our results to those on communication patterns in process algebras, versions of the π -calculus and I/O-automata [Lyn96]. Using a Petri net semantics of a suitable system description language, we could compare our class of distributed nets to the class of nets expressible in the language, especially when restricting the allowed communication patterns in the ways considered in [BP91, Bou88] or in [Lyn96]. A first step in that direction is [PNG13].

Acknowledgment. The authors gratefully thank the referees of this paper for their very thorough examination and helpful suggestions.

REFERENCES

- [BCD02] E. Badouel, B. Caillaud & P. Darondeau (2002): *Distributing Finite Automata Through Petri Net Synthesis*. *Formal Aspects of Computing* 13(6), pp. 447–470, doi:10.1007/s001650200022.
- [BKO87] J.A. Bergstra, J.W. Klop & E.-R. Olderog (1987): *Failures without chaos: a new process semantics for fair abstraction*. In M. Wirsing, editor: *Formal Description of Programming Concepts – III, Proceedings of the 3th IFIP WG 2.2 working conference*, Ebberup 1986, Amsterdam, pp. 77–103.
- [BD12] E. Best & Ph. Darondeau (2012): *Petri Net Distributability*. In E.M. Clarke, I. Virbitskaite & A. Voronkov, editors: *Perspectives of Systems Informatics - Revised Selected Papers presented at the 8th International Andrei Ershov Memorial Conference*, PSI 2011, Novosibirsk, LNCS 7162, Springer, pp. 1–18, doi:10.1007/978-3-642-29709-0_1.
- [BP91] 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: *Proc. 2nd International Conference on Concurrency Theory (CONCUR'91)*, Amsterdam, The Netherlands, LNCS 527, Springer, pp. 127–141, doi:10.1007/3-540-54430-5_85.
- [Bou88] L. Bougé (1988): *On the existence of symmetric algorithms to find leaders in networks of communicating sequential processes*. *Acta Inf.* 25(2), pp. 179–201, doi:10.1007/BF00263584.
- [BHR84] 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, doi:10.1145/828.833.
- [EHH10] D. El Hog-Benzina, S. Haddad & R. Hennicker (2010): *Process Refinement and Asynchronous Composition with Modalities*. In N. Sidorova & A. Serebrenik, editors: *Proceedings of the 2nd Intern. Workshop on Abstractions for Petri Nets and Other Models of Concurrency (APNOC'10)*, Braga, Portugal. Available at <http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/EHH-apnoc10.pdf>.
- [vG93] R.J. van Glabbeek (1993): *The Linear Time - Branching Time Spectrum II*. In: *Proceedings of the 4th International Conference on Concurrency Theory (CONCUR'93)*, Springer, London, UK, pp. 66–81, doi:10.1007/3-540-57208-2_6.
- [GG01] R.J. van Glabbeek & U. Goltz (2001): *Refinement of actions and equivalence notions for concurrent systems*. *Acta Informatica* 37(4/5), pp. 229–327, doi:10.1007/s002360000041.

- [GGS08] R.J. van Glabbeek, U. Goltz & J.-W. Schicke (2008): *On Synchronous and Asynchronous Interaction in Distributed Systems*. In E. Ochmański & J. Tyszkiewicz, editors: *Mathematical Foundations of Computer Science 2008*, LNCS 5162, Springer, pp. 16–35, doi:10.1007/978-3-540-85238-4_2. Full version: Technical Report 2008-03, TU-Braunschweig; <http://arxiv.org/abs/0901.0048>.
- [GGS11] R.J. van Glabbeek, U. Goltz & J.-W. Schicke (2011): *Abstract Processes of Place/Transition Systems*. *Information Processing Letters* 111(13), pp. 626 – 633, doi:10.1016/j.ip1.2011.03.013.
- [GGS12] R.J. van Glabbeek, U. Goltz & J.-W. Schicke-Uffmann (2012): *On Distributability of Petri Nets*. Informatik Bericht Nr. 2011-10, Institut für Programmierung und Reaktive Systeme, TU Braunschweig, Germany. Available at <http://arxiv.org/abs/1207.3597>. Ext. abstract in L. Birkedal, ed.: *Proc. 15th Int. Conf. on Foundations of Software Science and Computational Structures*, FoSSaCS'12, LNCS 7213, Springer, 2012, pp. 331-345, doi:10.1007/978-3-642-28729-9_22.
- [GLT09] R.J. van Glabbeek, B. Luttik & N. Trčka (2009): *Branching Bisimilarity with Explicit Divergence*. *Fundamenta Informaticae* 93(4), pp. 371–392. Archived at <http://arxiv.org/abs/0812.3068>.
- [GV87] R.J. van Glabbeek & F.W. Vaandrager (1987): *Petri net models for algebraic theories of concurrency (extended abstract)*. In: *Proceedings PARLE '87*, LNCS 259, Springer, pp. 224–242, doi:10.1007/3-540-17945-3_13. Available at <http://kilby.stanford.edu/~rvg/pub/petri.pdf>.
- [GW89] R.J. van Glabbeek & W.P. Weijland (1989): *Branching Time and Abstraction in Bisimulation Semantics (extended abstract)*. In G.X. Ritter, editor: *Information Processing 89*, Proceedings of the IFIP 11th World Computer Congress, San Francisco 1989, North-Holland, pp. 613–618. Full version appeared as [GW89].
- [GW96] R.J. van Glabbeek & W.P. Weijland (1996): *Branching Time and Abstraction in Bisimulation Semantics*. *Journal of the ACM* 43(3), pp. 555–600, doi:10.1145/233551.233556.
- [Gor06] D. Gorla (2006): *On the Relative Expressive Power of Asynchronous Communication Primitives*. In L. Aceto & A. Ingólfssdóttir, eds.: *Proc. 9th Int. Conf. on Foundations of Software Sc. and Comput. Structures (FoSSaCS'06)*, LNCS 3921, Springer, pp. 47–62, doi:10.1007/11690634_4.
- [Ho91] R.P. Hopkins (1991): *Distributable nets*. In: *Advances in Petri Nets 1991*, LNCS 524, Springer, pp. 161–187, doi:10.1007/BFb0019974.
- [KP13] J.-P. Katoen & D. Peled (2013): *Taming Confusion for Modeling and Implementing Probabilistic Concurrent Systems*. In M. Felleisen & P. Gardner, editors: *Programming Languages and Systems - Proceedings 22nd European Symposium on Programming, ESOP 2013, Rome, Italy, March 2013*, LNCS 7792, Springer, pp. 411-430, doi:10.1007/978-3-642-37036-6_23.
- [Lam78] L. Lamport (1978): *Time, Clocks, and the Ordering of Events in a Distributed System*. *Commun. ACM* 21(7), pp. 558–565, doi:10.1145/359545.359563.
- [Lam03] L. Lamport (2003): *Arbitration-free synchronization*. *Distrib. Comput.* 16(2-3), pp. 219–237, doi:10.1007/s00446-002-0076-2.
- [Lyn96] N.A. Lynch (1996): *Distributed Algorithms*. Morgan Kaufmann Publishers.
- [Mi89] R. Milner (1989): *Communication and Concurrency*. Prentice Hall, Englewood Cliffs.
- [Nes00] U. Nestmann (2000): *What Is a ‘Good’ Encoding of Guarded Choice? Information and Computation* 156, pp. 287–319, doi:10.1006/inco.1999.2822.
- [Pal97] C. Palamidessi (1997): *Comparing the Expressive Power of the Synchronous and the Asynchronous pi-calculus*. In: *Conf. Record of the 24th ACM SIGPLAN-SIGACT Symp. on Principles of Programming Languages (POPL'97)*, ACM Press, pp. 256–265, doi:10.1145/263699.263731.
- [PNG13] K. Peters, U. Nestmann & U. Goltz (2013): *On Distributability in Process Calculi*. In M. Felleisen & Ph. Gardner, editors: *Programming Languages and Systems - Proceedings 22nd European Symposium on Programming, ESOP 2013, Rome, Italy, March 2013*, LNCS 7792, Springer, pp. 310–329, doi:10.1007/978-3-642-37036-6_18.
- [PSN11] K. Peters, J.-W. Schicke & U. Nestmann (2011): *Synchrony vs Causality in the Asynchronous Pi-Calculus*. In B. Luttik & F. Valencia, editors: *Proceedings 18th International Workshop on Expressiveness in Concurrency*, Aachen, Germany, 5th September 2011, *Electronic Proceedings in Theoretical Computer Science* 64, pp. 89–103, doi:10.4204/EPTCS.64.7.
- [RL94] M.O. Rabin & D.J. Lehmann (1994): *On the Advantages of Free Choice: A Symmetric and Fully Distributed Solution to the Dining Philosophers Problem*. In A.W. Roscoe, editor: *A Classical Mind: Essays in Honour of C.A.R. Hoare*, chapter 20, Prentice Hall, pp. 333–352. Extended abstract in: *Proceedings of POPL'81*, pages 133–138, doi:10.1145/567532.567547

- [Re82] W. Reisig (1982): *Deterministic Buffer Synchronization of Sequential Processes*. *Acta Informatica* 18, pp. 115–134, doi:10.1007/BF00264434.
- [Ro98] A.W. Roscoe (1998): *The Theory and Practice of Concurrency*. Prentice Hall. Available at <http://web.comlab.ox.ac.uk/oucl/work/bill.roscoe/publications/68b.pdf>.
- [SPG11] J.-W. Schicke, K. Peters & U. Goltz (2011): *Synchrony vs. Causality in Asynchronous Petri Nets*. In B. Luttik & F. Valencia, editors: Proceedings 18th International Workshop on Expressiveness in Concurrency, Aachen, Germany, 5th September 2011, *Electronic Proceedings in Theoretical Computer Science* 64, pp. 119–131, doi:10.4204/EPTCS.64.9.
- [Sel97] P. Selinger (1997): *First-Order Axioms for Asynchrony*. In: Proc. 8th International Conference on Concurrency Theory (CONCUR'97), Warsaw, Poland, LNCS 1243, Springer, pp. 376–390, doi:10.1007/3-540-63141-0_26.
- [Tau88] D. Taubner (1988): *Zur verteilten Implementierung von Petrinetzen*. *Informationstechnik* 30(5), pp. 357–370. Technical report, TUM-I 8805, TU München.
- [TV89] D. Taubner & W. Vogler (1989): *Step Failures Semantics and a Complete Proof System*. *Acta Informatica* 27(2), pp. 125–156, doi:10.1007/BF00265151.
- [Vo93] W. Vogler (1993): *Bisimulation and Action Refinement*. *Theoretical Computer Science* 114(1), pp. 173–200, doi:10.1016/0304-3975(93)90157-0.