Causal Structures for General Concurrent Behaviours

Ryszard Janicki[1], Jetty Kleijn[2], Maciej Koutny[3], and Łukasz Mikulski[3,4]

[1] McMaster University, Canada
[2] LIACS, Leiden University, The Netherlands
[3] Newcastle University, U.K.
[4] Nicolaus Copernicus University, Poland

Abstract. Non-interleaving semantics of concurrent systems is often expressed using posets, where causally related events are ordered and concurrent events are unordered. Each causal poset describes a unique concurrent history, i.e., a set of executions, expressed as sequences or step sequences, that are consistent with it. Moreover, a poset captures all precedence-based invariant relationships between the events in the executions belonging to its concurrent history. However, concurrent histories in general may be too intricate to be described solely in terms of causal posets. In this paper, we introduce and investigate generalised mutex order structures which can capture the invariant causal relationships in any concurrent history consisting of step sequence executions. Each such structure comprises two relations, viz. interleaving/mutex and weak causality. As our main result we prove that each generalised mutex order structure is the intersection of the step sequence executions which are consistent with it.

Keywords: concurrent history, causal poset, weak causal order, mutex relation, interleaving, step sequence, causality semantics.

1 Introduction

In order to design and validate complex concurrent systems, it is essential to understand the fundamental relationships between events occurring in their executions. However, looking at sequential descriptions of executions in the form of sequences or step sequences is insufficient when it comes to providing faithful information about causality and independence between events. To address this drawback, one may resort to using partially ordered sets of events providing explicit representation of causality in the executions of a concurrent system. In particular, the order in which independent events are observed may be accidental and those executions which only differ in the order of occurrences of independent events may be regarded as belonging to the same *concurrent history*, underpinned by a causal poset [1, 13, 16, 17, 21].

In general, concurrent behaviours can be investigated at the level of individual executions as well as at the level of order structures, like causal posets, capturing the essential invariant dependencies between events. The key link between these

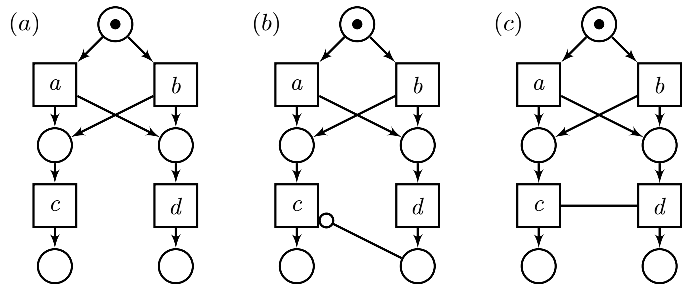

Fig. 1. A safe Petri net (a), extended with an inhibitor arc implying that when c is executed the output place of d must be empty (b), and extended with a mutex arc implying that c and d cannot be executed simultaneously (c).

two levels is the notion of a concurrent history [6], an *invariant closed* set Δ of executions. The latter means that Δ is fully determined by invariant relationships over X, its set of events: causality ($e \prec_\Delta f$ if, in all executions of Δ, e precedes f); weak causality ($e \sqsubset_\Delta f$ if, in all executions of Δ, e either precedes or is simultaneous with f); and interleaving/mutex ($e \rightleftharpoons_\Delta f$ if, in all executions of Δ, e is not simultaneous with f). In the case of safe Petri nets with sequential executions, $\prec_\Delta$ is the only invariant we need (as then, e.g., $\prec_\Delta \; = \; \sqsubset_\Delta$ and $\rightleftharpoons_\Delta \; = \; \prec_\Delta \cup \prec_\Delta^{-1}$). In particular, Δ is the set of all sequential executions corresponding to the linearisations of $\prec_\Delta$. The soundness of this approach is validated by Szpilrajn's Theorem [20] which states that each poset is equal to the intersection of its linearisations.

In this paper, executions are observed as step sequences, i.e., sequences of finite sets (steps) of simultaneously executed events. As an example, consider the safe Petri net depicted in Figure 1(a) which generates three step sequences involving a, c and d, viz. $\sigma = \{a\}\{c,d\}$, $\sigma' = \{a\}\{c\}\{d\}$ and $\sigma'' = \{a\}\{d\}\{c\}$. They can be seen as forming a single concurrent history $\Delta = \{\sigma, \sigma', \sigma''\}$ underpinned by a causal poset $\prec_\Delta$ satisfying $a \prec_\Delta c$ and $a \prec_\Delta d$. Moreover, such a Δ adheres to the following *true concurrency paradigm*:

> Given two events (c and d in Δ), they can be observed as simultaneous (in σ) $\Longleftrightarrow$ they can be observed in both orders (c before d in σ', and d before c in σ''). (TrueCon)

Concurrent histories adhering to TrueCon are underpinned by *causal partial orders*, in the sense that each such history comprises *all* step sequence executions consistent with a unique causal poset on events involved in the history.

In [6] fundamental concurrency paradigms are identified, including (TrueCon). Another paradigm is characterised by (TrueCon) with $\Longleftrightarrow$ replaced by $\Longleftarrow$. This paradigm has a natural system model interpretation provided by safe Petri nets with inhibitor arcs. Figure 1(b) depicts such a net generating two step sequences involving a, c and d, viz. $\sigma = \{a\}\{c,d\}$ and $\sigma' = \{a\}\{c\}\{d\}$. They form a concurrent history $\Delta' = \{\sigma, \sigma'\}$ adhering to the paradigm that unorderedness implies simultaneity, but *not* to the true concurrency paradigm as Δ' has no step sequence in which d precedes c although in σ, c and d occur in a single step.

As a result, histories adhering to the weaker paradigm are *not* underpinned by causal partial orders, but rather by causality structures $(X, \prec, \sqsubset)$ introduced in [7] — called *stratified order structures* (SO-structures) — based on causality and an additional weak causality ('not later than') relation. A version of Szpilrajn's Theorem can be shown to hold also for SO-structures and the concurrent histories they generate. Stratified order structures were independently introduced in [3] (as 'prossets'). Their comprehensive theory was developed in e.g. [8, 9, 12, 15]. As shown in this paper, SO-structures can be represented in a one-to-one manner by mutex order structures, or MO-structures, $(X, \rightleftharpoons, \sqsubset)$ based on interleaving/mutex and weak causality. The first, symmetric, relation defines the events that never occur simultaneously. Hence strict event precedence (causality) can be captured as a combination of mutex and weak causality.

This paper focuses on the least restrictive paradigm. i.e., there are no constraints imposed on concurrent histories. It admits all (invariant closed) concurrent histories comprising step sequence executions. As shown in [6], it is now sufficient to consider only two invariant relations, viz. mutex and weak causality. Figure 1(c) depicts a safe Petri net with mutex arcs (see [11]) generating two step sequences involving a, c and d, viz. $\sigma' = \{a\}\{c\}\{d\}$ and $\sigma'' = \{a\}\{d\}\{c\}$. We first observe that they form a concurrent history $\Delta'' = \{\sigma', \sigma''\}$ in which the executions of c and d interleave, and are both preceded by a; in other words, $c \rightleftharpoons_{\Delta''} d$, $a \sqsubset_{\Delta''} c$, $a \sqsubset_{\Delta''} d$ and $c \rightleftharpoons_{\Delta''} a \rightleftharpoons_{\Delta''} d$. That Δ'' is a concurrent history then follows from the observation that Δ'' contains *all* step sequences involving a, c and d which obey these invariant relationships. However, Δ'' does *not* conform to the two earlier considered paradigms as there is no step sequence in Δ'' in which c and d occur simultaneously. To summarise, a nonempty set Δ of step sequence executions over a common set of events X, is a concurrent history iff Δ consists of all step sequences σ over X such that for all $e, f \in X$: $e \rightleftharpoons_{\Delta} f$ implies that e and f are not simultaneous in σ, and $e \sqsubset_{\Delta} f$ implies that e precedes or is simultaneous with f in σ.

The aim of this paper is to provide a structural characterisation of general concurrent histories (consisting of step sequence executions). An early attempt to describe structures of this kind was made in [4]. The there proposed generalised stratified order structures (or GSO-structures) do however not always capture all implied invariant relationships involving the mutex relation. Here, we will show that *generalised mutex order structures* (or GMO-structures) describe exactly all general concurrent histories. The main result is a version of Szpilrajn's Theorem, formulated and proven to hold for GMO-structures and concurrent histories. For this we develop a notion of GMO-closure which is the GMO-structure counterpart of transitive closure of an acyclic relation.
First, we recall key notions and notations used throughout the paper. In Section 3, we introduce MO-structures and establish their relationship with stratified order structures. Then, Section 4 introduces GMO-structures and proves

their main properties, including GMO-closure and the GMO-structure version of Szpilrajn's Theorem. Section 5 presents concluding remarks.
Proofs of all the results can be found in the Technical Report available at http://www.cs.ncl.ac.uk/publications/trs/papers/1378.pdf.

2 Preliminary Definitions

Throughout the paper we use the standard notions of set theory and formal language theory. In particular, $\uplus$ denotes disjoint set union. The identity relation on a set X is defined as $Id_X = \{\langle a, a\rangle \mid x \in X\}$, the index X may be omitted if it is clear from the context.

Composing relations. The composition of two binary relations, R and Q, over X is given by $R \circ Q = \{\langle a, b\rangle \mid \exists x \in X : \ aRx \wedge xQb\}$. Moreover, if $P \subseteq X \times X$, then we define $R \circ_P Q = \{\langle a, b\rangle \mid \exists \langle x, y\rangle \in P : \ aRxQb \wedge aRyQb\}$ (see Figure 2).

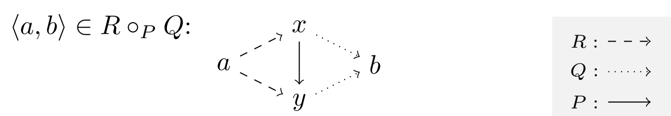

Fig. 2. A visualisation of $\circ_P$ composition.

Note that $\circ = \circ_{Id}$, and the associativity of relation composition holds for the extended notion. We will also denote $a_1 \dots a_k R b_1 \dots b_m$ whenever $a_i R b_j$, for all i, j. For example, $aRbcQd$ means that $aRbQd$ and $aRcQd$.

Given a relation $R \subseteq X \times X$, $R^0 = Id$ and $R^n = R^{n-1} \circ R$, for all $n \geq 1$. Then: (i) the reflexive closure of R is defined by $R \cup Id$; (ii) the transitive closure by $R^+ = \bigcup_{i \geq 1} R^i$; (iii) the reflexive transitive closure by $R^* = Id \cup R^+$; and (iv) the irreflexive transitive closure by $R^\curlywedge = R^+ \setminus Id = R^* \setminus Id$. Moreover, the inverse of R is given by $R^{-1} = \{\langle a, b\rangle \mid \langle b, a\rangle \in R\}$, and the symmetric closure by $R^{sym} = R \cup R^{-1}$.

Order relations. A relation $R \subseteq X \times X$ is: (i) symmetric if $R = R^{-1}$; (ii) antisymmetric if $R \cap R^{-1} \subseteq Id$; (iii) reflexive if $Id \subseteq R$; (iv) irreflexive if $Id \cap R = \varnothing$; (v) transitive if $R \circ R \subseteq R \cup Id$; and (vi) total if $R \cup R^{-1} = X \times X$.

A relation $R \subseteq X \times X$ is: (i) an *equivalence relation* if it is symmetric, transitive and reflexive; (ii) a *pre-order* if it is transitive and irreflexive; (iii) a *partial order* if it is an antisymmetric pre-order; and (iv) a *total order* if it is a partial order and $R \cup Id$ is total; (v) a *stratified order* if it is a partial order such that $X \times X \setminus R^{sym}$ is an equivalence relation.

Every irreflexive relation $R \subseteq X \times X$ induces a least pre-order containing R defined by $R^\curlywedge$. Following E. Schröder [19], we define the largest equivalence relation contained in R^* as $R^\circledast = R^* \cap (R^*)^{-1} = (R^\curlywedge \cap (R^\curlywedge)^{-1}) \uplus Id$.

For a stratified order $R \subseteq X \times X$ we define two relations, $\sqsubset_R$ and $\rightleftharpoons_R$, such that, for all distinct $a, b \in X$:

$$a \sqsubset_R b \iff \neg(bRa) \quad \text{and} \quad a \rightleftharpoons_R b \iff \neg(a \sqsubset_R^{\circledast} b) \iff aRb \vee bRa \;.$$

Intuitively, if R represents a stratified order execution, aRb means 'a *occurred earlier than* b'. In such a case $a \sqsubset_R b$ means 'a *occurred not later than* b', $a \rightleftharpoons_R b$ means 'a *did not occur simultaneously with* b', and $a \sqsubset_R^{\circledast} b$ means 'a *occurred simultaneously with* b'.

Relational structures. A tuple $S = (X, R_1, R_2, \ldots, R_n)$, where $n \geq 1$ and each $R_i \subseteq X \times X$ is a binary relation on X, is an (n-ary) *relational* structure. By the *domain* of a relational structure S we mean the set X. An *extension* of S is any relational structure $S' = (X, R'_1, R'_2 \ldots, R'_n)$ satisfying $R_i \subseteq R'_i$, for every $1 \leq i \leq n$. We denote this by $S \subseteq S'$. The *intersection* of a nonempty family $\mathcal{R} = \{(X, R_1^i, \ldots, R_n^i) \mid i \in I\}$ of relational structures with the same domain and arity is given by $\bigcap \mathcal{R} = (X, \bigcap_{i \in I} R_1^i, \ldots, \bigcap_{i \in I} R_n^i)$. In what follows, we consider only relational structures that contain two relations, while the set X is finite.

A relational structure $S = (X, Q, R)$ is: (i) *separable* if $Q \cap R^{\circledast} = \varnothing$, Q is symmetric and R is irreflexive; and (ii) *saturated* in a family of relational structures $\mathscr{X}$ if it belongs to $\mathscr{X}$ and for every extension $S' \in \mathscr{X}$ of S, we have $S = S'$. It is easily seen that an intersection of separable relational structures is also separable. Intuitively, if Q represents 'mutex' and R 'weak precedence', then separability means that simultaneous events cannot be in the mutex relation.

A *stratified order structure* (or SO-structure) is defined as a relational structure $sos = (X, \prec, \sqsubset)$, where $\prec$ and $\sqsubset$ are binary relations on X such that, for all $a, b, c \in X$:

$$\begin{array}{ll} \mathit{S1}: a \not\sqsubset a & \mathit{S3}: \; a \sqsubset b \sqsubset c \wedge a \neq c \implies a \sqsubset c \\ \mathit{S2}: a \prec b \implies a \sqsubset b & \mathit{S4}: \; a \sqsubset b \prec c \vee a \prec b \sqsubset c \implies a \prec c \,. \end{array}$$

A *generalized stratified order structure* [4] (or GSO-structure) is a relational structure $gsos = (X, \rightleftharpoons, \sqsubset)$ such that $\sqsubset$ is irreflexive, $\rightleftharpoons$ is irreflexive and symmetric, and $(X, \rightleftharpoons \cap \sqsubset, \sqsubset)$ is an SO-structure. A comprehensive treatment of GSO-structures can be found in [5].

Properties. For every binary relation $R \subseteq X \times X$ and all $a, b \in X$, we have:

$$(R \cup \langle a, b\rangle)^* = R^* \cup \{\langle c, d\rangle \mid cR^*a \wedge bR^*d\} \,. \tag{1}$$

$$\neg(bR^*a) \implies (R \cup \langle a, b\rangle)^{\circledast} = R^{\circledast} \tag{2}$$

$$R^{\circledast} = (R^{\circledast})^{-1} \subseteq R^* \tag{3}$$

$$(R^{\curlywedge})^{\curlywedge} = R^{\curlywedge} \qquad (R^{\curlywedge})^* = R^* \qquad (R^*)^* = R^* \qquad (R^{\curlywedge})^{\circledast} = R^{\circledast} \tag{4}$$

$$R^{\circledast} \circ R^{\circledast} = R^{\circledast} \qquad R^* \circ R^{\circledast} = R^{\circledast} \circ R^* = R^* \circ R^* = R^* \tag{5}$$

If $R \subseteq X \times X$ is a stratified order, then $\rightleftharpoons_R$ is irreflexive and symmetric, while $\sqsubset_R$ is a pre-order such that:

$$\sqsubset_R = \sqsubset_R^+ \setminus Id = \sqsubset_R^{\curlywedge} \quad \text{and} \quad \sqsubset_R^{\circledast} \setminus Id = \sqsubset_R \cap \sqsubset_R^{-1} \;. \tag{6}$$

Moreover, for all distinct $a, b \in X$, we have:

$$\neg(a \rightleftharpoons_R b) \iff a \sqsubset_R b \wedge b \sqsubset_R a \tag{7}$$

$$\neg(a \sqsubset_R b) \implies b \sqsubset_R a \tag{8}$$

$$aRb \iff a \rightleftharpoons_R b \wedge a \sqsubset_R b \,. \tag{9}$$

and exactly one of the following holds:

$$\begin{array}{ccccccccc} a & \rightleftharpoons & b & \rightleftharpoons & a & \sqsubset & b & \not\sqsubset & a \\ a & \rightleftharpoons & b & \rightleftharpoons & a & \not\sqsubset & b & \sqsubset & a \\ a & \not\rightleftharpoons & b & \not\rightleftharpoons & a & \sqsubset & b & \sqsubset & a \,. \end{array} \tag{10}$$

Intuitively, (9) means that 'a occurred earlier than b' iff 'a and b were not simultaneous' and 'a occurred not later than b'.

3 Separable Order Structures

In this section we take another look at the stratified order structures, substantially different from that of, e.g., [8, 12, 15]. We provide for them a new representation more suitable for further generalisation. The new representation replaces causal orders by mutex relations between events. While SO-structures allow for more compact representation (strict precedence involves fewer pairs of events than mutex), the new order structures are easier to generalise to cater for general interleaving/mutex requirements and their properties.

In the rest of this paper, we will be concerned with order structures of the form $S = (X, \rightleftharpoons, \sqsubset)$. Intuitively, X is a set of events involved in some history of a concurrent system, $\rightleftharpoons$ is a 'mutex' (or 'interleaving') relation which identifies pairs of events which cannot occur simultaneously, and $\sqsubset$ is a 'weak precedence' relation between events. The latter means, in particular, that if $a \sqsubset b \sqsubset a$ then a and b must occur simultaneously in any execution belonging to the history represented by S; in other words, S must be *separable* (i.e., $\rightleftharpoons \cap \sqsubset^{\circledast} = \varnothing$).

Mutex order structures The definition of the first class of order structures based on mutex and weak precedence relations is motivated by the observation that the 'precedence' (or 'causality') relation is nothing but 'mutex'+'weak precedence', c.f. (9). Therefore, the axioms defining stratified order structures can easily be rendered in terms of the latter relations.

Definition 1 (mutex order structure). *A* mutex order structure *(*MO*-structure) is a relational structure $mos = (X, \rightleftharpoons, \sqsubset)$, where $\rightleftharpoons$ and $\sqsubset$ are binary relations on X such that, for all $a, b, c \in X$:*

$$\begin{array}{ll} M1: & a \rightleftharpoons b \implies b \rightleftharpoons a \\ M2: & a \not\sqsubset a \\ M3: & a \rightleftharpoons b \implies a \sqsubset b \vee b \sqsubset a \\ M4: & a \sqsubset b \sqsubset c \wedge a \neq c \implies a \sqsubset c \\ M5: & a \sqsubset b \sqsubset c \wedge (a \rightleftharpoons b \vee b \rightleftharpoons c) \implies a \rightleftharpoons c \,. \end{array}$$

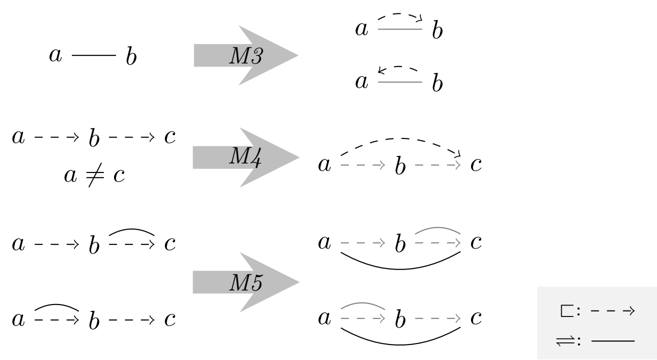

Fig. 3. A visualisation of axioms *M3* − *M5*.

Axioms *M3* − *M5* are illustrated in Figure 3, and some relevant properties of MO-structures are given below.

Proposition 1. *Let $mos = (X, \rightleftharpoons, \sqsubset)$ be an* MO*-structure. Then mos is separable and, for all $a, b, c, d \in X$, we have:*

$$a \not\rightleftharpoons a \tag{11}$$

$$a \sqsubset b \sqsubset a \ \wedge\ a \rightleftharpoons c \quad\Longrightarrow\quad b \rightleftharpoons c \tag{12}$$

$$a \sqsubset c \sqsubset b \ \wedge\ a \sqsubset d \sqsubset b \ \wedge\ c \rightleftharpoons d \quad\Longrightarrow\quad a \rightleftharpoons b\,. \tag{13}$$

The next results demonstrate that MO-structures are in a one-to-one relationship with SO-structures. Below, we use two mappings between these two classes of order structures. For every SO-structure $sos = (X, \prec, \sqsubset)$, we define $\mathsf{so2mo}(sos) = (X, \prec^{sym}, \sqsubset)$, and for every MO-structure $mos = (X, \rightleftharpoons, \sqsubset)$, we define $\mathsf{mo2so}(mos) = (X, \rightleftharpoons \cap \sqsubset, \sqsubset)$.

Theorem 1. *The mappings* mo2so *and* so2mo *are inverse bijections.*

Layered order structures In general, order structures like MO-structures are not saturated, and may capture histories comprising several executions (like a single causal partial order may have numerous total order linearisations). However, there is also a class of order structures which correspond in a one-to-one way to step sequences.

Definition 2. *Let $R \subseteq X \times X$ be a stratified order. Then $los = (X, \rightleftharpoons_R, \sqsubset_R)$ is the* layered order structure *(or* LO*-structure) induced by R.*

For a separable relational structure $sr = (X, \rightleftharpoons, \sqsubset)$, we will denote by $\mathsf{sr2los}(sr)$ the set of all LO-structures los extending sr, i.e., $sr \subseteq los$. With this notation, a nonempty set LOS of LO-structures is a concurrent history if $LOS = \mathsf{sr2los}\left(\bigcap LOS\right)$.

Proposition 2. *Every layered order structure is separable and saturated in the set of all separable order structures.*

Proposition 3. *Every* LO*-structure is an* MO*-structure.*

An MO-structure is linked with LO-structures (step sequences) through the set $\mathsf{sr2los}(mos)$ of all LO-structures extending mos. Similarly, for every SO-structure sos we can define $\mathsf{so2los}(sos) = \mathsf{sr2los}(\mathsf{so2mo}(sos))$. It can then be seen ([8]) that $\mathsf{so2los}(sos)$ is a nonempty set and (in the notation used in this paper):

$$sos = \bigcap \mathsf{mo2so}(\mathsf{so2los}(sos)) \ . \tag{14}$$

That result corresponds to Szpilrajn's Theorem that any partial order is the intersection of its linearisations (c.f. [5, 8]). Such a result extends to MO-structures and we obtain

Theorem 2. *For every* MO*-structure* mo, $\mathsf{sr2los}(mo) \neq \varnothing$ *and* $mo = \bigcap \mathsf{sr2los}(mo)$.

We can therefore conclude that the saturated extensions of an MO-structure mos form a concurrent history represented by mos. It is then important to ask which concurrent histories can be derived in this way; in other words, which concurrent histories can be represented by MO-structures.

Consider now a nonempty set $LOS = \{(X, \rightleftharpoons_i, \sqsubset_i) \mid i \in I\}$ of LO-structures forming a concurrent history, and their intersection $S = \bigcap LOS = (X, \rightleftharpoons, \sqsubset)$. Since every LO-structure is also an MO-structure, we immediately obtain that S is an order structure satisfying axioms *M1*, *M2*, *M4* and *M5*. However, *M3* in general does not hold although it holds for histories in which the possibility of executing two events in either order always implies also simultaneous execution, meaning that, for all distinct $a, b \in X$,

$$\left.\begin{array}{l} \exists i \in I: \ \langle a,b\rangle \ \in \rightleftharpoons_i \cap \sqsubset_i \\ \exists j \in I: \ \langle b,a\rangle \ \in \rightleftharpoons_j \cap \sqsubset_j \end{array}\right\} \implies \exists k \in I: \ \langle a,b\rangle \ \in \sqsubset_j^{sym} \ .$$

One might now wonder what happens if we do not assume any special properties of a concurrent history. As we will show in the rest of the paper, Proposition 1 in combination with the observation that it always holds for $S = \bigcap LOS$, yields axioms for order structures underpinning general concurrent histories.

4 Generalised Order Structures

In this section, we provide a complete characterisation of general concurrent histories where executions are represented by layered order structures; in other words, histories comprising step sequence executions. We achieve this by retaining all those MO-structure axioms which hold in general, and then replacing the only dropped axiom *M3* by Proposition 1.

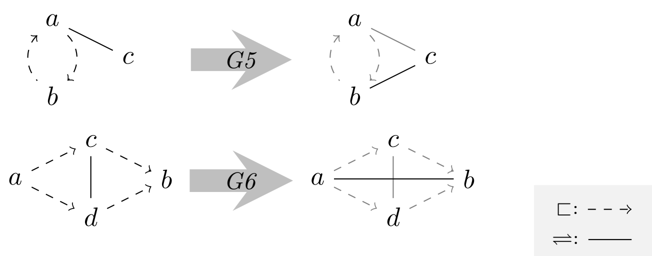

Fig. 4. A visualisation of axioms *G5* and *G6*.

Definition 3 (generalised mutex order structure). *A* generalised mutex order structure *(or* GMO*-structure) is a relational structure gmos* $= (X, \rightleftharpoons, \sqsubset)$*, where* $\rightleftharpoons$ *and* $\sqsubset$ *are binary relations on* X *such that, for all* $a, b, c, d \in X$*:*

G1 :	$a \rightleftharpoons b \implies b \rightleftharpoons a$	*M1*
G2 :	$a \not\sqsubset a \;\wedge\; a \not\rightleftharpoons a$	*M2* & (11)
G3 :	$a \sqsubset b \sqsubset c \;\wedge\; a \neq c \implies a \sqsubset c$	*M4*
G4 :	$a \sqsubset b \sqsubset c \;\wedge (a \rightleftharpoons b \;\vee\; b \rightleftharpoons c) \implies a \rightleftharpoons c$	*M5*
G5 :	$a \sqsubset b \sqsubset a \;\wedge\; a \rightleftharpoons c \implies b \rightleftharpoons c$	(12)
G6 :	$a \sqsubset c \sqsubset b \;\wedge\; a \sqsubset d \sqsubset b \;\wedge\; c \rightleftharpoons d \implies a \rightleftharpoons b$	(13)

The set of axioms in Definition 3 is minimal (see Figure 5). Moreover, GMO-structures enjoy a number of useful properties.

Proposition 4. *Let gmos* $= (X, \rightleftharpoons, \sqsubset)$ *be a* GMO*-structure. Then gmos is separable and, for all* $a, b \in X$*, we have:*

$$a \sqsubset^{\curlywedge} b \implies a \sqsubset b \qquad\qquad a \sqsubset b \sqsubset a \implies a \not\rightleftharpoons b\,.$$

Proposition 5. *Each* MO*-structure is a* GMO*-structure.*

The converse of Proposition 5 does not hold; for example, as *M3* does not hold, $(\{a, b\}, \{\langle a, b\rangle, \langle b, a\rangle\}, \varnothing)$ is a GMO-structure but *not* an MO-structure.

Proposition 6. *If gmos* $= (X, \rightleftharpoons, \sqsubset)$ *is a* GMO*-structure, then* $(X, \rightleftharpoons \cap \sqsubset, \sqsubset)$ *is an* SO*-structure.*

Proposition 6 states that every GMO-structure is a GSO-structure. We observe that the converse is not true, with suitable counterexamples provided by the GSO-structures S_{G5} and S_{G6} in Figure 5.

Closure operator for generalised mutex order structures We will now provide a method for deriving valid GMO-structures from other relational structures.

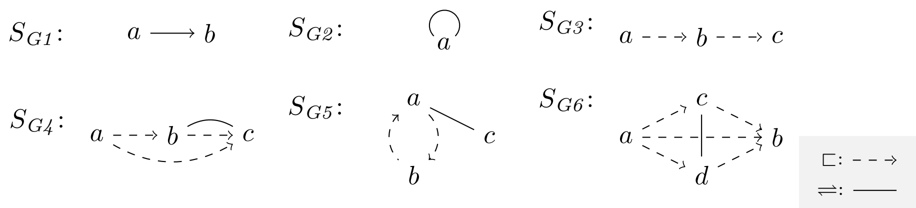

Fig. 5. Examples showing that the set of axioms in Definition 3 is minimal. Each relational structure S_{Gi} satisfies all axioms except for Gi.

Definition 4 (GMO-closure). *Let $S = (X, Q, R)$ be a relational structure and $Q^{[R]} = R^{\circledast} \circ \big(Q \cup (R^* \circ_Q R^*)^{sym}\big) \circ R^{\circledast}$. Then the* GMO*-closure of S is given by:*

$$S^{\blacklozenge} = \big(X\ ,\ Q^{[R]}\ ,\ R^{\curlywedge}\big)\ .$$

The GMO-closure operator introduced here can be seen as related to two different closure operators: (i) the transitive closure operator for acyclic reflexive binary relations; and (ii) the $\Diamond$-operator for $\Diamond$-acyclic order structures introduced in [7] in order to obtain SO-structures. It is also be seen as a generalisation of the GSO-closure introduced in [10] for GSO-acyclic structures in order to obtain GSO-structures.

The main property we want from the notion of GMO-closure is that whenever $S = (X, Q, R)$ is a separable relational structure, $S^{\blacklozenge}$ is a GMO-structure. Furthermore, if S is already a GMO-structure, then we want $S^{\blacklozenge} = S$. The form of $Q^{[R]}$ follows from the requirement that $S^{\blacklozenge}$ should be a GMO-structure and the axioms for GMO-structures. In particular the factor $(R^* \circ_Q R^*)^{sym}$ follows from axioms *G4* and *G6*, while the factor $R^{\circledast} \circ_Q R^{\circledast}$ corresponds to *G5*.

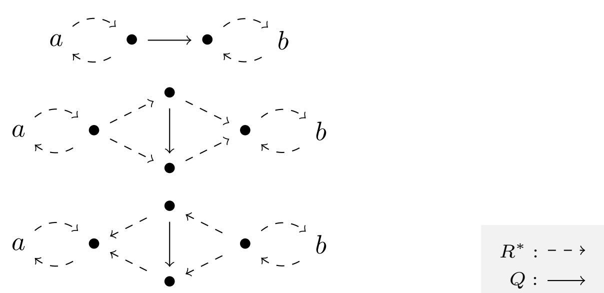

Fig. 6. A visualisation of the three cases of $\langle a, b\rangle \in Q^{[R]}$.

The next four results respectively correspond to saying that: (*i*) the transitive closure of an acyclic relation is also acyclic; (*ii*) GMO-closure is a closure operation in the usual sense; (*iii*) the transitive closure of an acyclic relation yields a poset; and (*iv*) posets are transitively closed.

Proposition 7. *If S is a separable relational structure, then $S^\blacklozenge$ is also separable, $S \subseteq S^\blacklozenge$ and $(S^\blacklozenge)^\blacklozenge = S^\blacklozenge$. Moreover, $S^\blacklozenge$ is a* GMO*-structure.*

Proposition 8. *If gmos is a* GMO*-structure, then $gmos^\blacklozenge = gmos$.*

As layered order structures and mutex order structures are special cases of generalised mutex order structures, we obtain an immediate

Corollary 1. *Let los be an* LO*-structure and mos be an* MO*-structure. Then $los^\blacklozenge = los$ and $mos^\blacklozenge = mos$.*

The following technical lemma describes a single stage of the saturation process for a GMO-structure leading to a LOS-structure. In such a process, we may add an arbitrary link between two elements that do not yet satisfy (10). We only need to remember that in the case of extending the relation Q, together with $\langle a, b\rangle$ we have to add $\langle b, a\rangle$. After such an addition, we get a separable order structure that may be closed. As a result, we obtain one of possible extensions of an initial *gmos*. The above process terminates after obtaining an LO-structure and it is central to the proof of the main Theorem 3.

In what follows, we denote $R_{xy} = R \cup \{\langle x, y\rangle\}$ and $Q_{xyx} = Q \cup \{\langle x, y\rangle, \langle y, x\rangle\}$.

Lemma 1. *Let $gmos = (X, Q, R)$ be a* GMO*-structure, $a, b \in X$ and $a \neq b$.*

$$\begin{array}{lll}
\langle a,b\rangle \notin R \ \wedge\ \langle b,a\rangle \notin R & \Longrightarrow & (X,Q,R_{ab})^\blacklozenge \quad \textit{is a } \textsc{gmo}\textit{-structure} \\
\langle a,b\rangle \notin R \ \wedge\ \langle a,b\rangle \notin Q & \Longrightarrow & (X,Q,R_{ab})^\blacklozenge \quad \textit{is a } \textsc{gmo}\textit{-structure} \\
\langle a,b\rangle \notin R \ \wedge\ \langle a,b\rangle \notin Q & \Longrightarrow & (X,Q_{aba},R)^\blacklozenge \quad \textit{is a } \textsc{gmo}\textit{-structure}\,.
\end{array}$$

To complete the properties of the saturation process described in Lemma 1 and used in the proof of Theorem 3, we formulate the following

Lemma 2. *Let $gmos = (X, Q, R)$ be a* GMO*-structure such that $a, b \in X$, $a \neq b$, $\langle a, b\rangle \notin R$ and $\langle a, b\rangle \notin Q$ and $S' = (X, Q, R_{ab})^\blacklozenge = (X, Q', R^{\curlywedge}_{ab})$. Then $\langle a, b\rangle \notin Q'$.*

In Lemmas 1 and 2 we have captured a method of saturating GMO-structures that are not LO-structures. It moreover allows us to formulate an immediate

Corollary 2. *Every relational structure saturated among all separable relational structures is a layered order structure.*

General concurrent histories We now return to our original goal which was to provide a structural characterisation of all histories comprising step sequence executions. Recall that $\mathsf{sr2los}(gmos)$ is the set of all LO-structures associated with a GMO-structure *gmos*. Then we obtain a result corresponding to Szpilrajn's Theorem:

Theorem 3. *For every* GMO*-structure gmos,*

$$\mathsf{sr2los}(gmos) \neq \varnothing \quad and \quad gmos = \bigcap \mathsf{sr2los}(gmos)\ .$$

Together with the fact that, for every nonempty set LOS of LO-structures with the same domain, $\bigcap LOS$ is a GMOS-structure, this leads to the conclusion that all concurrent histories are represented by GMO-structures.

5 Concluding Remarks

We can finally clarify the relationship between GSO-structures and GMO-structures. In general, in order to accept an order structure $os = (X, \rightleftharpoons, \sqsubset)$ as an invariant representation of a concurrent history, we require that

$$\mathsf{sr2los}(os) \neq \varnothing \ \text{ and } \ os = \bigcap \mathsf{sr2los}(os)\ .$$

We demonstrated that this property holds whenever os is a GMO-structure, and that it may fail to hold for a GSO-structure. We have further shown that GMO-structures are GSO-structures, but that the converse does not hold. However, what is the case is that each GSO-structure $gsos$ is separable, and so its GMO-closure $gsos^{\blacklozenge}$ is a GMO-structure satisfying $\mathsf{sr2los}(gsos^{\blacklozenge}) = \mathsf{sr2los}(gsos)$. In other words, concurrent histories described by separable order structures and their GMO-closures are the same. The importance of GSO-structures comes from the fact that they paved the way for GMO-structures, by exposing the fundamental property that causal ordering is a combination of mutex and weak ordering.

A key motivation for the research presented in this paper comes from concurrent behaviours as exhibited by safe Petri nets with mutex arcs. The resulting semantical approach — which has been meticulously worked out above — is based on GMO-structures which characterise all concurrent histories comprising step sequence executions. A natural direction for further work is to provide a compatible language-theoretic representation of concurrent histories, by generalising Mazurkiewicz traces [13] which correspond to causal posets, and comtraces [7] which correspond to SO-structures (or MO-structures). This development would also allow to link the dynamic notions of mutex and weak causality with the static properties of Petri nets with mutex arcs. The resulting semantics can also support efficient verification techniques [2, 14, 18].

Acknowledgements

We would like to thank the anonymous reviewers for useful comments and suggestions. This research was supported by a fellowship funded by the "Enhancing Educational Potential of Nicolaus Copernicus University in the Disciplines of Mathematical and Natural Sciences" Project POKL.04.01.01-00-081/10, the EPSRC GAELS and UNCOVER projects, and by an NSERC of Canada grant.

References

1. Best, E., Devillers, R.: Sequential and concurrent behaviour in Petri net theory. TCS 55(1), 87–136 (1987)
2. Esparza, J., Heljanko, K.: Unfoldings: A Partial-Order Approach to Model Checking. Monographs in Theoretical Computer Science, Springer (2008)
3. Gaifman, H., Pratt, V.R.: Partial order models of concurrency and the computation of functions. In: LICS. pp. 72–85 (1987)
4. Guo, G., Janicki, R.: Modelling concurrent behaviours by commutativity and weak causality relations. In: AMAST. pp. 178–191 (2002)
5. Janicki, R.: Relational structures model of concurrency. Act. Inf. 45, 279–320 (2008)
6. Janicki, R., Koutny, M.: Structure of concurrency. TCS 112(1), 5–52 (1993)
7. Janicki, R., Koutny, M.: Semantics of inhibitor nets. Inf.&Comp. 123(1), 1 – 16 (1995)
8. Janicki, R., Koutny, M.: Fundamentals of modelling concurrency using discrete relational structures. Acta Inf. 34, 367–388 (1997)
9. Juhás, G., Lorenz, R., Mauser, S.: Causal semantics of algebraic Petri nets distinguishing concurrency and synchronicity. Fundam. Inf. 86(3), 255–298 (2008)
10. Kleijn, J., Koutny, M.: The mutex paradigm of concurrency. In: Petri Nets. pp. 228–247 (2011)
11. Kleijn, J., Koutny, M.: Mutex causality in processes and traces. Fundam. Inf. 122, 119–146 (2013)
12. Le, D.T.M.: On three alternative characterizations of combined traces. Fundam. Inf. 113(3), 265–293 (2011)
13. Mazurkiewicz, A.: Concurrent program schemes and their interpretations. DAIMI Rep. PB 78, Aarhus University (1977)
14. McMillan, K.L.: Using unfoldings to avoid the state explosion problem in the verification of asynchronous circuits. In: CAVC. pp. 164–177 (1992)
15. Mikulski, Ł., Koutny, M.: Hasse diagrams of combined traces. In: Brandt, J., Heljanko, K. (eds.) ACSD. pp. 92–101 (2012)
16. Nielsen, M., Plotkin, G.D., Winskel, G.: Petri nets, event structures and domains, Part I. TCS 13, 85–108 (1981)
17. Pratt, V.R.: Some constructions for order-theoretic models of concurrency. In: Logic of Programs. pp. 269–283 (1985)
18. Rodriguez, C., Schwoon, S., Khomenko, V.: Contextual merged processes. In: Petri Nets. LNCS, vol. 7927, pp. 29–48 (2013)
19. Schröder, E.: Vorlesungen über die Algebra der Logik (Exakte Logik). Dritter Band: Algebra und Logik der Relative. B. G. Teubner (1895)
20. Szpilrajn, E.: Sur l'extension de l'ordre partiel. Fundam. Math. 16, 386–389 (1930)
21. Winkowski, J., Maggiolo-Schettini, A.: An algebra of processes. Journal of Computer and System Sciences 35(2), 206–228 (1987)