

Branching Pomsets for Choreographies

Luc Edixhoven

Open University (Heerlen) and
CWI (Amsterdam), Netherlands
{led,ssj}@ou.nl

Sung-Shik Jongmans

José Proença

CISTER, ISEP,
Polytechnic Institute of Porto, Portugal
pro@isep.ipp.pt

Guillermina Cledou

HASLab, INESC TEC & University of Minho, Portugal
mgc@inesctec.pt

Choreographic languages describe possible sequences of interactions among a set of agents. Typical models are based on languages or automata over sending and receiving actions. Pomsets provide a more compact alternative by using a partial order over these actions and by not making explicit the possible interleaving of concurrent actions. However, pomsets offer no compact representation of choices. For example, if an agent Alice can send one of two possible messages to Bob three times, one would need a set of $2 \times 2 \times 2$ distinct pomsets to represent all possible branches of Alice’s behaviour. This paper proposes an extension of pomsets, named *branching pomsets*, with a branching structure that can represent Alice’s behaviour using $2 + 2 + 2$ ordered actions. We encode choreographies as branching pomsets and show that the pomset semantics of the encoded choreographies are bisimilar to their operational semantics.

1 Introduction

Choreographic languages describe possible sequences of interactions, or communication protocols, among a set of agents. Their use is well established [11, 1, 9, 10, 2, 6], and it typically includes:

1. reasoning statically over interaction properties;
2. generating code that facilitates the implementation of the concurrent protocols.

Regarding 1, static properties include deadlock absence or the equivalence between global protocols and the parallel composition of local protocols for each agent. Regarding 2, the code generated from choreographic languages include skeleton code for concurrent code, generated behavioural types that can be used to type-check agents, or dedicated orchestrators that dictate how the agents can interact. In this work we focus on how to analyse choreographies by proposing a new structure to compactly represent their behaviour, based on *partial-ordered multisets* (pomsets).

We foresee applications of this work in both aforementioned areas. Regarding 1, in static analysis, a more compact model of choreographies could reduce the complexity of the analysis of protocols featuring both concurrency and choices. Our work in progress in this area focuses on realisability. Regarding 2, our recent work includes API generation using an approach based on traditional sets of pomsets [5]. We are keen to extend it to take full advantage of the branching structure presented in this paper.

We use two simple running examples to motivate our approach.




Figure 1: An automaton (left) and a pomset (right) representing the master-workers protocol.

- 1. Master-workers (MW) protocol [13].** A *master* (m) concurrently sends *tasks* (t) to a number of *workers* (w_1, \dots, w_n). Once workers finish their task, they inform the master that they are *done* (d). This protocol is expressed in our choreographic language as follows for the case of two workers.

$$(m \rightarrow w_1 : t ; w_1 \rightarrow m : d) \parallel (m \rightarrow w_2 : t ; w_2 \rightarrow m : d).$$

Here, $m \rightarrow w_1 : t$ represents an asynchronous communication from m to w_1 of a message of type t , ‘;’ represents sequential composition and ‘ \parallel ’ represents parallel composition.

- 2. Distributed voting (DV) protocol.** Three participants – Alice (a), Bob (b) and Carol (c) – send their vote (*yes* (y) or *no* (n)) to every other participant in parallel. This is expressed as follows, where $+$ indicates nondeterministic choice.

$$\left(\begin{array}{l} (a \rightarrow b : y \parallel a \rightarrow c : y)^+ \\ (a \rightarrow b : n \parallel a \rightarrow c : n) \end{array} \right) \parallel \left(\begin{array}{l} (b \rightarrow a : y \parallel b \rightarrow c : y)^+ \\ (b \rightarrow a : n \parallel b \rightarrow c : n) \end{array} \right) \parallel \left(\begin{array}{l} (c \rightarrow a : y \parallel c \rightarrow b : y)^+ \\ (c \rightarrow a : n \parallel c \rightarrow b : n) \end{array} \right)$$

A protocol can evolve by performing sequences of sending and receiving actions. E.g., $ab!x$ denotes a sending action from a to b with a message of type x , and $ab?x$ denotes the dual receiving action on b . Protocols with parallel interactions can have an explosion of states, such as our MW protocol, whose full state machine can be found on the left of Figure 1. To avoid this explosion, the state space can be represented more compactly using so-called *partially ordered multisets*, or simply pomsets [12, 8]. The right of Figure 1 shows a graphical pomset representation of the same MW protocol. The pomset contains eight events, whose labels are shown. The arrows visualise the partial order: an event precedes any other event to which it has an outgoing arrow, either directly or transitively. In this example, the event with label $mw_1!t$ precedes the event with label $mw_1?t$ directly and the events with labels $w_1m!d$ and $w_1m?d$ transitively. However, it is independent of the events involving w_2 .

The behaviour represented by a pomset is the set of all its linearisations, i.e., all sequences of the labels of its events that respect their partial order. The set of linearisations of the pomset in Figure 1 consists of all interleavings of the two threads $mw_1!t \ mw_1?t \ w_1m!d \ w_1m?d$ and $mw_2!t \ mw_2?t \ w_2m!d \ w_2m?d$. This explicit concurrency yields a compact representation of the possible interleavings using just $4 + 4$ events, whereas the state machine needs 5×5 states to represent all interleavings. If we were to add a third worker, the automaton would grow by another factor 5, while the pomset would expand by just four additional events.

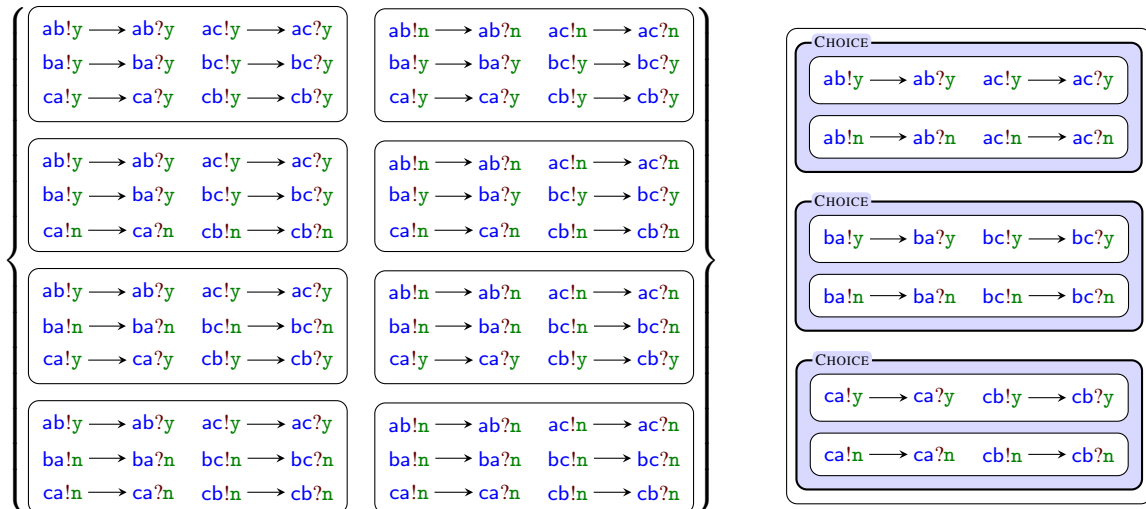


Figure 2: A set of pomsets (left) and a branching pomset (right) representing a three-participant distributed vote.

While pomsets can compactly represent concurrent behaviour, **choices** need to be represented as *sets* of pomsets: one for every branch. As a consequence, one might need an exponential number of pomsets to represent a protocol with many choices. The exponential growth is visible in our DV protocol with three participants, depicted on the left side of Figure 2. This diagram represents a set of pomsets that capture the protocol's possible behaviour, counting $2 \times 2 \times 2$ different pomsets. If we were to add a fourth participant, the set would grow by another factor 2.

This paper proposes an extension to pomsets, named *branching pomsets*, with a branching structure that can compactly represent choices. A branching pomset initially contains all branches of choices, and discards non-chosen branches when firing events that require resolving a choice. The right side of Figure 2 depicts an example of a branching pomset for our DV protocol: where we would traditionally need $2 \times 2 \times 2$ pomsets (with six pairs of events each), we can represent the same behaviour as a single branching pomset with $2 + 2 + 2$ choices (with four pairs of events each). Adding an additional participant would double the number of pomsets in the set of pomsets, while it would add a single choice to the branching pomset.

To aid in the understanding of branching pomsets and their semantics, we provide a prototype tool to visualise them, available at <https://arca.di.uminho.pt/b-pomset/>. The tool provides a web interface where one can submit an input choreography, which is then visualised as a branching pomset and can be simulated. The examples and figures in the paper are already available as preset inputs. We note that the pomset simulation in our prototype currently does not support loops, for reasons which will become apparent later in the paper; however, all other operators are supported and we are most interested in (combinations of) choice and parallel composition.

Contribution This paper provides three core contributions: (1) an extension of pomsets with a branching structure, named branching pomsets, (2) an encoding from a choreographic language into branching pomsets, and (3) a formal proof that the operational semantics of a choreography and of its encoded branching pomset are equivalent, i.e., bisimilar.

$$c ::= \mathbf{0} \mid a \rightarrow b : x \mid \boxed{ab?x} \mid c ; c \mid c + c \mid c \parallel c \mid c^*$$

Figure 3: Syntax of choreographies, where a and b are participants ($a \neq b$) and x is a message type.

Structure of the paper Section 2 presents the syntax of our choreography language and its operational semantics. Section 3 formalises branching pomsets and their semantics. Section 4 formalises how to obtain a branching pomset from a choreography and shows that a choreography and its derived branching pomset are behaviourally equivalent. Finally, Section 5 presents our conclusions and a brief discussion about future work and related work.

2 Choreographies

In this section, we formally define the syntax and semantics of our choreographic language, examples of which have been shown in the previous section.

Let \mathcal{A} be the set of all participants a, b, \dots . Let \mathcal{X} be the set of all message types x, y, \dots . Let $\mathcal{L} = \bigcup_{a, b \in \mathcal{A}, x \in \mathcal{X}} \{ab!x, ab?x\}$ be the set of actions $\{ab!x, ab?x\}$ for all participants $a \neq b$ and message types x . For all actions the *subject* of the action is its active participant: the subject of a send action $ab!x$ is a , written $subj(ab!x) = a$, and the subject of a receive action $ab?x$ is b .

The syntax is formally defined in Figure 3. Its components are standard: ‘ $\mathbf{0}$ ’ is the empty choreography; ‘ $a \rightarrow b : x$ ’ is the asynchronous communication from a to b of a message of type x ; the boxed term ‘ $ab?x$ ’ represents a pending receive on b from a of a message of type x (it is boxed in Figure 3 to indicate that it is only used internally to formalise behaviour but the box is not part of the syntax); ‘ $c_1 ; c_2$ ’, ‘ $c_1 + c_2$ ’ and ‘ $c_1 \parallel c_2$ ’ are respectively the weak sequential composition, nondeterministic choice and parallel composition of choreographies c_1 and c_2 ; finally, ‘ c^* ’ is the finite repetition (or, more informally, loop) of choreography c . The semantics for choice, parallel composition and loop are standard. We note that our sequential composition is weak. More traditionally, when sequencing c_1 and c_2 , the choreography c_1 must fully terminate before proceeding to c_2 . With weak sequential composition, however, actions in c_2 can already be executed as long as they do not interfere with c_1 . For example, in $a \rightarrow b : x ; c \rightarrow d : x$ we can execute the action $cd!x$ as it does not affect the participants of $a \rightarrow b : x$: there is no dependency and thus no need to wait for $a \rightarrow b : x$ to go first. However, in $a \rightarrow b : x ; a \rightarrow c : x$ the action $ac!x$ cannot be executed first as its subject (a) must first execute $ab!x$. This is the common interpretation of sequential composition in the context of message sequence charts [12], multiparty session types [9] and choreographic programming [2].

The reduction rules of our choreographic language are formally defined in Figure 4a and its termination rules in Figure 4b. To formalise the reduction of weak sequential composition, we follow Rensink and Wehrheim [15], who define a notion of *partial termination*.

Partial termination In a weak sequential composition $c_1 ; c_2$, an action ℓ in c_2 can be executed if c_1 can *partially terminate* for the subject of ℓ . Conceptually, a choreography c_1 can partially terminate for the subject of ℓ by discarding all branches of its behaviour which would conflict with it, i.e., in which the subject of ℓ occurs. This is written $c_1 \xrightarrow{\ell} c'_1$, where c'_1 is the remainder of c_1 after discarding all branches involving the subject of ℓ . For example, if $c_1 = a \rightarrow b : x + a \rightarrow c : x$ then $c_1 \xrightarrow{cd!x} a \rightarrow b : x$, as this branch does not contain c . An exception is when the subject of ℓ occurs in *every* branch of c_1 , in which case c_1 cannot partially terminate for the subject of ℓ , i.e., $c_1 \not\xrightarrow{\ell}$. In the above example, $c_1 \not\xrightarrow{ad!x}$.

$$\begin{array}{c}
\frac{}{a \rightarrow b : x \xrightarrow{ab!x} ab?x} \quad \frac{}{ab?x \xrightarrow{ab?x} \mathbf{0}} \quad \frac{c_1 \xrightarrow{\ell} c'_1}{c_1 ; c_2 \xrightarrow{\ell} c'_1 ; c_2} \quad \frac{c_1 \xrightarrow{\check{\ell}} c'_1 \quad c_2 \xrightarrow{\ell} c'_2}{c_1 ; c_2 \xrightarrow{\ell} c'_1 ; c'_2} \\
\frac{c_1 \xrightarrow{\ell} c'_1}{c_1 \parallel c_2 \xrightarrow{\ell} c'_1 \parallel c_2} \quad \frac{c_2 \xrightarrow{\ell} c'_2}{c_1 \parallel c_2 \xrightarrow{\ell} c_1 \parallel c'_2} \quad \frac{c_1 \xrightarrow{\ell} c'_1}{c_1 + c_2 \xrightarrow{\ell} c'_1} \quad \frac{c_2 \xrightarrow{\ell} c'_2}{c_1 + c_2 \xrightarrow{\ell} c_2} \quad \frac{c \xrightarrow{\ell} c'}{c^* \xrightarrow{\ell} c' ; c^*}
\end{array}$$

(a) Reduction rules.

$$\frac{}{\mathbf{0} \downarrow} \quad \frac{}{c^* \downarrow} \quad \frac{c_1 \downarrow \quad c_2 \downarrow \quad \dagger \in \{;, \parallel\}}{c_1 \dagger c_2 \downarrow} \quad \frac{c_i \downarrow \quad i \in \{1, 2\}}{c_1 + c_2 \downarrow}$$

(b) Termination rules.

$$\frac{}{\mathbf{0} \xrightarrow{\check{\ell}} \mathbf{0}} \quad \frac{c \xrightarrow{\check{\ell}} c}{c^* \xrightarrow{\check{\ell}} c^*} \quad \frac{c \not\xrightarrow{\check{\ell}} c}{c^* \xrightarrow{\check{\ell}} \mathbf{0}} \quad \frac{c_1 \xrightarrow{\check{\ell}} c'_1 \quad c_2 \xrightarrow{\check{\ell}} c'_2 \quad \dagger \in \{;, \parallel, +\}}{c_1 \dagger c_2 \xrightarrow{\check{\ell}} c'_1 \dagger c'_2} \\
\frac{c_1 \xrightarrow{\check{\ell}} c'_1 \quad c_2 \not\xrightarrow{\check{\ell}}}{c_1 + c_2 \xrightarrow{\check{\ell}} c'_1} \quad \frac{c_1 \not\xrightarrow{\check{\ell}} \quad c_2 \xrightarrow{\check{\ell}} c'_2}{c_1 + c_2 \xrightarrow{\check{\ell}} c'_2} \quad \frac{\text{subj}(\ell) \notin \{a, b\}}{a \rightarrow b : x \xrightarrow{\check{\ell}} a \rightarrow b : x} \quad \frac{\text{subj}(\ell) \neq b}{ab?x \xrightarrow{\check{\ell}} ab?x}$$

(c) Partial termination rules.

Figure 4: Operational semantics of choreographies.

The rules for partial termination are deterministic and only discard the absolutely necessary. In the example above, $c_1 \xrightarrow{da!x} c_1$ since the subject d does not occur in either branch: dropping one of the branches would be unnecessary and is thus not allowed. The rules for partial termination are defined in Figure 4c. We highlight the rules for operators:

- Sequential composition $c_1 ; c_2$ and parallel composition $c_1 \parallel c_2$ can partially terminate if both c_1 and c_2 can.
- A choice $c_1 + c_2$ can partially terminate if at least one of its branches can. If both branches can partially terminate then both are kept, otherwise only the partially terminated one is kept.
- Following Rensink and Wehrheim, a loop c^* can partially terminate if its body (c) can partially terminate without discarding any branches, i.e., if $c \xrightarrow{\check{\ell}} c$. In that case also $c^* \xrightarrow{\check{\ell}} c^*$. Otherwise we allow c^* to be skipped entirely, represented as partial termination to $\mathbf{0}$, i.e., $c^* \xrightarrow{\check{\ell}} \mathbf{0}$. This can happen either if c can partially terminate to c' but $c' \neq c$, or if c cannot partially terminate at all. We use $c \not\xrightarrow{\check{\ell}} c$ as a shorthand to cover both these cases. Skipping a loop is necessary, for example, in a modified master-workers protocol where the master can send an arbitrary number of tasks to the workers, followed by an `end` message to indicate termination. With one worker, this protocol is expressed as $(m \rightarrow w_1 : t ; w_1 \rightarrow m : d)^* ; m \rightarrow w_1 : \text{end}$. In this choreography, the loop has to eventually partially terminate to $\mathbf{0}$ to allow for the action $mw_1 ! \text{end}$.

Example 1. Let $c_1 = (a \rightarrow b:x + a \rightarrow c:x); (d \rightarrow b:x + d \rightarrow e:x)$. Let $c_2 = (a \rightarrow b:x + c \rightarrow b:x)^* \parallel (c \rightarrow a:x + c \rightarrow b:x)$.

- $c_1 \xrightarrow{\sqrt{be!x}} a \rightarrow c:x; d \rightarrow e:x$. The subject b of $be!x$ occurs in one branch of each of both choices. While the recipient e also occurs in the second branch of the second choice, since it is not the actual subject it does not create a conflict.
- $c_1 \not\xrightarrow{\sqrt{ab!x}}$. While the second choice can partially terminate without reducing, the first choice contains the subject a of $ab!x$ in both of its branches. Since one of the choices cannot reduce, neither can their sequential composition.
- $c_2 \xrightarrow{\sqrt{ad!x}} \mathbf{0} \parallel c \rightarrow b:x$. The subject a of $ad!x$ only occurs in one branch of the loop body, but the loop can only reduce to $\mathbf{0}$. On the right hand side of the parallel composition, a occurs only in the first branch.
- $c_2 \not\xrightarrow{\sqrt{cd!x}}$. While the loop can again reduce to $\mathbf{0}$, the subject c of $cd!x$ occurs in both branches of the right hand side of the parallel composition. Since its right hand side cannot partially terminate, neither can it as a whole.

As already discovered by Rensink and Wehrheim [15], an unwanted consequence of these rules for partial termination is that unfolding iterations of loops no longer preserves behaviour. We would like c^* and $(c; c^*) + \mathbf{0}$ to behave the same, but this is not the case. For example, if $c = a \rightarrow b:x + c \rightarrow d:x$, then $c^* \xrightarrow{\sqrt{ab!x}} \mathbf{0}$ but $(c; c^*) + \mathbf{0} \xrightarrow{\sqrt{ab!x}} (c \rightarrow d:x; \mathbf{0}) + \mathbf{0}$. Then $c^* \xrightarrow{ab!x} ab?x$ by skipping the loop; however, $((c; c^*) + \mathbf{0}); c$ has no way to match this as it can skip the loop but it can only reduce the already unfolded iteration c to $c \rightarrow d:x$ — it cannot discard it entirely. We borrow the solution that Rensink and Wehrheim offer, which is the concept *dependent guardedness*.

Dependent guardedness A loop c^* is *dependently guarded* if, for all actions ℓ , the loop body c can only partially terminate for the subject of ℓ if it does not occur in c at all. In other words: any participant that occurs in some branch of c must also occur in every other branch of c . It then follows that c can either partially terminate for the subject of ℓ without having to reduce, or it cannot partially terminate at all. Formally: if $c \xrightarrow{\sqrt{\ell}} c'$ then $c' = c$. A choreography \hat{c} is then dependently guarded if all of its loops are.

As a consequence, we avoid the problem above: if $c^* \xrightarrow{\sqrt{\ell}} \mathbf{0}$ then $c \not\xrightarrow{\sqrt{\ell}}$ and $(c; c^*) + \mathbf{0}$ is also forced to reduce to the second branch of the choice, which is $\mathbf{0}$. More precisely, let c^* be some dependently guarded expression. If $c \xrightarrow{\sqrt{\ell}} c'$ for some ℓ, c' , then $c' = c$. It follows that $c^* \xrightarrow{\sqrt{\ell}} c^*$ and $(c; c^*) + \mathbf{0} \xrightarrow{\sqrt{\ell}} (c; c^*) + \mathbf{0}$. Similarly, if $c \not\xrightarrow{\sqrt{\ell}}$ then $c^* \xrightarrow{\sqrt{\ell}} \mathbf{0}$ and $(c; c^*) + \mathbf{0} \xrightarrow{\sqrt{\ell}} \mathbf{0}$.

Example 2. Let $c_1 = a \rightarrow b:x + a \rightarrow c:x$. Let $c_2 = a \rightarrow b:x + b \rightarrow a:x$.

- c_1^* is not dependently guarded as $c_1 \xrightarrow{\sqrt{cd!x}} a \rightarrow b:x \neq c_1$. However, c_1 itself is dependently guarded as it does not contain any loop.
- c_2^* is dependently guarded since both a and b occur in both branches of c_2 . However, $(c_2^*)^*$ is *not* dependently guarded, since $c_2^* \xrightarrow{\sqrt{ab!x}} \mathbf{0}$.

3 Branching pomsets

In this section, we formally define the syntax and semantics of branching pomsets. Additionally, we define a pomset interpretation of expressions in our choreographic language and we show this interpretation to be faithful by proving that it is bisimilar to the original choreography.




Figure 5: A branching pomset representing the choreography $a \rightarrow b:x; (b \rightarrow c:x + b \rightarrow d:x); c \rightarrow d:x$.

A partially ordered multiset [14], or pomset for short, consists of a set of nodes E (events), a labelling function λ to map events to some set of labels (e.g., send and receive actions), and a partial order \leq to define dependencies between pairs of events (e.g., an event, or rather its corresponding action, can only fire if all events preceding it in the partial order have already fired). Its behaviour is the set of all sequences of the labels of its events that abide by \leq .

For example, for the pomset in Figure 1, $E = \{e_1, \dots, e_8\}$, $\lambda = \{e_1 \mapsto mw_1!t, e_2 \mapsto mw_1?t, e_3 \mapsto w_1m!d, e_4 \mapsto w_1m?d, e_5 \mapsto mw_2!t, e_6 \mapsto mw_2?t, e_7 \mapsto w_2m!d, e_8 \mapsto w_2m?d\}$, and $\leq = \{(e_i, e_j) \mid (i, j \in [1, 4] \vee i, j \in [5, 8]) \wedge i \leq j\}$. Its behaviour consists of all interleavings of $mw_1!t mw_1?t w_1m!d w_1m?d$ and $mw_2!t mw_2?t w_2m!d w_2m?d$.

As illustrated in Figure 2, however, traditional pomsets suffer from the same problem when representing choices that automata suffer from when representing concurrency: there is no explicit representation of choices in pomsets, and they are represented only implicitly as a set of possible pomsets. We tackle this by extending pomsets with an explicit representation of choices: a branching structure on events.

Branching structure The general idea of a branching pomset is that all possible events are initially part of the pomset, but some are defined as being part of a choice. To fire these, all relevant choices must first be resolved by replacing the choice with one of its branches, thereby discarding the other branch. This same idea governs the operational semantics of choreographies (Figure 4): both branches of a choice are initially part of the choreography but, to proceed in one of them, the other must be dropped.

The branching structure does not interrupt the partial order and all events still participate in it, as shown in Figure 5, where arrows flow both into and out of the branches of the choice. As such, a choice can also be resolved to fire an event which is only preceded by one of the branches, reminiscent of the partial termination of choices (Figure 4c). For example, in Figure 5 the upper branch ($b \rightarrow c:x$) can be discarded to fire the event labelled $cd!x$, as it is not dependent on the lower branch. As shown in Figure 6, nested choices are supported as well.

Formally, the branching structure is defined below as a tree with root node \mathcal{B} , whose children are either a single event e or a choice node \mathcal{C} with children (branches) $\mathcal{B}_1, \mathcal{B}_2$. All leaves are events.

$$\begin{aligned} \mathcal{B} &::= \{\mathcal{C}_1, \dots, \mathcal{C}_n\} \\ \mathcal{C} &::= e \mid \{\mathcal{B}_1, \mathcal{B}_2\} \end{aligned}$$

For example, for the pomset in Figure 5, if $E = \{e_1, \dots, e_8\}$ and $\lambda = \{e_1 \mapsto ab!x, e_2 \mapsto ab?x, e_3 \mapsto cd!x, e_4 \mapsto cd?x, e_5 \mapsto bc!x, e_6 \mapsto bc?x, e_7 \mapsto bd!x, e_8 \mapsto bd?x\}$, then its branching structure is $\{e_1, e_2, e_3, e_4, \{\{e_5, e_6\}, \{e_7, e_8\}\}\}$. For the pomset in Figure 6, if $E = \{e_1, \dots, e_{14}\}$ and $\lambda = \{e_1 \mapsto ab!x, e_2 \mapsto ab?x, e_3 \mapsto ba!x, e_4 \mapsto ba?x, e_5 \mapsto bd!x, e_6 \mapsto bd?x, e_7 \mapsto ac!x, e_8 \mapsto ac?x, e_9 \mapsto ca!x, e_{10} \mapsto ca?x, e_{11} \mapsto cd!x, e_{12} \mapsto cd?x, e_{13} \mapsto da!x, e_{14} \mapsto da?x\}$, then its branching structure is $\{e_{13}, e_{14}, \{\{e_1, e_2, \{\{e_3, e_4\}, \{e_5, e_6\}\}\}, \{e_7, e_8, \{\{e_9, e_{10}\}, \{e_{11}, e_{12}\}\}\}\}\}$. By resolving the outer choice and picking its upper branch ($a \rightarrow b:x$), we drop events e_7, \dots, e_{12} and obtain the middle branching pomset in Figure 8, with events $e_1, \dots, e_6, e_{13}, e_{14}$ and branching structure $\{e_1, e_2, e_{13}, e_{14}, \{\{e_3, e_4\}, \{e_5, e_6\}\}\}$.




Figure 6: A branching pomset representing the choreography $((a \rightarrow b:x; (b \rightarrow a:x + b \rightarrow d:x)) + (a \rightarrow c:x; (c \rightarrow a:x + c \rightarrow d:x))) ; d \rightarrow a:x$.

We now formally define branching pomsets.

Definition 1 (Branching pomset). A branching pomset is a four-tuple $R = \langle E, \leq, \lambda, \mathcal{B} \rangle$, where E is a set of events, $\leq \subseteq E \times E$ is such that \leq^* (the transitive closure of \leq) is a partial order on events, $\lambda : E \mapsto \mathcal{L}$ is a labelling function assigning an action to every event, and \mathcal{B} is a branching structure such that the set of leaves of \mathcal{B} is E and no event in E occurs in \mathcal{B} more than once. We use $R.E$, $R.\leq$, $R.\lambda$ and $R.\mathcal{B}$ to refer to the components of R .

Semantics To fire an event in a branching pomset, on top of being minimal it must also be *active*, i.e., it must not be inside any choice. In other words: it must be a child of the branching structure's root node. We thus define a set of refinement rules in Figure 7a, written $R \sqsupseteq R'$, which can be used to resolve choices and move events upwards in the branching structure.

The first two rules, REFL and TRANS, are straightforward. The third rule, CHOICE, resolves choices. It states that we can replace a choice with one of its branches. This rule serves a dual purpose: by applying it to the outer choice of the pomset in Figure 6 we can fire the event $ab!x$ in its first branch; alternatively, by applying it to the pomset in Figure 5 we can discard one branch of the choice and then fire the event $cd!x$, which is now minimal. The latter use corresponds with the partial termination rules for choreographies. The fourth rule, CONGR is used for more fine-grained partial termination. To make the event $da!x$ minimal in Figure 6 we could resolve two choices with CHOICE (and TRANS). However, as the rules for partial termination tell us, it is unnecessary to resolve the outer choice. Instead, we can apply CHOICE to both inner choices and apply CONGR to the outer choice to update it without unnecessarily resolving it. Finally, the fifth rule overloads the refinement notation to also apply to branching pomsets themselves: if $R.\mathcal{B}$ can refine to some \mathcal{B}' , then R itself can refine to a derived branching pomset with branching structure \mathcal{B}' , whose events are restricted to those occurring in \mathcal{B}' and likewise for \leq and λ .

The reduction and termination rules are defined in Figure 7b. The first rule simply states that a pomset can terminate if its branching structure can reduce to the empty set. The second rule defines the conditions for *enabling* an event e , written $R \xrightarrow{e} R'$. A branching pomset R can enable e by refining to R' if e is both minimal and active in R' ($e \in \text{a-min}(R')$), and if there is no other refinement in between in which e is already minimal and active. In other words, R may only refine as far as strictly necessary to enable e . This rule implements the same idea as partial termination, with the subtle difference that, whereas partial termination tries to remove any occurrence of a participant, in this case e is actually

$$\begin{array}{c} \overline{\mathcal{B} \sqsupseteq \mathcal{B}}^{\text{[REFL]}} \quad \frac{\mathcal{B} \sqsupseteq \mathcal{B}' \sqsupseteq \mathcal{B}''}{\mathcal{B} \sqsupseteq \mathcal{B}''}^{\text{[TRANS]}} \quad \frac{i \in \{1,2\}}{\{\{\mathcal{B}_1, \mathcal{B}_2\}\} \cup \mathcal{B} \sqsupseteq \mathcal{B}_i \cup \mathcal{B}}^{\text{[CHOICE]}} \\ \frac{\mathcal{B}_1 \sqsupseteq \mathcal{B}'_1 \quad \mathcal{B}_2 \sqsupseteq \mathcal{B}'_2}{\{\{\mathcal{B}_1, \mathcal{B}_2\}\} \cup \mathcal{B} \sqsupseteq \{\{\mathcal{B}'_1, \mathcal{B}'_2\}\} \cup \mathcal{B}}^{\text{[CONGR]}} \quad \frac{R.\mathcal{B} \sqsupseteq \mathcal{B}'}{R \sqsupseteq R[\mathcal{B}']} \end{array}$$

(a) Refinement rules, where we assume for CHOICE and CONGR that $\{\mathcal{B}_1, \mathcal{B}_2\} \notin \mathcal{B}$.

$$\frac{R \sqsupseteq R' \quad e \in \text{a-min}(R')}{R.\mathcal{B} \sqsupseteq \emptyset} \quad \frac{\forall R'' : R \sqsupseteq R'' \sqsupseteq R' \Rightarrow e \notin \text{a-min}(R'')}{R \xrightarrow{e} R'} \quad \frac{R \xrightarrow{e} R'}{R \xrightarrow{e} R' - e} \quad \frac{R \xrightarrow{e} R'}{R \xrightarrow{\lambda(e)} R'}$$

(b) Reduction and termination rules.

$$\begin{array}{l} \langle E, \leq, \lambda, \mathcal{B} \rangle[\mathcal{B}'] = \langle E|_{\mathcal{B}'}, \leq|_{\mathcal{B}'}, \lambda|_{\mathcal{B}'}, \mathcal{B}' \rangle \\ X|_{\mathcal{B}} = \text{restricts } X \text{ only to the events in } \mathcal{B} \\ \text{a-min}(R) = \{e \in R.E \mid \nexists e' \in R.E : e' < e\} \wedge e \in R.\mathcal{B} \\ \hat{e} - e = \hat{e} \\ \{\mathcal{C}_1, \dots, \mathcal{C}_n\} - e = \begin{cases} \{\mathcal{C}_1, \dots, \mathcal{C}_{i-1}, \mathcal{C}_{i+1}, \dots, \mathcal{C}_n\} & \text{if } \mathcal{C}_i = e \\ \{\mathcal{C}_1 - e, \dots, \mathcal{C}_n - e\} & \text{otherwise} \end{cases} \\ \{\mathcal{B}_1, \mathcal{B}_2\} - e = \{\mathcal{B}_1 - e, \mathcal{B}_2 - e\} \\ R - e = R[R.\mathcal{B} - e] \end{array}$$


(c) Operations on branching pomsets.

Figure 7: Semantics of branching pomsets.


an event in R itself. As the two notions are very similar, we use the same notation for enabling events in branching pomsets as for partial termination. Finally, the last two rules state that, if R can enable e by refining to R' , then it can fire e by reducing to $R' - e$, which is the branching pomset obtained by removing e from R' (Figure 7c). This reduction is defined both on e 's label and on the event itself, the latter for internal use in proofs since $\lambda(e)$ is typically not unique but e is.

Example 3.


- $R \xrightarrow{e} R'$, where R is the branching pomset in Figure 5, R' is the topmost branching pomset in Figure 8 and e is the event with label $cd!x$.
- $R \xrightarrow{e} R'$, where R is the branching pomset in Figure 6, R' is the middle branching pomset in Figure 8 and e is the event with label $ab!x$.
- $R \xrightarrow{e} R'$, where R is the branching pomset in Figure 6, R' is the middle branching pomset in Figure 8 and e is the event with label $da!x$.



Obtained by applying CHOICE to the pomset in Figure 5.



Obtained by applying CHOICE to the outer choice of the pomset in Figure 6.



Obtained by applying CONGR to the outer and CHOICE to both inner choices of the pomset in Figure 6.

Figure 8: Three refined pomsets.

4 Branching pomsets for choreographies

In this section we formalise the construction of a branching pomset for a choreography c and we show that the pomset semantics for the branching pomset are bisimilar to the operational semantics for c .

We have given examples of choreographies and corresponding branching pomsets in Figures 5 and 6. Formally, the rules for the construction of a branching pomset for a choreography c , written $\llbracket c \rrbracket$, are defined in Figure 9. Most rules are as expected. We highlight the rules for operators.

- The rule for parallel composition ($\llbracket c_1 \parallel c_2 \rrbracket$) takes the pairwise union of all components.
- The rule for sequential composition ($\llbracket c_1 ; c_2 \rrbracket$) also adds dependencies to ensure that, for every a , all events with subject a in $\llbracket c_1 \rrbracket$ (denoted E_{1_a}) must precede all events with subject a in $\llbracket c_2 \rrbracket$. This matches the reduction rule for weak sequential composition of choreographies (Figure 4a), as events in $\llbracket c_2 \rrbracket$ are only required to wait for events in $\llbracket c_1 \rrbracket$ whose subject is the same.
- The rule for choice ($\llbracket c_1 + c_2 \rrbracket$) adds a single top-level choice in the branching structure to choose between the pomsets for c_1 and c_2 .
- The rule for loops ($\llbracket c^* \rrbracket$) encodes a loop as a choice between terminating ($\mathbf{0}$) and unfolding one iteration of the loop ($c ; c^*$). This results in a pomset of infinite size. We note that our theoretical results still hold even on infinite pomsets, but that any analysis of an infinite pomset will have to

$$\begin{aligned}
\llbracket \mathbf{0} \rrbracket &= \langle \emptyset, \emptyset, \emptyset, \emptyset \rangle \\
\llbracket a \rightarrow b : x \rrbracket &= \langle \{e_1, e_2\}, \{e_1 \leq e_1, e_1 \leq e_2, e_2 \leq e_2\}, \{e_1 \mapsto ab!x, e_2 \mapsto ab?x\}, \{e_1, e_2\} \rangle \\
\llbracket ab?x \rrbracket &= \langle \{e\}, \{e \leq e\}, \{e \mapsto ab?x\}, \{e\} \rangle \\
\llbracket c_1 \dagger c_2 \rrbracket &= \llbracket c_1 \rrbracket \dagger \llbracket c_2 \rrbracket \text{ for } \dagger \in \{;, +, \parallel\} \\
\llbracket c^* \rrbracket &= \llbracket (c; c^*) + \mathbf{0} \rrbracket \\
R_1 \parallel R_2 &= \langle E_1 \cup E_2, \leq_1 \cup \leq_2, \lambda_1 \cup \lambda_2, \mathcal{B}_1 \cup \mathcal{B}_2 \rangle \\
R_1 ; R_2 &= \langle E_1 \cup E_2, \leq_1 \cup \leq_2 \cup \bigcup_{a \in \mathcal{A}} E_{1_a} \times E_{2_a}, \lambda_1 \cup \lambda_2, \mathcal{B}_1 \cup \mathcal{B}_2 \rangle \\
R_1 + R_2 &= \langle E_1 \cup E_2, \leq_1 \cup \leq_2, \lambda_1 \cup \lambda_2, \{\{\mathcal{B}_1, \mathcal{B}_2\}\} \rangle
\end{aligned}$$

Figure 9: Pomset interpretation of choreographies, where $R_i = \langle E_i, \leq_i, \lambda_i, \mathcal{B}_i \rangle$ for $i \in \{1, 2\}$, \mathcal{A} is the set of all participants (a, b, \dots) and E_{i_a} is the subset of events in E_i with subject a .

be symbolic. However, since the focus of this paper is on supporting choices, we do not discuss this further and leave symbolic analyses for loops for future work.

As an example, we construct part of the branching pomset in Figure 5: $(b \rightarrow c : x + b \rightarrow d : x); c \rightarrow d : x$ (thus omitting $a \rightarrow b : x$). Let $\llbracket b \rightarrow c : x \rrbracket = \langle \{e_1, e_2\}, \leq_1, \lambda_1, \{e_1, e_2\} \rangle$, $\llbracket b \rightarrow d : x \rrbracket = \langle \{e_3, e_4\}, \leq_2, \lambda_2, \{e_3, e_4\} \rangle$ and $\llbracket c \rightarrow d : x \rrbracket = \langle \{e_5, e_6\}, \leq_3, \lambda_3, \{e_5, e_6\} \rangle$ as in Figure 9. First, $\llbracket b \rightarrow c : x + b \rightarrow d : x \rrbracket = \langle \{e_1, \dots, e_4\}, \leq_1 \cup \leq_2, \lambda_1 \cup \lambda_2, \{\{\{e_1, e_2\}, \{e_3, e_4\}\}\} \rangle$; this is the pairwise union of the first three components, with the branching structure adding a choice between the two branches. Then $\llbracket (b \rightarrow c : x + b \rightarrow d : x); c \rightarrow d : x \rrbracket = \langle \{e_1, \dots, e_6\}, \leq_1 \cup \leq_2 \cup \leq_3 \cup \{e_2 \leq e_5, e_4 \leq e_6\}, \lambda_1 \cup \lambda_2 \cup \lambda_3, \{e_5, e_6, \{\{e_1, e_2\}, \{e_3, e_4\}\}\} \rangle$; again, this is the pairwise union of all components, with the addition of two dependencies: $e_2 \leq e_5$ represents the arrow in Figure 5 from $bc?x$ to $cd!x$ as they both have subject c , $e_4 \leq e_6$ represents the arrow from $bd?x$ to $cd?x$ as they both have subject d . There are no direct dependencies between e_1 ($bc!x$) or e_3 ($bd!x$) and either e_5 or e_6 , as the latter two do not have subject b .

Bisimulation For any given a choreography c we can derive two labelled transition systems: one from the operational semantics in Figure 4 over c , and one from the pomset semantics in Figure 7 over the branching pomset $\llbracket c \rrbracket$ produced by the rules in Figure 9. In the remainder of this section we show that the two transition systems are bisimilar.

Two systems are language equivalent (or trace equivalent) if their languages are the same, i.e., if they accept the same set of words (or traces). Two systems are bisimilar if each of them can simulate the other, i.e., if they cannot be distinguished from each other just by looking at their behaviour. This is a stronger notion of equivalence than language equivalence: if two systems are bisimilar then they are also language equivalent, but the inverse is not necessarily true.

Example 4.

- $a \rightarrow b : x; (b \rightarrow a : x + b \rightarrow a : y)$ is language equivalent but not bisimilar to $(a \rightarrow b : x; b \rightarrow a : x) + (a \rightarrow b : x; b \rightarrow a : y)$. In the former the choice between $b \rightarrow a : x$ and $b \rightarrow a : y$ is made only after $a \rightarrow b : x$, while in the latter the choice is made up front. As a result, it is possible in the latter system to fire $ab!x; ab?x$ and then end up in a state where $ba!x$ cannot be fired because the branch with $b \rightarrow a : y$ was chosen — or the other way around; in the former system it is always possible to fire both $ba!x$ and $ba!y$.

- $a \rightarrow b : x$ is bisimilar to $a \rightarrow b : x + a \rightarrow b : x$. While the latter contains a choice, the two systems cannot be distinguished by their behaviour. In both cases, the only allowed action is $ab!x$ and then $ab?x$.

Formally, two transition systems A_1, A_2 are bisimilar, written $A_1 \sim A_2$, if there exists a bisimulation relation \mathcal{R} relating the states of A_1 and A_2 which relates their initial states [16]. The relation \mathcal{R} is a bisimulation relation if, for every pair of states $\langle p, q \rangle \in \mathcal{R}$:

- If $p \xrightarrow{\ell} p'$ then $q \xrightarrow{\ell} q'$ and $\langle p', q' \rangle \in \mathcal{R}$ for some q' , and vice-versa.
- If $p \downarrow$ then $q \downarrow$, and vice-versa.

In other words: if one of the two can perform a step, then the other can perform a matching step such that the resulting states are again in the bisimulation relation.

This is also the approach we follow when proving that $c \sim \llbracket c \rrbracket$ for all (dependently guarded) choreographies c : we define a relation $\mathcal{R} = \{ \langle c, \llbracket c \rrbracket \rangle \mid c \text{ is a dependently guarded choreography} \}$ relating all dependently guarded choreographies with their interpretation as branching pomset by the rules in Figure 9. We then show that:

- If $c \xrightarrow{\ell} c'$ then $\llbracket c \rrbracket \xrightarrow{\ell} \llbracket c' \rrbracket$ (Lemma 2).
- If $\llbracket c \rrbracket \xrightarrow{\ell} R'$ then $c \xrightarrow{\ell} c'$ such that $R' = \llbracket c' \rrbracket$ (Lemma 3).
- If $c \downarrow$ then $\llbracket c \rrbracket \downarrow$ (Lemma 4).
- If $\llbracket c \rrbracket \downarrow$ then $c \downarrow$ (Lemma 5).

Together these lemmas prove that $c \sim \llbracket c \rrbracket$ for all dependently guarded c (Theorem 6). Most of the proofs are straightforward by structural induction on c . Of particular interest, however, are the two reduction lemmas in the case of weak sequential composition, i.e., if $c_1 ; c_2 \xrightarrow{\ell} c'_1 ; c'_2$ in Lemma 2 and if $\llbracket c_1 ; c_2 \rrbracket \xrightarrow{e} R'$ where e is an event in $\llbracket c_2 \rrbracket$ in Lemma 3. To prove these specific cases we need to show a correspondence between partial termination and enabling events. We do this with Lemma 1, in which we show two directions simultaneously. If the choreography c_1 can partially terminate for the subject of an action ℓ in c_2 then the branching pomset $\llbracket c_1 ; c_2 \rrbracket$ can enable the corresponding event. Conversely, if $\llbracket c_1 ; c_2 \rrbracket$ can enable some event in $\llbracket c_2 \rrbracket$ then the choreography c_1 can partially terminate for the subject of its label. When proving these cases in Lemmas 2 and 3, we then only have to show that the preconditions of Lemma 1 hold.

In the following, a number of technical lemmas and most of the proofs are omitted in favour of informal proof sketches or highlights. The omitted proofs and technical lemmas can be found in a separate technical report [7].

Lemma 1. *Let c_1 and c_2 be dependently guarded choreographies. Let $c_2 \xrightarrow{\ell} c'_2$ and $\llbracket c_2 \rrbracket \xrightarrow{e} R'_2$ such that $\lambda(e) = \ell$ and $\llbracket c'_2 \rrbracket = R'_2 - e$.*

- If $c_1 \xrightarrow{\ell} c'_1$ then $\llbracket c_1 ; c_2 \rrbracket \xrightarrow{e} \llbracket c'_1 \rrbracket ; R'_2$.*
- If $\llbracket c_1 ; c_2 \rrbracket \xrightarrow{e} R'_1 ; R'_2$ then $c_1 \xrightarrow{\lambda(e)} c'_1$ and $\llbracket c'_1 \rrbracket = R'_1$.*

Proof sketch. This proof is by structural induction on c_1 . Although the details require careful consideration, it is conceptually straightforward: every case in (a) consists of showing that e is minimal and active in $\llbracket c'_1 \rrbracket ; R'_2$ and that $\llbracket c'_1 \rrbracket ; R'_2$ is the first refinement for which this is true, and then applying the second rule in Figure 7b; every case in (b) consists of showing that $\llbracket c_3 ; c_2 \rrbracket \xrightarrow{e} \llbracket c'_3 \rrbracket ; R'_2$ for some subexpression c_3 of c_1 and similarly for c_4 (e.g., when $c_1 = c_3 + c_4$), then applying the induction hypothesis (b) to obtain $c_3 \xrightarrow{\ell} c'_3$ and $c_4 \xrightarrow{\ell} c'_4$, and finally applying the partial termination rules in Figure 4c. \square

Lemma 2. *Let c be a dependently guarded choreography. If $c \xrightarrow{\ell} c'$ then $\llbracket c \rrbracket \xrightarrow{\ell} \llbracket c' \rrbracket$.*

Proof sketch. This proof is by structural induction on c . We note that, if $c = c_1 ; c_2$ and $c' = c'_1 ; c'_2$, i.e., when partial termination is applied, then the premises of Lemma 1 hold by the induction hypothesis and the result swiftly follows. All other cases are straightforward. \square

Lemma 3. *Let c be a dependently guarded choreography. If $\llbracket c \rrbracket \xrightarrow{\ell} R'$ for some R' then $c \xrightarrow{\ell} c'$ such that $R' = \llbracket c' \rrbracket$.*

Proof sketch. This proof is by structural induction on c . We highlight two cases:

- If $c = c_1^*$ then we use a technical lemma to show that $R' = R'_1 ; \llbracket c_1^* \rrbracket$ such that $\llbracket c_1 \rrbracket \xrightarrow{\ell} R'_1$. It then follows from the induction hypothesis that $c_1 \xrightarrow{\ell} c'_1$ such that $\llbracket c'_1 \rrbracket = R'_1$. The remainder is straightforward.
- If $c = c_1 ; c_2$ then $\llbracket c \rrbracket = \llbracket c_1 \rrbracket ; \llbracket c_2 \rrbracket$. If e is an event in $\llbracket c_2 \rrbracket$ then we proceed to show that $\llbracket c_2 \rrbracket \xrightarrow{\ell} R'_2$, at which point we can apply the induction hypothesis. We have then satisfied the premises of Lemma 1. The remainder is straightforward.

All other cases are straightforward. \square

Lemma 4. *Let c be a dependently guarded choreography. If $c \downarrow$ then $\llbracket c \rrbracket \downarrow$.*

Proof sketch. This proof is by structural induction on c . All cases are straightforward. \square

Lemma 5. *Let c be a dependently guarded choreography. If $\llbracket c \rrbracket \downarrow$ then $c \downarrow$.*

Proof sketch. This proof is by structural induction on c . All cases are straightforward. \square

Theorem 6. *Let c be a dependently guarded choreography. Then $c \sim \llbracket c \rrbracket$.*

Proof. Recall the relation $\mathcal{R} = \{\langle c, \llbracket c \rrbracket \rangle \mid c \text{ is a dependently guarded choreography}\}$. Let $\langle c, R \rangle \in \mathcal{R}$.

- If $c \xrightarrow{\ell} c'$ then $R \xrightarrow{\ell} R'$ and $\langle c', R' \rangle \in \mathcal{R}$ (Lemma 2).
- If $R \xrightarrow{\ell} R'$ then $c \xrightarrow{\ell} c'$ and $\langle c', R' \rangle \in \mathcal{R}$ (Lemma 3).
- If $c \downarrow$ then $R \downarrow$ (Lemma 4).
- If $R \downarrow$ then $c \downarrow$ (Lemma 5).

Then \mathcal{R} is a bisimulation relation and $c \sim \llbracket c \rrbracket$ ([16]). \square

5 Conclusion

We have defined a choreography language and its operational semantics (Figures 3 and 4) using the weak sequential composition and partial termination of Rensink and Wehrheim [15], which is novel in the context of choreographies. We have defined a model, branching pomsets (Definition 1), which can compactly represent both concurrency and choices, and have defined its semantics (Figure 7). We have shown that we can use branching pomsets to model choreographies (Figure 9) and that this model is behaviourally equivalent to the operational semantics (Theorem 6).

We believe that branching pomsets can be further improved. We mention three points in particular and then discuss related work.

Binary choices Our branching structure \mathcal{B} only supports binary choices. This matches the structure of choreographies, but it would be more natural to represent $c_1 + (c_2 + c_3)$ as a single choice between the pomsets $\llbracket c_1 \rrbracket$, $\llbracket c_2 \rrbracket$ and $\llbracket c_3 \rrbracket$ instead of as two nested binary choices. However, supporting arbitrary n -ary choices also requires some thought about how to change the rules for refinement (Figure 7a), in particular CHOICE. A naive change would be to simply have this rule use $i \in \{1, \dots, n\}$ and $\{\{\mathcal{B}_1, \dots, \mathcal{B}_n\}\}$ instead of its current binary rules, but this is not sufficient as this naive n -ary choice would not be equivalent to the same branches composed as nested binary choices. For example, $c_1 + (c_2 + c_3)$ can partially terminate to $c_1 + c_2$ and its interpretation as a branching pomset can refine to $\llbracket c_1 + c_2 \rrbracket$, but a branching pomset whose branching structure consists of a single ternary choice $\{\{\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3\}\}$ would not be able to refine to $\{\{\mathcal{B}_1, \mathcal{B}_2\}\}$ as the rules would only allow it to refine all of its branches or discard all but one of them. Properly supporting n -ary choices would thus also require a new rule that allows $\{\{\mathcal{B}_1, \dots, \mathcal{B}_m\}\}$ to refine to choice between an arbitrary (non-empty) subset of its branches.

Partial order In Definition 1, \leq is defined as a relation on events such that its transitive closure is a partial order, rather than \leq being a partial order itself as it is in traditional pomsets. The need for this change arises from the update rule $R[\mathcal{B}]$ (Figure 7c) in our use case as choreographies. Consider the branching pomset in Figure 5. To match the operational semantics, we should be able to refine this pomset by discarding the $b \rightarrow c : x$ branch of the choice, after which $cd!x$ should be minimal. In our current rules the events $bc!x$ and $bc?x$ are removed along with their entries in \leq and then $cd!x$ is minimal. However, if \leq is a partial order, then since a partial order is transitive \leq would also contain the entries $ab!x \leq cd!x$ and $ab?x \leq cd!x$ and, since these entries do not contain $bc!x$ or $bc?x$ but are obtained by transitivity, they are not removed. Consequently, there would be no refinement that enables $cd!x$.

In general, if $R_1 \sqsupseteq R'_1$ and $R_2 \sqsupseteq R'_2$ then it would not necessarily be true that $R_1 ; R_2 \sqsupseteq R'_1 ; R'_2$, as $R_1 ; R_2$ may contain dependencies obtained by transitivity which would still be present in its updated version but which cannot be derived in $R'_1 ; R'_2$. We have no ready alternative. In the case of choreographies it may suffice to provide a more sophisticated update rule which properly trims these unwanted dependencies, but since this relies on knowledge of how these dependencies were derived from choreographies it is difficult to see how this could be applied to branching pomsets in general.

Loops In Figure 9 a loop c^* is encoded by infinitely unfolding it. As such, branching pomsets do not currently provide a finite representation of infinite choreographies. This remains a topic for future work, for which we envision two possible directions. One possibility would be to add an explicit repetition construct to the branching structure (e.g., change the second grammatical rule to $\mathcal{C} = e \mid \{\mathcal{B}_1, \mathcal{B}_2\} \mid \mathcal{B}^*$) and expand the semantics and proofs accordingly. Another possibility would be to explore the approach used in message sequence chart graphs [1] and add a graph structure on top of the branching structure.

Related work Choreographies are typically used in a top-down workflow: the developer writes a global view C and decomposes it into its projections, such that the behaviour of C is *behaviourally equivalent* to the parallel composition of its projections. Examples of this approach include workflows based on message sequence charts [11, 1], multiparty session types [9, 10], and choreographic programs [2, 6]. The choreographic language used in this paper assumes asynchronous communication between agents and includes a finite loop operator, borrowing from this literature the same notion of actions as interactions and their (parallel, sequential, and choice) composition.

Pomsets were initially introduced by Pratt [14] for concurrent models and have been widely used, e.g., in the context of message sequence charts by Katoen and Lambert [12]. Recently Guanciale and

Tuosto proposed two semantic frameworks for choreographies, one of which uses sets of pomsets [17]. They also note that the pomset framework exhibits exponential growth in the number of choices in a choreography, and they propose an alternative semantic framework using hypergraphs, which can compactly represent choices. While the hypergraph framework is more compact, their pomset framework is simpler and, they believe, more elegant. We agree with this analysis, and we aim to preserve the simplicity and elegance of the pomset framework by proposing a framework that avoids exponential growth in the number of choices while still being based on pomsets. In another recent paper Guanciale and Tuosto use pomsets to reason over choreography realisability [8]. This demonstrates the potential of using pomsets for semantic analysis, and we are investigating how to use our framework for similar analysis.

Other related work includes the usage of event structures in the context of binary session types by Castellán and Yoshida [3] and multiparty by Castellani et al. [4]. Event structures and branching pomsets both feature a set of events with a causality relation and a choice mechanism. The main difference between the two approaches is in the latter. Choices in event structures are based on a conflict relation on events, where two events in conflict cannot occur together in an execution and one of the two must be chosen. In contrast, we structure events in branching pomsets hierarchically. In the latter work, Castellani et al. use prime event structures to encode global types. The conflict heredity property of prime event structures leads to event duplication, which is avoided in flow event structures and branching pomsets. We leave a detailed comparison between flow event structures and branching pomsets for future work. We note that, given a branching pomset, one may construct an event structure by defining its conflict relation as all pairs of events that belong to different branches of some choice in the branching structure.

Acknowledgments

Funding by G. Cledou and J. Proença: European Regional Development Fund (ERDF), Operational Programme for Competitiveness and Internationalisation (COMPETE 2020), and Fundação para a Ciência e a Tecnologia (FCT): POCI-01-0145-FEDER-029946 (DaVinci). *Funding by S. Jongmans:* Netherlands Organisation of Scientific Research: 016.Veni.192.103. *Funding by J. Proença:* FCT, within the CISTER Research Unit: UIDB/04234/2020. FCT, I.P. (Portuguese Foundation for Science and Technology): PTDC/CCI-COM/4280/2021 (IBEX). ERDF and FCT, Portugal 2020 Partnership Agreement, Norte Portugal Regional Operational Programme (NORTE 2020): NORTE-01-0145-FEDER-028550 (REASSURE). ECSEL Joint Undertaking (JU): grant agreement No 876852 (VALU3S).

References

- [1] Rajeev Alur, Kousha Etessami & Mihalis Yannakakis (2005): *Realizability and verification of MSC graphs*. *Theor. Comput. Sci.* 331(1), pp. 97–114, doi:10.1016/j.tcs.2004.09.034.
- [2] Marco Carbone & Fabrizio Montesi (2013): *Deadlock-freedom-by-design: multiparty asynchronous global programming*. In Roberto Giacobazzi & Radhia Cousot, editors: *The 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '13, Rome, Italy - January 23 - 25, 2013*, ACM, pp. 263–274, doi:10.1145/2429069.2429101.
- [3] Simon Castellán & Nobuko Yoshida (2019): *Two sides of the same coin: session types and game semantics: a synchronous side and an asynchronous side*. *Proc. ACM Program. Lang.* 3(POPL), pp. 27:1–27:29, doi:10.1145/3290340.
- [4] Ilaria Castellani, Mariangiola Dezani-Ciancaglini & Paola Giannini (2019): *Event Structure Semantics for Multiparty Sessions*. In Michele Boreale, Flavio Corradini, Michele Loreti & Rosario Pugliese, editors: *Models, Languages, and Tools for Concurrent and Distributed Programming - Essays Dedicated to Rocco*

- De Nicola on the Occasion of His 65th Birthday, Lecture Notes in Computer Science* 11665, Springer, pp. 340–363, doi:10.1007/978-3-030-21485-2_19.
- [5] Guillermina Cledou, Luc Edixhoven, Sung-Shik Jongmans & José Proença (2022): *API Generation for Multiparty Session Types, Revisited and Revised Using Scala 3*. In Karim Ali & Jan Vitek, editors: *36th European Conference on Object-Oriented Programming, ECOOP 2022, June 6-10, 2022, Berlin, Germany, LIPIcs 222*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 27:1–27:28, doi:10.4230/LIPIcs.ECOOP.2022.27.
- [6] Luís Cruz-Filipe & Fabrizio Montesi (2020): *A core model for choreographic programming*. *Theor. Comput. Sci.* 802, pp. 38–66, doi:10.1016/j.tcs.2019.07.005.
- [7] Luc Edixhoven, Sung-Shik Jongmans, José Proença & Guillermina Cledou (2022): *Branching pomsets for choreographies (technical report)*. Technical Report OUNL-CS-2022-04, Open University of the Netherlands.
- [8] Roberto Guanciale & Emilio Tuosto (2019): *Realisability of pomsets*. *J. Log. Algebraic Methods Program.* 108, pp. 69–89, doi:10.1016/j.jlamp.2019.06.003.
- [9] Kohei Honda, Nobuko Yoshida & Marco Carbone (2008): *Multiparty asynchronous session types*. In George C. Necula & Philip Wadler, editors: *Proceedings of the 35th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2008, San Francisco, California, USA, January 7-12, 2008*, ACM, pp. 273–284, doi:10.1145/1328438.1328472.
- [10] Kohei Honda, Nobuko Yoshida & Marco Carbone (2016): *Multiparty Asynchronous Session Types*. *J. ACM* 63(1), pp. 9:1–9:67, doi:10.1145/2827695.
- [11] ITU (2011): *ITU-T recommendation Z.120. Message Sequence Chart (MSC)*.
- [12] Joost-Pieter Katoen & Lennard Lambert (1998): *Pomsets for MSC*. In Hartmut König & Peter Langendörfer, editors: *Formale Beschreibungstechniken für verteilte Systeme, 8. GI/ITG-Fachgespräch, Cottbus, 4. und 5. Juni 1998*, Verlag Shaker, pp. 197–207.
- [13] Romyana Neykova & Nobuko Yoshida (2017): *Let it recover: multiparty protocol-induced recovery*. In Peng Wu & Sebastian Hack, editors: *Proceedings of the 26th International Conference on Compiler Construction, Austin, TX, USA, February 5-6, 2017*, ACM, pp. 98–108, doi:10.1145/3033019.3033031. Available at <http://dl.acm.org/citation.cfm?id=3033031>.
- [14] Vaughan R. Pratt (1986): *Modeling concurrency with partial orders*. *Int. J. Parallel Program.* 15(1), pp. 33–71, doi:10.1007/BF01379149.
- [15] Arend Rensink & Heike Wehrheim (2001): *Process algebra with action dependencies*. *Acta Informatica* 38(3), pp. 155–234, doi:10.1007/s002360100070.
- [16] Davide Sangiorgi (2011): *Introduction to bisimulation and coinduction*. Cambridge University Press, doi:10.1017/CBO9780511777110.
- [17] Emilio Tuosto & Roberto Guanciale (2018): *Semantics of global view of choreographies*. *J. Log. Algebraic Methods Program.* 95, pp. 17–40, doi:10.1016/j.jlamp.2017.11.002.