



HAL
open science

Weakly Synchronous Systems with Three Machines Are Turing Powerful

Cinzia Di Giusto, Davide Ferré, Etienne Lozes, Nicolas Nisse

► **To cite this version:**

Cinzia Di Giusto, Davide Ferré, Etienne Lozes, Nicolas Nisse. Weakly Synchronous Systems with Three Machines Are Turing Powerful. RP 2023 - 17th International Conference on Reachability Problems, 2023, Nice, France. pp.28-41, 10.1007/978-3-031-45286-4_3. hal-04273451

HAL Id: hal-04273451

<https://hal.science/hal-04273451>

Submitted on 7 Nov 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Weakly synchronous systems with three machines are Turing powerful

Cinzia Di Giusto¹[0000–0003–1563–6581], Davide Ferre²[0000–0002–6578–6804],
Etienne Lozes¹[0000–0001–8505–585X], and Nicolas Nisse²[0000–0003–4500–5078]

¹ Université Côte d’Azur, CNRS, I3S, France

² Université Côte d’Azur, Inria, CNRS, I3S, France

Abstract. Communicating finite-state machines (CFMs) are a Turing powerful model of asynchronous message-passing distributed systems. In weakly synchronous systems, processes communicate through phases in which messages are first sent and then received, for each process. Such systems enjoy a limited form of synchronization, and for some communication models this restriction is enough to make the reachability problem decidable. In particular, we explore the intriguing case of p2p (FIFO) communication, for which the reachability problem is known to be undecidable for four processes, but decidable for two. We show that the configuration reachability problem for weakly synchronous systems of three processes is undecidable. This result is heavily inspired by our study on the treewidth of the Message Sequence Charts (MSCs) that might be generated by such systems. In this sense, the main contribution of this work is a weakly synchronous system with three processes that generates MSCs of arbitrarily large treewidth.

Keywords: Distributed Systems · Message-Passing · Treewidth

1 Introduction

Systems of communicating finite-state machines (CFMs) are a simple, yet expressive, model of asynchronous message-passing distributed systems. In this model, each machine performs a sequence of send and receive actions, where a send action can be matched by a receive action of another machine. For instance, the system in Fig. ?? (left), models a protocol between three processes a , b , and r .

A computation of such a system can be represented graphically by a Message Sequence Chart (MSC), a simplified version of the ITU recommendation [?]. Each machine of the system has its own “timeline” on the MSC, where actions are listed in the order in which they are executed, and message arrows link a send action to its matching receive action. For instance, the MSC of Fig. ?? (right) represents one of the many computations of the system in Fig. ?? (left). The set of all MSCs that the system may generate is determined both by the machines, since the sequence of actions of each timeline must be a sequence of action in the corresponding CFM, and by the “transport layer” or “communication model” employed by the machines. Roughly speaking, a communication model is a class of




Fig. 1: Example of a system of 3 CFMs (left) and of an MSC generated by it (right). $p!q(a)$ (resp., $p?q(a)$) denotes the sending (reception) of message a from (by) process a to (from) process b .

MSCs that are considered “realizable” within that model of communications. For instance, for rendezvous synchronization, an MSC is considered to be realizable with synchronous communication if the only path between a sending and its matching receipt is the direct one through the message arrow that relates them. Among the various communication models that have been considered, we can cite p2p (or FIFO) model, where each ordered pair of machines defines a dedicated FIFO queue; causal ordering (CO), where a message cannot overtake the messages that were sent causally before it; the mailbox model, where each machine holds a unique FIFO queue for all incoming messages; the bag (or simply asynchronous) model, where a message can overtake any other message (see [?, ?, ?] for various presentations of these communication models).

The configuration reachability problem for a system of CFMs consists in checking whether a control state, together with a given content of the queues, is reachable from the initial state. This problem is decidable for synchronous communication, as the state space of the system is finite, and also for bag communication, by reduction to Petri nets [?]. For other communication models, as soon as two machines are allowed to exchange messages through two FIFO queues, reachability becomes undecidable [?]. Due to this strong limitation, there has been a wealth of work that tried to recover decidability of reachability by considering systems of CFMs that are “almost synchronous”.

In weakly synchronous systems, processes communicate through phases in which messages are first sent and then received, for each process; graphically, the MSCs of such systems are the concatenation of smaller, independent MSCs, within which no send happens after a receive. For instance, the MSC in Fig. ?? (right) is weakly synchronous, as it is the concatenation of three “blocks” (namely $\{m_1\}$, $\{m_2\}$, and $\{m_2, m_3, m_4\}$), within which all sends of a given machine happen before all the receives of this same machine. It is known that reachability is decidable for mailbox weakly synchronous systems [?], whereas it is undecidable for either p2p or CO weakly synchronous systems with at least four machines. On the other hand, reachability is decidable for two machines (since any p2p MSC with two

machines is also mailbox). In this work, we deal with weakly synchronous systems with three machines, and conclude that reachability is undecidable for these systems. Our result is based on a study of the unboundedness of the treewidth for MSCs that may be generated by these systems. The first contribution of this work is a weakly synchronous system with only three machines that is “treewidth universal”, in the sense that it may generate MSCs of arbitrarily large treewidth. The second contribution, strongly inspired by the treewidth universal system, is showing that weakly synchronous systems with three processes are Turing powerful. To do so, we establish a one-to-one correspondence between the computations of a FIFO automaton (a finite state machine that may push and pop from a FIFO queue, which is known to be a Turing powerful computational model) on the one hand, and a subset of the MSCs of the treewidth universal system on the other hand.

Related work. Beyond weakly synchronous systems, several similar notions have been considered to try to capture the intuition of an “almost synchronous” synchronous. Reachability of existentially bounded systems [?,?] is decidable for FIFO, CO, p2p, or bag communications. Synchronizable systems [?] were an attempt to define a class of systems with good decidability properties, however reachability for such systems with FIFO communications is undecidable [?]. The status of reachability for k -stable systems [?] is unknown. Finally, reachability for k -synchronous systems [?] is decidable for FIFO, CO, p2p, or bag communications.

Another form of under-approximation of the full behaviour of a system of CFMs is the bounded context-switch reachability problem, which is known to be decidable for systems of CFMs, even with a controlled form of function call [?,?].

Finally, weak synchronisability share some similarities with reversal-bounded counter machines [?,?]: in the context of bag communications, a send is a counter increment, a receive a decrement, and weak synchronisability is a form of bounding the number of reversals of increment and decrement phases.

Outline. Section ?? introduces the necessary terminology. Section ?? presents the weakly synchronous system with three machines that may generate MSCs of arbitrarily large treewidth. Then, Section ?? discusses the undecidability of the configuration reachability problem for weakly-synchronous systems with three machines. Finally, Section ?? concludes with some final remarks. An Appendix with proofs and additional material is added for the reviewer convenience.

2 MSCs and communicating automata

We recall here concepts and definitions related to MSCs and communicating automata. Assume a finite set of processes \mathbb{P} and a finite set of messages \mathbb{M} . A send action is of the form $p!q(m)$ where $p, q \in \mathbb{P}$ and $m \in \mathbb{M}$; it is executed by p and sends message m to process q . The corresponding receive action, executed by q , is $p?q(m)$. Let $Send(p, q, _) = \{p!q(m) \mid m \in \mathbb{M}\}$ and $Rec(p, q, _) = \{p?q(m) \mid m \in \mathbb{M}\}$. For $p \in \mathbb{P}$, we set $Send(p, _, _) = \{p!q(m) \mid q \in \mathbb{P} \setminus \{p\} \text{ and } m \in \mathbb{M}\}$, etc. Moreover, $\Sigma_p = Send(p, _, _) \cup Rec(_, p, _) \cup \{\varepsilon\}$ will denote the set of all actions that are executed by p . Finally, $\Sigma = \bigcup_{p \in \mathbb{P}} \Sigma_p$ is the set of all the actions.

Definition 1 (p2p MSC). A (p2p) MSC over \mathbb{P} and \mathbb{M} is a tuple $\text{MSC} = (\mathcal{E}, \rightarrow, \triangleleft, \lambda)$ where \mathcal{E} is a finite (possibly empty) set of events and $\lambda : \mathcal{E} \rightarrow \Sigma$ is a labeling function. For $p \in \mathbb{P}$, let $\mathcal{E}_p = \{e \in \mathcal{E} \mid \lambda(e) \in \Sigma_p\}$ be the set of events that are executed by p . \rightarrow (the process relation) is the disjoint union $\bigcup_{p \in \mathbb{P}} \rightarrow_p$ of relations $\rightarrow_p \subseteq \mathcal{E}_p \times \mathcal{E}_p$ such that \rightarrow_p is the direct successor relation of a total order on \mathcal{E}_p . $\triangleleft \subseteq \mathcal{E} \times \mathcal{E}$ (the message relation) satisfies the following:

- (1) for every pair $(s, r) \in \triangleleft$, there is a send action $p!q(m) \in \Sigma$ such that $\lambda(s) = p!q(m)$, $\lambda(r) = p?q(m)$, and $p \neq q$;
- (2) for all $r \in \mathcal{E}$ with $\lambda(r) = p?q(m)$, there is a unique $s \in \mathcal{E}$ such that $s \triangleleft r$;
- (3) letting $\leq_{\text{MSC}} = (\rightarrow \cup \triangleleft)^*$, we require that \leq_M is a partial order;
- (4) for every $s_1 \in \mathcal{E}$ and pair $(s_2, r_2) \in \triangleleft$ with $\lambda(s_1) = p!q(m_1)$ and $\lambda(s_2) = p!q(m_2)$, if $s_1 \rightarrow_p^+ s_2$, then there exists r_1 such that $(s_1, r_1) \in \triangleleft$ and $r_1 \rightarrow_q^+ r_2$.

Condition (1) above ensures that message arrows relate a send event to a receive event on a distinct machine. By Condition (2), every receive event has a matching send event. Note that, however, there may be unmatched send events in an MSC. An MSC is called *orphan free* if all send events are matched. Condition (3) ensures that there exists at least one scheduling of all events such that each receive event happens after its matching send event. Condition (4) captures the p2p communication model: a message cannot overtake another message that has the same sender and same receiver as itself.

Let $M = (\mathcal{E}, \rightarrow, \triangleleft, \lambda)$ be an MSC, then $\text{SendEv}(M) = \{e \in \mathcal{E} \mid \lambda(e) \text{ is a send action}\}$, $\text{RecEv}(M) = \{e \in \mathcal{E} \mid \lambda(e) \text{ is a receive action}\}$, $\text{Matched}(M) = \{e \in \mathcal{E} \mid \text{there is } f \in \mathcal{E} \text{ such that } e \triangleleft f\}$, and $\text{Unm}(M) = \{e \in \mathcal{E} \mid \lambda(e) \text{ is a send action and there is no } f \in \mathcal{E} \text{ such that } e \triangleleft f\}$. We do not distinguish isomorphic MSCs. Let $E \subseteq \mathcal{E}$ such that E is \leq_M -downward-closed, i.e, for all $(e, f) \in \leq_M$ such that $f \in E$, we also have $e \in E$. Then the MSC $M' = (E, \rightarrow, \triangleleft, \lambda)$ obtained by restriction to E is called a *prefix* of M . If $M_1 = (\mathcal{E}_1, \rightarrow_1, \triangleleft_1, \lambda_1)$ and $M_2 = (\mathcal{E}_2, \rightarrow_2, \triangleleft_2, \lambda_2)$ are two MSCs, their *concatenation* $M_1 \cdot M_2 = (\mathcal{E}, \rightarrow, \triangleleft, \lambda)$ is as expected: \mathcal{E} is the disjoint union of \mathcal{E}_1 and \mathcal{E}_2 , $\triangleleft = \triangleleft_1 \cup \triangleleft_2$, λ is the “union” of λ_1 and λ_2 , and $\rightarrow = \rightarrow_1 \cup \rightarrow_2 \cup R$. Here, R contains, for all $p \in \mathbb{P}$ such that $(\mathcal{E}_1)_p$ and $(\mathcal{E}_2)_p$ are non-empty, the pair (e_1, e_2) , where e_1 and e_2 are the last and the first event executed by p in M_1 and M_2 , respectively. Due to condition (4), concatenation is a partially defined operation: $M_1 \cdot M_2$ is defined if for all $s_1 \in \text{Unm}(M_1)$ and $s_2 \in \text{SendEv}(M_2)$ that have the same sender and destination ($\lambda(s_1) \in \text{Send}(p, q, _)$ and $\lambda(s_2) \in \text{Send}(p, q, _)$), we have $s_2 \in \text{Unm}(M_2)$. In particular, $M_1 \cdot M_2$ is defined when M_1 is orphan-free. Concatenation is associative.

We recall from [?] the definition of weakly synchronous MSC. We say that an MSC is weakly synchronous if it can be broken into phases where all sends are scheduled before all receives.

Definition 2 (weakly synchronous). We say that $M \in \text{MSC}$ is weakly synchronous if it is of the form $M = M_1 \cdot M_2 \cdots M_n$ such that for every $M_i = (\mathcal{E}, \rightarrow, \triangleleft, \lambda)$ $\text{SendEv}(M_i)$ is a \leq_{M_i} -downward-closed set.

We now recall the definition of communicating system, which consists of finite-state machines A_p (one per process $p \in \mathbb{P}$) that can exchange messages.

Definition 3 (communicating system). A (communicating) system over \mathbb{P} and \mathbb{M} is a tuple $\mathcal{S} = ((A_p)_{p \in \mathbb{P}})$. For each $p \in \mathbb{P}$, $A_p = (Loc_p, \delta_p, \ell_p^0, \ell_p^{acc})$ is a finite transition system where: Loc_p is the finite set of (local) states of p , $\delta_p \subseteq Loc_p \times \Sigma_p \times Loc_p$ (also denoted $\ell \xrightarrow[A_p]{a} \ell'$) is the transition relation of p , $\ell_p^{acc} \in Loc_p$ is the final state of p .

Given $p \in \mathbb{P}$ and a transition $t = (\ell, a, \ell') \in \delta_p$, we let $source(t) = \ell$, $target(t) = \ell'$, $action(t) = a$, and $msg(t) = m$ if $a \in Send(_, _, m) \cup Rec(_, _, m)$.

An *accepting run* of \mathcal{S} on an MSC M is a mapping $\rho : \mathcal{E} \rightarrow \bigcup_{p \in \mathbb{P}} \delta_p$ that assigns to every event e the transition $\rho(e)$ that is executed at e by A_p . Thus, we require that (i) for all $e \in \mathcal{E}$, we have $action(\rho(e)) = \lambda(e)$, (ii) for all $(e, f) \in \rightarrow$, $target(\rho(e)) \xrightarrow[A_p]{\varepsilon}^* source(\rho(f))$, (iii) for all $(e, f) \in \triangleleft$, $msg(\rho(e)) = msg(\rho(f))$, (iv) for all $p \in \mathbb{P}$ and $e \in \mathcal{E}_p$ such that there is no $f \in \mathcal{E}$ with $f \rightarrow e$, we have $source(\rho(e)) = \ell_p^0$, (v) for all $p \in \mathbb{P}$ and $e \in \mathcal{E}_p$ such that there is no $f \in \mathcal{E}$ with $e \rightarrow f$, we have $target(\rho(e)) = \ell_p^{acc}$ and, (vi) $Unm(M) = \emptyset$. Essentially, in an accepting run of \mathcal{S} every A_p takes a sequence of transitions that lead to its final state ℓ_p^{acc} , and such that each send action will have a matching receive action (i.e., there are no unmatched messages). The *language* of \mathcal{S} is $L(\mathcal{S}) = \{MSC \in MSC \mid \text{there is an accepting run of } \mathcal{S} \text{ on } M\}$. We say that \mathcal{S} is weakly synchronous if for all $M \in L(\mathcal{S})$, M is weakly synchronous.

The *emptiness problem* is the decision problem that takes as input a system \mathcal{S} and addresses the question “is $L(\mathcal{S})$ empty?”. This problem is a configuration reachability problem, and under several circumstances, its decidability is closely related to the one of the control state reachability problem. In this work, we will study the emptiness problem with the additional hypothesis that \mathcal{S} is a weakly synchronous system with three machines only.

Finally, we recall the less known notion of “FIFO automaton”, a finite state machine that can push into and pop from a FIFO queue. This is a system of communicating machines with just one machine, whose semantics is a set of MSCs with a single timeline, for which we exceptionally relax condition (1) of Definition ??, so to allow a send event and its matching receive event to occur on the same machine. The following result is proved in [?, Lemma 4].

Lemma 1 ([?]). *The emptiness problem for FIFO automata is undecidable.*

3 Treewidth of weakly synchronous p2p MSCs

There is a strong correlation between MSCs and graphs. An MSC is a directed graph (digraph in the following) where the nodes are the events of the MSC and the arcs are represented by the \rightarrow and the \triangleleft relations. We are, therefore, able to use some tools and techniques from graph theory to possibly derive some interesting results about MSCs. A graph parameter which is particularly important in this context is the *treewidth* [?] mostly due to Courcelle’s theorem that, roughly, states that many properties can be checked in classes of MSCs

with bounded treewidth³. For instance, in [?], it is shown that the class of weakly synchronous mailbox MSCs has bounded treewidth. Interestingly enough, it is also shown that the bigger class of weakly synchronous p2p MSCs has unbounded treewidth, by means of a reduction to the Post correspondence problem. Here we give a more direct proof, for all weakly synchronous systems that have at least three processes. We begin with some terminology:

Definition 4. To contract an arc (u, v) in a (di)graph G means replacing u and v by a single vertex whose neighbourhood is the union of the neighbourhoods of u and v . A (di)graph H is a minor of a (di)graph G if H can be obtained from a subgraph of G by contracting some edges/arcs.

Next, we show how to build a family of weakly synchronous MSCs with three processes (a , b and c) and unbounded treewidth. We want to find a class of MSCs that admit grids of unbounded size as a minor. The idea is illustrated in Fig. ??, and it consists in bouncing groups of messages between processes so to obtain the depicted shape. The class of MSCs is indexed by two non-zero natural numbers: h and ℓ . Intuitively, h represents the number of consecutive events in a group, and ℓ is the number of groups per process, divided by 2. The graph depicted on the top left of Fig. ?? is not an MSC, because it is undirected and there are multiple actions associated to the same event. Nonetheless, the connection with MSCs is quite intuitive, and formalized in Lemma ??.

We, now, specify how to build a digraph $G_{h,\ell} = (V(G_{h,\ell}), E(G_{h,\ell}))$, from which our MSC $G_{h,\ell}^*$ will be obtained. The set of vertices $V(G_{h,\ell}) = \mathcal{A} \cup \mathcal{B} \cup \mathcal{C}$ contains all the events of each process: $\mathcal{A} = \{s_a^{i,j}, r_a^{i,j} \mid 1 \leq i \leq h, 1 \leq j \leq \ell\}$, $\mathcal{B} = \{s_b^{i,j}, r_b^{i,j} \mid 1 \leq i \leq h, 1 \leq j \leq \ell\}$, and $\mathcal{C} = \{s_c^{i,j}, r_c^{i,j} \mid 1 \leq i \leq h, 1 \leq j \leq \ell\}$.

For $x \in \{a, b, c\}$ and $y \in \{r, s\}$, we add the following arcs to $E(G_{h,\ell})$, which will represent the “timelines” connecting events of each process:

1. for each group of h events/messages and $1 \leq j \leq \ell$, $Col_{x,y,j} = \{(y_x^{i,j}, y_x^{i+1,j}) \mid 1 \leq i < h\}$;
2. then, to link groups together $\{(y_x^{h,j}, y_x^{1,j+1}) \mid 1 \leq j < \ell\}$;
3. and finally, to link the phase of sendings with the one of receptions: $(s_x^{h,\ell}, r_x^{1,1})$.

It remains to add the arcs that correspond to the messages exchanged by the processes. Intuitively, each vertex $s_x^{i,j}$ corresponds to two messages sent by process x to the other two processes (except for $j = 1$ and $x = a$, in which case it will correspond to a single message), and each vertex $r_x^{i,j}$ will correspond to two messages received by process x from the other two processes (except for $j = \ell$ and $x = c$, in which case it will correspond to a single message). Formally:

$$E_{\mathcal{M}} = \{(s_a^{i,j}, r_b^{i,j}), (s_c^{i,j}, r_b^{i,j}), (s_c^{i,j}, r_a^{i,j}), (s_b^{i,j}, r_a^{i,j}), (s_b^{i,j}, r_c^{i,j}) \mid 1 \leq i \leq h, 1 \leq j \leq \ell\} \\ \cup \{(s_a^{i,j+1}, r_c^{i,j}) \mid 1 \leq i \leq h, 1 \leq j < \ell\}. \quad (1)$$

³ Since we do not explicitly use tree-decompositions, we refer to [?] for their formal definitions. Definitions are recalled in Appendix ?? for the reviewer convenience.

Lemma 2. For any $h, \ell \in \mathbb{N}^*$, $G_{h,\ell}$ is the minor of a graph arising from a weakly synchronous p2p MSC $G_{h,\ell}^*$ with 3 processes and a single phase.

Proof. Fig. ?? exemplifies the transformation below. Note that some vertices of $G_{h,\ell}$ have degree 4 while any MSC is a subcubic graph (i.e., every vertex has degree at most 3). For every $s_x^{i,j}$ with degree 4, let α (resp., β) be the in-neighbour (resp., out-neighbour) of $s_x^{i,j}$ in \mathcal{P}_x and let γ and δ be the other two neighbours of $s_x^{i,j}$. Replace $s_x^{i,j}$ by two vertices $su_x^{i,j}$ and $sd_x^{i,j}$, with the 5 arcs $(\alpha, su_x^{i,j})$, $(su_x^{i,j}, sd_x^{i,j})$, $(sd_x^{i,j}, \beta)$, $(su_x^{i,j}, \gamma)$ and $(sd_x^{i,j}, \delta)$. Do a similar transformation for every $r_x^{i,j}$ with degree 4. A similar transformation is done for the four vertices (with degree 3) $s_b^{1,1}$, $s_c^{1,1}$, $r_a^{h,\ell}$ and $r_b^{h,\ell}$. Let $G_{h,\ell}^*$ be the obtained digraph. It is clear that $G_{h,\ell}^*$ is an MSC and that $G_{h,\ell}$ is a minor of $G_{h,\ell}^*$.

Note that for any $x \in \{a, b, c\}$, $\mathcal{X} \in \{\mathcal{A}, \mathcal{B}, \mathcal{C}\}$ induces a directed path \mathcal{P}_x with first the vertices $s_x^{i,j}$ (in increasing lexicographical order of (j, i)) and then the vertices $r_x^{i,j}$ (in increasing lexicographical order of (i, j)). The fact that $G_{h,\ell}^*$ is weakly synchronous with one phase directly follows the fact that, for every $x \in \{a, b, c\}$, the vertices s, su and sd (corresponding to sendings) appear before the vertices r, ru and rd (corresponding to receptions) in the directed path \mathcal{P}_x .

Moreover, for every $x, y \in \{a, b, c\}$, $x \neq y$, the arcs from \mathcal{P}_x to \mathcal{P}_y are all parallel (i.e., for every arc (u, v) and (u', v') from \mathcal{P}_x to \mathcal{P}_y , if u is a predecessor of u' in \mathcal{P}_x , then v is a predecessor of v' in \mathcal{P}_y). This implies that $G_{h,\ell}^*$ is p2p. \square

Note that, for fixed $i \leq h$ and $j < \ell$, $P_{i,j} = (s_a^{i,j}, r_b^{i,j}, s_c^{i,j}, r_a^{i,j}, s_b^{i,j}, r_c^{i,j}, s_a^{i+1,j})$ is a (undirected) path with 6 arcs linking $s_a^{i,j}$ to $s_a^{i+1,j}$. From this, it is not difficult to see that $G_{h,\ell}$ admits a grid of size $h \times 6\ell$ as a minor, which is the content of next lemma (see Fig. ?? for an example).

Let $tw(G_{h,\ell})$ be the treewidth of the underlying undirected graph of $G_{h,\ell}$.

Lemma 3. For any $h, \ell \in \mathbb{N}^*$, $tw(G_{h,\ell}) \geq \min\{h, 6\ell\}$.

Proof. The subgraph obtained from $G_{h,\ell}$ by keeping the arcs in item 1 and Equation ??: $G'_{h,\ell} = (V(G_{h,\ell}), E_{\mathcal{M}} \cup \bigcup_{x \in \{a,b,c\}, y \in \{r,s\}, 1 \leq j \leq \ell} Col_{x,y,j})$, is a $h \times 6\ell$ grid. From [?], we know that $tw(G'_{h,\ell}) \geq \min\{h, 6\ell\}$ and, since the treewidth is closed under subgraphs [?], $tw(G_{h,\ell}) \geq tw(G'_{h,\ell}) \geq \min\{h, 6\ell\}$. \square

We can then easily derive the main result for this section.

Theorem 1. The class of weakly synchronous p2p MSCs with three processes (and a single phase) has unbounded treewidth.

Proof. From Lemma ??, $G_{h,\ell}^*$ is a weakly synchronous p2p MSC with 3 processes and $G_{h,\ell}$ is a minor of $G_{h,\ell}^*$. Hence, from Lemma ?? and the fact that the treewidth is minor-closed [?], we get that $tw(G_{h,\ell}^*) \geq \min\{h, 6\ell\}$. \square




Fig. 2: Transformation of Lemma ??.

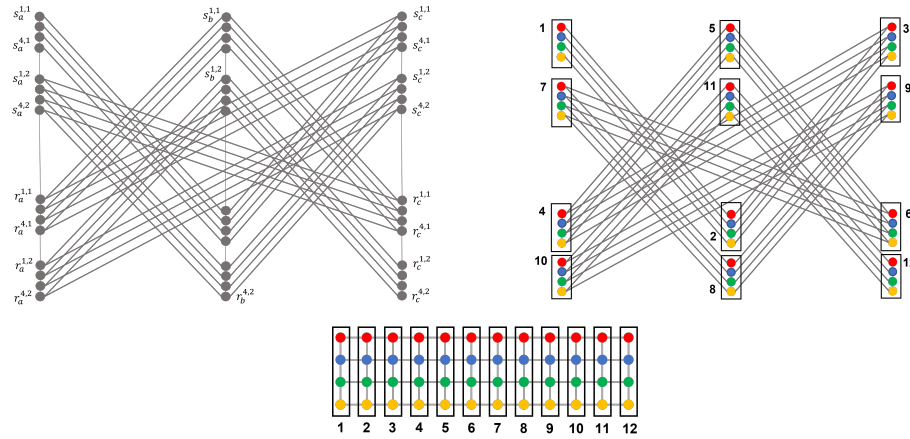


Fig. 3: The undirected graph of $G_{4,2}$ (top left) with a 4×12 grid as a minor (top right and bottom). All arcs go from top to bottom.

Notice that, a similar technique, this time exploiting four processes instead of three, can be used to show that we can build a weakly synchronous p2p MSC that can be contracted to whatever graph.

Theorem 2. *Let H be any graph. There exists a weakly synchronous p2p MSCs with four processes that admits H as minor.*


Proof. Let $V(H) = \{v_1, \dots, v_h\}$ and $E(H) = \{e_1, \dots, e_\ell\}$. Take graph $G_{h,\ell}$ defined above. Add a new directed path (d_1, \dots, d_ℓ) (which corresponds to the fourth process). Finally, for every $1 \leq j \leq \ell$, and edge $e_j = \{v_i, v_{i'}\} \in E(H)$, add two arcs $(r_a^{i,j}, d_j)$ and $(r_a^{i',j}, d_j)$. Let G be the obtained graph.

Using similar arguments as in the proof of Lemma ??, G arises from a weakly synchronous p2p MSC with 4 processes.

To see that H is a minor of G : first remove all “vertical” arcs. Then, for every $1 \leq i \leq h$, contract the path $\bigcup_{1 \leq j \leq \ell} P_{i,j}$ into one single vertex (corresponding to v_i), and finally contract the arc $(r_a^{i',j}, d_j)$ for every edge $e_j = \{v_i, v_{i'}\}$. These operations lead to H . \square

4 Reachability for weakly synchronous p2p systems with 3 machines

In [?], it is shown that the control state reachability problem for weakly p2p synchronous systems with at least 4 processes is undecidable. The result is obtained via a reduction of the Post correspondence problem. In the same paper, following from the boundedness of treewidth, it is also shown that reachability is decidable for systems with 2 processes. The arguments easily adapt to show the same results for the emptiness problem instead. The decidability of reachability,


 Fig. 4: Automata A_b and A_c of \mathcal{S}_3 .

or emptiness, remained open for systems with 3 processes. In the previous section, we showed that the treewidth of weakly synchronous **p2p** MSCs is unbounded for 3 processes. Here, inspired by the proof of the unboundedness of the treewidth, we provide a reduction from the emptiness problem for a FIFO automaton \mathcal{S}_1 to the emptiness problem for a weakly synchronous system \mathcal{S}_3 with three machines. The reduction makes sure that there is an accepting run of \mathcal{S}_1 if and only if there is one for \mathcal{S}_3 , which shows the undecidability of the emptiness problem for weakly synchronous systems with three machines.

Let $\mathcal{S}_1 = (A)$, with $A = (Loc, \delta, \ell^0, \ell^{acc})$ be a communicating system with a single process over \mathbb{M} . We will consider only automata that, from any state, have at most one non-epsilon outgoing transition, and no self loops (i.e., transitions that start and land in the same state). More precisely, we prove that any system can be encoded into one that satisfies this additional property while accepting the same language (see the corresponding encoding in Appendix ??).

We provide an encoding of the FIFO automaton \mathcal{S}_1 into the system $\mathcal{S}_3 = (A_a, A_b, A_c)$ over $\mathbb{M} \cup \{D\}$, where D is an additional special message called the *dummy message*. We show that \mathcal{S}_3 is weakly synchronous, and that $L(\mathcal{S}_1) \neq \emptyset$ if and only if $L(\mathcal{S}_3) \neq \emptyset$. Processes b and c (see Fig. ??) are used as forwarders so that messages circulate as in Fig. ?. Basically, process b (resp., process c) goes through two phases, the first one in which messages are sent to a and c (resp., a and b), and the second in which messages can be received. In Fig. ??, there is one state $\ell_{b_0}^m$ (resp., $\ell_{b_?}^m$), which is the in and out-neighbour of $\ell_{b_0}^m$ (resp., $\ell_{b_?}^m$), per message $m \in \mathbb{M} \cup \{D\}$. Formally, $A_b = (Loc_b, \delta_b, \ell_b^0, \ell_b^{acc})$ where

$$\begin{aligned} Loc_b &= \{\ell_b^0, \ell_b^?, \ell_b^{acc}\} \cup \{\ell_{b_0}^m, \ell_{b_?}^m \mid m \in \mathbb{M} \cup \{D\}\} \\ \delta_b &= \{(\ell_b^0, \varepsilon, \ell_b^?), (\ell_b^?, \varepsilon, \ell_b^{acc})\} \cup \\ &\quad \{(\ell_b^0, b!a(m), \ell_{b_0}^m), (\ell_{b_0}^m, b!c(m), \ell_b^0), \\ &\quad (\ell_b^?, b?a(m), \ell_{b_?}^m), (\ell_{b_?}^m, b?c(m), \ell_b^?) \mid m \in \mathbb{M} \cup \{D\}\} \end{aligned}$$

and symmetrically $A_c = (Loc_c, \delta_c, \ell_c^0, \ell_c^{acc})$ where

$$\begin{aligned} Loc_c &= \{\ell_c^0, \ell_c^?, \ell_c^{acc}\} \cup \{\ell_{c_0}^m, \ell_{c_?}^m \mid m \in \mathbb{M} \cup \{D\}\} \\ \delta_c &= \{(\ell_c^0, \varepsilon, \ell_c^?), (\ell_c^?, \varepsilon, \ell_c^{acc})\} \cup \{(\ell_c^0, c!b(m), \ell_{c_0}^m), (\ell_{c_0}^m, c!a(m), \ell_c^0), \\ &\quad (\ell_c^?, c?b(m), \ell_{c_?}^m), (\ell_{c_?}^m, c?a(m), \ell_c^?) \mid m \in \mathbb{M} \cup \{D\}\}. \end{aligned}$$

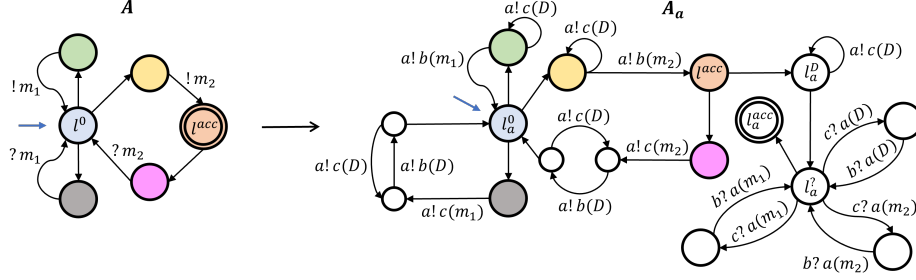


Fig. 5: The automaton A_a for the system \mathcal{S}_3 , built from the automaton A of \mathcal{S}_1 . Arcs without actions represent ϵ transitions.

Process a mimics the behavior of A . Fig. ?? shows an example of how A_a is built, starting from A . At a high level, A_a is composed of two parts: the first simulates A , and the second (after state ℓ_a^D) receives all messages sent by b and c . In the first part of A_a , each send action of A is replaced by a send action addressed to process b , and each reception of A is replaced by a send action to process c . We then use some dummy messages to ensure that our encoding works properly. Roughly, we force A_a to send a dummy message to b after each message sent to c , and we let A_a send any number of dummy messages to c right before each message sent to b . Similarly, after A_a sends a dummy message to b , it is also allowed to send two other dummy messages (the first one to c and the second one to b) an unbounded number of times. Formally, $A_a = (Loc_a, \delta_a, \ell_a^0, \ell_a^{acc})$, where:

$$\begin{aligned}
Loc_a &= Loc \cup \{\ell_{t_1}, \ell_{t_2} \mid t = (\ell, ?m, \ell') \in \delta\} \cup \\
&\quad \{\ell_a^D, \ell_a^?, \ell_a^{acc}\} \cup \{\ell_{a?}^m \mid m \in \mathbb{M} \cup \{D\}\} \\
\delta_a &= \{(\ell, a!b(m), \ell'), (\ell, a!b(D), \ell) \mid (\ell, !m, \ell') \in \delta\} \cup \\
&\quad \{(\ell, a!c(m), \ell_{t_1}), (\ell_{t_1}, a!b(D), \ell_{t_2}), \\
&\quad (\ell_{t_2}, a!c(D), \ell_{t_1}), (\ell_{t_2}, \varepsilon, \ell') \mid t = (\ell, ?m, \ell') \in \delta\} \cup \\
&\quad \{(\ell, \varepsilon, \ell') \mid (\ell, \varepsilon, \ell') \in \delta\} \cup \\
&\quad \{(\ell^{acc}, \varepsilon, \ell_a^D), (\ell_a^D, a!c(D), \ell_a^D)\} \cup \\
&\quad \{(\ell_a^D, \varepsilon, \ell_a^?), (\ell_a^?, \varepsilon, \ell_a^{acc})\} \cup \\
&\quad \{(\ell_a^?, a?c(m), \ell_{a?}^m), (\ell_{a?}^m, a?b(m), \ell_a^?) \mid m \in \mathbb{M} \cup \{D\}\}
\end{aligned}$$

In Fig. ??, colors show the mapping of states from an instance of A to the corresponding automaton A_a . Fig. ?? illustrates an accepting run of some system \mathcal{S}_1 and one of the corresponding accepting runs of the associated \mathcal{S}_3 .

Given a sequence of actions $!m$ and $?m$, where m can be any message, we call it a FIFO sequence if (i) all messages are received in the order in which they are sent, and (ii) no message is received before being sent. We relax this definition to talk about sequences of send actions $a!b(m)$ and $a!c(m)$ taken by A_a (in the first




Fig. 6: Above, a run with two messages for some system \mathcal{S}_1 . Below, one possible corresponding MSC realized by the associated \mathcal{S}_3 . Gray lines correspond to dummy messages.

part of the automaton); in particular, we say that such a sequence γ' is FIFO if, when interpreting each $a!b(m)$ and $a!c(m)$ action as $!m$ and $?m$, respectively, γ' is a FIFO sequence. Dummy messages are used to enforce that the sequence of send actions taken by A_a in an accepting run of \mathcal{S}_3 is FIFO.

Theorem 3. *There is an accepting run of \mathcal{S}_1 if and only if there is an accepting run of \mathcal{S}_3 .*

Sketch of proof. We only provide a sketch of the proof, which is quite convoluted and requires several intermediate lemmata. The full proof is in Appendix ??.

(\Rightarrow) We design Algorithm ??, which takes an accepting run σ of \mathcal{S}_1 , and returns an accepting run μ for \mathcal{S}_3 . At a high level, Algorithm ?? takes the sequence of actions taken by A in σ , rewrites each $!m$ and $?m$ action as $a!b(m)$ and $a!c(m)$, and then adds some actions related to dummy messages. We first show that the sequence of actions γ' returned by Algorithm ?? is a sequence of send actions that takes A_a of \mathcal{S}_3 from state ℓ^0 to ℓ^{acc} (note that this is not the final state of A_a , see Fig. ?? for an example). We then show that γ' is a FIFO sequence, and prove that there exists an accepting run of \mathcal{S}_3 in which A_a starts by executing exactly the sequence of actions in γ' . Finally, we show that Algorithm ?? always terminates.

(\Leftarrow) Let μ be an accepting run of \mathcal{S}_3 , from which we show that it is easy to build a sequence of actions γ taken by A in an accepting run of \mathcal{S}_1 . Let γ' be the sequence of send actions taken by A_a in the accepting run μ . The first step is to show that γ' is a FIFO sequence. The three automata A_a , A_b , and A_c are built so to ensure that γ' is always a FIFO sequence. This is closely related to the shape of the MSCs associated to accepting runs of \mathcal{S}_3 ; these MSCs exploit the same kind of pattern seen in Section ?? to bounce messages back and forth between the three processes. We then prove that, if we ignore actions related to dummy messages in γ' and interpret each $a!b(m)$ and $a!c(m)$ action as $!m$ and

? m , we get a sequence of actions γ that takes A from its initial state ℓ^0 to its final state ℓ^{acc} in an accepting run of \mathcal{S}_1 . \diamond

Algorithm 1 Let σ be an accepting run of \mathcal{S}_1 , and α^σ be the sequence of n actions taken by A in σ . We use $\alpha^\sigma(i)$ to denote the i -th action of α^σ .

```

1:  $\gamma' \leftarrow$  empty list
2: Queue  $\leftarrow$  empty queue
3: for  $i$  from 1 to  $n$  do
4:    $action \leftarrow \alpha^\sigma(i)$ 
5:   if  $action = !x$  then
6:     while  $first(Queue) = D$  do
7:       add  $a!c(D)$  to  $\gamma'$ 
8:       dequeue  $D$  from Queue
9:     end while
10:    add  $a!b(x)$  to  $\gamma'$ 
11:    enqueue  $x$  in Queue
12:  else if  $action = ?x$  then
13:    add  $a!c(x)$  to  $\gamma'$ 
14:    dequeue  $x$  from Queue
15:    add  $a!b(D)$  to  $\gamma'$ 
16:    enqueue  $D$  in Queue
17:  if Queue does not contain only  $D$ 
18:    then
19:      while  $first(Queue) = D$  do
20:        add  $a!c(D)$  to  $\gamma'$ 
21:        dequeue  $D$  from Queue
22:        add  $a!b(D)$  to  $\gamma'$ 
23:        enqueue  $D$  in Queue
24:      end while
25:    end if
26:  end for
27: while  $first(Queue) = D$  do
28:   add  $a!c(D)$  to  $\gamma'$ 
29:   dequeue  $D$  from Queue
30: end while
31: return  $\gamma'$ ;

```

The following result immediately follows from Lemma ?? and Theorem ??.

Theorem 4. *The emptiness problem for weakly synchronous communicating systems with three processes is undecidable.*

Notice that our results extend to causally ordered (CO) communication, since an MSC is weakly synchronous if and only if it is weakly synchronous CO (see Appendix ??).

Corollary 1. *The emptiness problem for causal order communicating systems with three processes is undecidable.*

5 Conclusion

We showed the undecidability of the reachability of a configuration for weakly synchronous systems with three processes or more. The main contribution lies in the technique used to achieve this result. We first show that the treewidth of the class of weakly synchronous MSCs is unbounded, by proving that it is always possible to build such an MSC with an arbitrarily large grid as minor. Then, a similar construction is employed to provide an encoding of a FIFO automaton into a weakly synchronous system with three processes, allowing to show that reachability of a configuration is undecidable.

A Tree-decomposition and treewidth

Definition 5. A tree-decomposition of a graph $G = (V, E)$ is a pair $(T, \mathcal{X} = \{X_t \mid t \in V(T)\})$ such that T is a tree, and \mathcal{X} is a set of subsets of V , one for each node of T , such that: (i) $\bigcup_{t \in V(T)} X_t = V(G)$; (ii) for every $\{u, v\} \in E(G)$, there exists $t \in V(T)$ such that $u, v \in X_t$; (iii) for every $v \in V(G)$, the set $\{t \in V(T) \mid v \in X_t\}$ induces a subtree of T .

The width of a tree decomposition (T, \mathcal{X}) is $\max_{t \in V(T)} |X_t| - 1$, i.e., the size of the largest set V minus one. The treewidth $tw(G)$ of G is the minimum width over all possible tree decompositions of G .

The following well-known result in graph theory gives a connection between the notions of treewidth and minor.

Claim. [?] If G is a minor of H , then $tw(G) \leq tw(H)$.

B Automata with epsilon transitions

Given a communicating automaton A , we build an equivalent one with epsilon-transitions such that (i) from each state there are either only epsilon-transitions or a single transition labeled with a letter from Σ , (ii) and there are no states with a transition that lands in the same state (i.e., a self-loop).

Let $A = (Loc, \delta, \ell^0, \ell^{acc})$ be a communicating automaton for process p . Its encoding into an automaton with single non-epsilon transitions is the automaton $A^\varepsilon = (Loc^\varepsilon, \delta^\varepsilon, \ell^0, \ell^{acc})$ where

- $Loc^\varepsilon = Loc \cup \{\ell^t \mid t \in \delta\}$
- $\delta^\varepsilon = \{(\ell, \varepsilon, \ell_t), (\ell_t, a, \ell') \mid t = (\ell, a, \ell') \in \delta\}$

Let \mathcal{S}^ε be the system obtained from \mathcal{S} where each for each of the processes $p \in \mathbb{P}$ we take the corresponding encoding A_p^ε .

Immediately from the definition of accepting run we can see that $L(\mathcal{S}^\varepsilon) = L(\mathcal{S})$.

C Proofs for Section ??

This section is devoted to the proof of Theorem ??, which we restate below.

Theorem ??. *There is an accepting run of \mathcal{S}_1 if and only if there is an accepting run of \mathcal{S}_3 .*

Proof. (\Rightarrow) Follows from Lemma ??, which uses Lemmata ?? to ??.

(\Leftarrow) Follows from Lemma ??, which uses Lemmata ?? and ??. All Lemmata are stated and proved below. \square

Let μ be an accepting run of a system \mathcal{S} , over a set \mathbb{M} of messages and a set \mathbb{P} of processes, on an MSC $M = (\mathcal{E}, \rightarrow, \triangleleft, \lambda)$. For $p \in \mathbb{P}$, we will use α_p^μ to denote the sequence of actions, ignoring ε -actions, taken by A_p in the run μ ; more formally, α_p^μ is the sequence of actions in $\{\text{action}(t) \mid t \in \mu(e), e \in \mathcal{E}_p, \text{action}(t) \neq \varepsilon\}$, ordered according to the \rightarrow_P relation, i.e., given $a_1 = \text{action}(\mu(e_1))$, $a_2 = \text{action}(\mu(e_2))$, such that $e_1 \rightarrow_p e_2$, then a_1 is right before a_2 in α_p^μ , which we abbreviate as $a_1 \xrightarrow{P} a_2$ (or simply $a_1 \dashrightarrow a_2$, when the process p is clear from the context). The lightweight notation α_p will be used when the run μ is clear from the context, and we omit p when the system only has one process. The j -th action in α_p^μ will be denoted by $\alpha_p^\mu(j)$. In this section, when talking about an accepting run of a system $\mathcal{S} = ((A_p)_{p \in \mathbb{P}})$ on an MSC M , we will often not even mention M , and only focus on the sequence of actions taken by the automata A_p .

Let $\mathcal{S}_1 = (A)$ be any communicating system with a single process and one queue, and $\mathcal{S}_3 = (A_a, A_b, A_c)$ be the weakly synchronous system obtained from \mathcal{S}_1 with the reduction described in Section ???. For $p \in \{a, b, c\}$, we use $!\alpha_p^\mu$ and $?\alpha_p^\mu$ to denote the sequence of send actions and, respectively, receive actions taken by A_p . Note that $\alpha_p^\mu = !\alpha_p^\mu + ?\alpha_p^\mu$, where $+$ is the concatenation of two sequences, since \mathcal{S}_3 is a weakly synchronous system with one phase.

Lemma 4. *Let $\gamma = a_1 \dots a_k$ be a sequence of send actions taken by A_a to get from ℓ_a^0 to $\ell_a^?$, where A_a is also allowed to take extra ε -actions in any state. If γ is a FIFO sequence, there is an accepting run μ of \mathcal{S}_3 such that $!\alpha_a^\mu = \gamma$.*

Proof. Suppose γ is a FIFO sequence. It follows that a must send messages to b in the same order $X = m_1 \dots m_k$ as a sends messages to c . We will now build an accepting run μ of \mathcal{S}_3 such that $!\alpha_a^\mu = \gamma$. For A_b (A_c), we can construct $!\alpha_b^\mu$ ($!\alpha_c^\mu$) by sending messages to a and c (b and a) in order X . For $p \in \{a, b, c\}$, $?\alpha_p^\mu$ is constructed by receiving messages from the other two processes in order X . α_a^μ , α_b^μ , and α_c^μ all lead the corresponding automaton to the final state, and all messages that are sent are also received, so μ is an accepting run of \mathcal{S}_3 . \square

We now present Algorithm ??, which essentially takes an accepting run of \mathcal{S}_1 , and returns an accepting run for \mathcal{S}_3 . The correctness of the algorithm is proved in a few steps.

Lemma 5. *Let γ' be the sequence of actions returned by one execution of Algorithm ??. Then, γ' is a sequence of actions that takes A_a of \mathcal{S}_3 from state ℓ^0 to ℓ^{acc} .*

Proof. Let σ be the accepting run of \mathcal{S}_1 that is given as an input to the algorithm. Suppose that both A in \mathcal{S}_1 and A_a in \mathcal{S}_3 start in their initial states and, every time we read an action from α^σ in the algorithm (line 4), we take that action in the current state of A^4 ; we know that this is always possible since α^σ is by definition a sequence of actions⁵ that takes A from its initial state ℓ^0 to its final

⁴ Or from a state of A that is reachable from the current one using only ε -transitions.

⁵ Possibly interleaved by some ε actions.

state ℓ^{acc} . Similarly, each time that an action is added by the algorithm to the sequence γ' , we take that action in the current state of A_a in $\mathcal{S}_3^{??}$, provided that there exists a transition with such an action; we show that there always is such a transition, and that the sequence of actions γ' built by the algorithm^{??} will take A_a from its initial state ℓ^0 to the state ℓ^{acc} (note that ℓ^{acc} is not the final state of A_a , by definition). We show by induction that, right before each iteration of the for loop, A and A_a can always be in the same state (the correspondence between states of A and the states of A_a is given by the definition of A_a); in particular, we show that before the i -th iteration of the for loop, they can always both be in state ℓ^{i-1} , which is the state from which A will take the action $\alpha^\sigma(i)$. For the base case, we are right before the first execution of the for loop. A is in a state ℓ that is either the initial state ℓ^0 or some state reachable with only ε -transitions from ℓ^0 . In both cases, by construction, A_a can also be in the same state ℓ . For the inductive step, we assume that A and A_a are both in the same state ℓ^{i-1} before executing the i -th iteration of the for loop (where ℓ^{i-1} is the state in which A is ready to take the $\alpha^\sigma(i)$ action), and we show that at the end of the i -th iteration both A and A_a end up in state ℓ^i . There are two main possibilities for the i -th iteration of the for loop:

- $\alpha^\sigma(i) = !x$, so the *if* at line 5 is entered. This means that ℓ^{i-1} in A is a state that has an outgoing transition with the send action $!x$. By construction, since A_a is also in ℓ^{i-1} , it can take an unlimited number of $a!c(D)$ actions, followed by a $a!b(x)$ action. These are exactly the kind of actions added to γ' by the *if* that starts at line 5.
- $\alpha^\sigma(i) = ?x$, so the *if* at line 12 is entered. A is then in a state ℓ^{i-1} that has an outgoing transition with the receive action $?x$. By construction, A_a in state ℓ^{i-1} can take the $a!c(x)$ action, followed by a $a!b(D)$ action; after that, A_a can take the consecutive pair of actions $a!c(D)$ and $a!b(D)$ any number of times. These are exactly the kind of actions added to γ' by the *if* that starts at line 12.

In both cases, after taking action $\alpha^\sigma(i)$, A gets to state ℓ^i (possibly using some additional ε -actions), ready for the next execution of the for loop; after taking the actions added by the algorithm to γ' , A_a can also get to state ℓ^i , by construction. After the last iteration of the for loop, A and A_a will both be in state $\ell^n = \ell^{acc}$. By construction, A_a can take an unlimited number of $a!c(D)$ actions in this state, which are the only kind of actions that can be added by the algorithm during the final while loop (line 27). \square

Lemma 6. *Let γ' be the sequence of actions returned by one execution of Algorithm ^{??}. γ' is a valid FIFO sequence.*

Proof. Each time that a $a!b(m)$ action gets added to γ' by the algorithm, m is enqueued, and each time a $a!c(m)$ action is added to γ' , m is dequeued. Our claim directly follows (the behavior of a queue is naturally FIFO).

Lemma 7. *Algorithm ^{??} always terminates.*

Proof. The only ways in which the algorithm does not terminate are either (i) if it blocks when trying to dequeue a message m that is not the first in *Queue*, or (ii) if a while loop runs forever. We show that neither ever happens. Let us first focus on the specific case of line 14, when a message x is dequeued. Each time the algorithm encounters a send action $!x$ in α^σ , message x is enqueued; each time it encounters a receive action $?x$, message x is dequeued. There are no other occasions in which a normal message (i.e., not a dummy message D) is enqueued or dequeued. By definition, α^σ is a valid FIFO sequence for a single queue, so each time that the algorithm reads a receive action $?x$ and gets to line 14, message x must be the first in the queue, unless there are some dummy messages D before. We show that this is impossible. Suppose, by contradiction, that during the i -th iteration of the for loop, $\alpha^\sigma(i) = ?b$ and the algorithm blocks at line 14, because there are some D messages before b in the queue; these D messages must have been enqueued during previous iterations of the for loop. Note that $i > 1$, since the first action in α^σ cannot be a receive action, and the algorithm gets to line 14 only when it reads a receive action. Consider the previous $(i - 1)$ -th iteration of the algorithm, where $\alpha^\sigma(i - 1)$ could either be a send or receive action:

- In the first case, we would have entered the while loop at line 6, which would have dequeued all the D messages on the top of the queue, leaving b as the first one when entering the i -th iteration of the for loop, therefore leading to a contradiction.
- In the second case, $\alpha^\sigma(i - 1)$ is a receive action, so we would have entered the *if* at line 17 (since b is in the queue by hypothesis), and the while loop right after at line 18; this loop also dequeues all D messages and puts them back in the queue, leaving b as the first one when entering the i -th iteration of the for loop, leading again to a contradiction.

We showed that the *dequeue* operation at line 14 never blocks.

Now, we consider the cases in which a D message is dequeued (line 8, 20, and 29), and show that the algorithm never blocks. In all of these cases, we are in a while loop that is entered only if the message at the top of the queue is D , therefore the algorithm will never block. The last thing to show is that no while loop will run forever. To do this, we first show that, at any point of the algorithm, the number of D messages in the queue is at most n . Note that a D message can only be added to the queue at lines 16 and 22. In the case of line 22, a D message is enqueued only after another D message was dequeued (line 20), so the total number of D messages in the queue does not change each time that an iteration of the while loop at line 18 is executed. Line 16 is therefore the only one that can effectively increase the number of D messages in the queue, and can only be executed at most once per iteration of the for loop. The number of D messages in the queue at any time can then be at most n (in particular, it is finite). It follows directly that the while loops at line 6 and 27 will never run forever. We get to the while loop at line 18 only if there is at least one non-dummy message x in the queue; since the number of D in the queue is finite, the loop will run a finite number of time before encountering message x at the top of the queue. \square

Lemma 8. *If there is an accepting run σ of \mathcal{S}_1 , then there is an accepting run μ of \mathcal{S}_3 .*

Proof. By Lemma ??, Algorithm ?? always terminates and returns γ' . By Lemma ??, γ' is a sequence of actions that takes A_a from ℓ_a^0 to $\ell_a^?$. Lemma ?? shows that γ' is a FIFO sequence, so we can finally use Lemma ?? to claim that there is an accepting run μ of \mathcal{S}_3 in which $! \alpha_a^\mu = \gamma'$. \square

Lemma 9. *Let μ be an accepting run of \mathcal{S}_3 . In $! \alpha_a^\mu$, there is an equal number of messages sent to b and to c . Moreover, in $! \alpha_a^\mu$, if x is the i -th message sent to b , and y the i -th message sent to c , then $x = y$.*

Proof. By construction, in an accepting run of \mathcal{S}_3 , A_b must send messages in the same order and in the same number to the other two processes, in order not to block and to reach the state $\ell_b^?$, in which it is ready to start receiving messages. The same goes for A_c (but not for A_a). Also, once A_b gets to state $\ell_b^?$, it must receive messages from the other two processes in the same number (let it be n) and in the same order, so not to block and to reach the final state; this means that a and c must send exactly n messages to b ⁶. The same kind of reasoning holds for A_a and A_c , i.e., each process receives messages in the same order and in the same number from the other two processes in an accepting run of \mathcal{S}_3 . We now show that the number of messages sent by a to b (let it be n_1) and by a to c (let it be n_2) is the same. Suppose, by contradiction, that $n_1 \neq n_2$ in an accepting run of \mathcal{S}_3 . Based on the above, n_1 is also the number of messages sent by c to b , and by c to a ; similarly, n_2 is the number of messages sent by b to c , and by b to a . We then have that a receives n_1 messages from c , and n_2 messages from b . We said that we must have $n_1 = n_2$ in an accepting run of \mathcal{S}_3 , hence the contradiction. The second part of the lemma essentially says that, for every accepting run of \mathcal{S}_3 , the order in which a sends messages to b is the same as the order in which a sends messages to c . By contradiction, suppose a sends messages to b following the order $X = m_1 \dots m_k$, and messages to c following another order Y , such that $X \neq Y$. Based on the above, X is also the order in which messages are sent by c to b , and by c to a ; similarly, Y is the order in which messages are sent by b to c , and by b to a . We then have that a receives messages in order X from c , and in order Y from b . We said that we must have $X = Y$ in an accepting run of \mathcal{S}_3 , hence the contradiction. \square

In order to make the following proofs more readable, we introduce some simplified terminology. Let μ be an accepting run of \mathcal{S}_3 . In $! \alpha_a^\mu$, we will often refer to send actions addressed to b as “sends”, and to send actions addressed to c as “receipts” (it follows from the way \mathcal{S}_3 was built from \mathcal{S}_1). Additionally, in $! \alpha_a^\mu$, we will refer to the i -th send action to c as the matching receipt for the i -th send action to b (which, in turn, will be referred to as the matching send for the i -th send action

⁶ By the definition of accepting run, A_b has to receive all messages that were sent by the other two processes before moving to the final state, since we cannot have some messages sent to b that are not received.

to c). For example, let $!\alpha_a^\mu = a!c(x) a!b(x) a!b(y) a!b(y) a!b(z) a!c(y) a!c(y) a!c(z)$ (note that it respects Lemma ??): we will refer to the first $a!c(x)$ action as the receipt of the first $a!b(x)$ action, and similarly to $a!c(z)$ as the receipt of the only $a!b(z)$ action in $!\alpha_a^\mu$ (note that $a!c(z)$ is the 4th send action to c , and $a!b(z)$ is the 4th send action to b).

Lemma 10. *Let μ be an accepting run of \mathcal{S}_3 . For every message x , we cannot have more $a!c(x)$ actions than $a!b(x)$ actions in any prefix of $!\alpha_a^\mu$.*

Proof. Using the above-mentioned simplified terminology, we could rephrase the lemma as: given an accepting run μ of \mathcal{S}_3 , in $!\alpha_a^\mu$ there cannot be a receipt that appears before its send. By contradiction, suppose there is an accepting run μ of \mathcal{S}_3 in which a receipt appears before its matching send in $!\alpha_a^\mu$. Let us uniquely identify as $lastR_x$ the first such receipt in $!\alpha_a^\mu$, and as $lastS_x$ its matching send (we have $lastR_x \dashrightarrow^+ lastS_x$ in $!\alpha_a^\mu$). According to our reduction rules, we must send a dummy message right after $lastR_x$ in $!\alpha_a^\mu$. We will uniquely identify this action as $lastS_D$, and its receipt as $lastR_D$. There are two possibilities: either (i) $lastS_D \dashrightarrow^+ lastR_D$ or (ii) $lastR_D \dashrightarrow^+ lastS_D$. The first case leads to a contradiction, because we would have $lastS_D \dashrightarrow^+ lastS_x$ and $lastR_x \dashrightarrow^+ lastR_D$, which violates Lemma ?? (message D is sent to b before message x , but D is sent to c after x). We then consider the second scenario, in which $lastR_D \dashrightarrow^+ lastS_D$. According to the implementation of \mathcal{S}_3 (see reduction rules), a receipt of a dummy message, such as $lastR_D$, can only happen either (i) before a send, or (ii) somewhere after a receipt (in any case, before the next non dummy-related action)⁷. The first case leads to a contradiction. Let s be the above-mentioned send action and r its receipt; we must have $s \dashrightarrow^+ r$, since $lastR_x$ was chosen as the first receipt that appears in $!\alpha_a^\mu$ before its send, but this violates again Lemma ?? (we would have $s \dashrightarrow^+ lastS_D$ and $lastR_D \dashrightarrow^+ r$). We then consider the second scenario, in which $lastR_D$ happens somewhere after a receipt of a message y , which we uniquely identify as R_y . Let S_y be the matching send of R_y . For the same reason as before, $S_y \dashrightarrow^+ R_y$. According to our reduction rules, between R_y and $lastR_D$ there could be an arbitrary large sequence of alternating $a!b(D)$ and $a!c(D)$ actions, where the last $a!c(D)$ is exactly $lastR_D$. In any case, if there are $k + 1$ dummy messages sent between R_y and $lastR_D$, there must be k dummy messages received (excluding the end points). Let S_D be any send of these dummy messages. Note that the matching receipt of S_D (uniquely identified as R_D) cannot be neither (i) after $lastR_D$, nor (ii) before R_y , since both would again violate Lemma ??: in the first case, we would have $S_D \dashrightarrow^+ lastS_D$ and $lastR_D \dashrightarrow^+ R_D$, whereas in the second case $S_y \dashrightarrow^+ S_D$ (since $S_y \dashrightarrow^+ R_y \dashrightarrow^+ S_D$) and $R_D \dashrightarrow^+ R_y$. This means that any of the k sends of dummy messages between R_y and $lastR_D$ must have its matching receipt also between R_y and $lastR_D$ (end points excluded); this is impossible, since between R_y and $lastR_D$ there are only k dummy messages

⁷ When not specified, actions do not refer to dummy messages. For example, “can only happen before a send” refers to the sending of a non-dummy message.

received and $k + 1$ dummy messages sent, so at least one send will not have its matching receipt. \square

Lemma 11. *If there is an accepting run μ of \mathcal{S}_3 , then there is an accepting run σ of \mathcal{S}_1 .*

Proof. Given an accepting run μ of \mathcal{S}_3 , Algorithm ?? always returns a sequence of actions α^σ for an accepting run σ of \mathcal{S}_1 . The proof is very similar to that of Lemma ??, but much easier; therefore, we only describe the main intuition without dealing with most of the formalism. First, the algorithm removes all actions related to dummy messages from $!\alpha_p^\mu$, and creates the sequence seq ; then, it returns the sequence γ , which is identical to seq , except that $a!b(x)$ and $a!c(x)$ actions are rewritten as $!x$ and $?x$, respectively. Let $l^0 \dots l^?$ be the sequence of states traversed by A_a while taking the actions in $!\alpha_a^\mu$, ignoring states in which the only outgoing transitions have a $a!b(D)$ or $a!c(D)$ action; more specifically, these are the intermediate states introduced by the first reduction rule, which do not have a one-to-one correspondence with states of A . Note that, by Lemmas ?? and ??, seq and, therefore, γ , are FIFO sequences. By construction, it is now not difficult to see that the sequence of actions γ takes A from l^0 to $l^?$ where $l^?$ is its final state. After all, just by looking at how A_a is constructed, it is clear that a sequence of send actions in A_a , when removing dummy messages actions and interpreting $a!b(x)$ and $a!c(x)$ as $!x$ and $?x$, also represents a valid sequence for A , as long as it is a valid FIFO sequence (otherwise some receive actions in A might block trying to read a message that is not at the top of the queue). \square

Algorithm 2 Let μ be an accepting run of \mathcal{S}_3 , and $!\alpha_a^\mu$ be the sequence of *size* send actions taken by A_a in μ . $!\alpha_a^\mu(i)$ denotes the i -th action of $!\alpha_a^\mu$.

<pre> 1: $seq \leftarrow !\alpha_a^\mu$ 2: for i from 1 to $size$ do 3: $action \leftarrow !\alpha_a^\mu(i)$ 4: if $action = a!b(D)$ or $action = a!c(D)$ then 5: remove $action$ from seq 6: end if 7: end for 8: $n \leftarrow length(seq)$ 9: $\gamma \leftarrow$ empty list </pre>	<pre> 10: for i from 1 to n do 11: $action \leftarrow seq(i)$ 12: if $action = a!b(x)$ then 13: add $!x$ to γ 14: else if $action = a!c(x)$ then 15: add $?x$ to γ 16: end if 17: end for 18: return γ; </pre>
--	--

D Weakly synchronous causally ordered MSCs

We recall here the definition of causally ordered (CO) MSC, borrowed from [?].

Definition 6 (CO MSC). An MSC $M = (\mathcal{E}, \rightarrow, \triangleleft, \lambda)$ is causally ordered if, for any two send events s and s' , such that $\lambda(s) \in \text{Send}(_, q, _)$, $\lambda(s') \in \text{Send}(_, q, _)$, and $s \leq_{hb} s'$:

- either $s, s' \in \text{Matched}(M)$ and $r \rightarrow^* r'$, with r and r' receive events such that $s \triangleleft r$ and $s' \triangleleft r'$.
- or $s' \in \text{Unm}(M)$.

An MSC is weakly synchronous CO if it is a weakly synchronous MSC and a CO MSC.

Theorem 5. An MSC is weakly synchronous CO if and only if it is weakly synchronous p2p.

Proof. (\Leftarrow) Let M be a weakly synchronous CO MSC. M is weakly synchronous and is also p2p, since each CO MSC is a p2p MSC.

(\Rightarrow) Let M be a weakly synchronous p2p MSC. By contradiction, suppose it is not causally ordered, which means that there exist two send events s and s' addressed to the same process, such that $s \leq_M s'$, and one of the following holds:

- $r' \rightarrow^+ r$, where $s \triangleleft r$ and $s' \triangleleft r'$. Note that s and s' cannot be executed by the same process, otherwise M would not even be p2p. Since $s \leq_M s'$, there is a 'chain' of events that causally links s to s' . Note that, in this chain, there must exist a receive event r'' and a send event s'' such that $r'' \rightarrow^+ s''$ (otherwise s and s' could not be causally related). We now have a send event s'' that is executed after a receive event r'' by the same process. Note that r'' and s'' cannot be in two distinct phases of the weakly synchronous MSC M , since s and r (matching events) must be in the same phase, and we have that $s \leq_M r'' \rightarrow^+ s'' \leq_M s' \triangleleft r' \rightarrow^+ r$ (i.e., all these events between s and r must be part of the same phase).
- s in unmatched, and $s' \triangleleft r'$. As before, note that s and s' cannot be executed by the same process, otherwise M would not even be p2p.