

Towards Generalised Half-Duplex Systems*

Cinzia Di Giusto

Loïc Germerie Guizouarn

Etienne Lozes

Université Nice Côte d'Azur

CNRS, I3S

Sophia Antipolis, France

{cinzia.di-giusto,loic.germerie-guizouarn,etienne.lozes}@univ-cotedazur.fr

FIFO automata are finite state machines communicating through FIFO queues. They can be used for instance to model distributed protocols. Due to the unboundedness of the FIFO queues, several verification problems are undecidable for these systems. In order to model-check such systems, one may look for decidable subclasses of FIFO systems. Binary half-duplex systems are systems of two FIFO automata exchanging over a half-duplex channel. They were studied by Cece and Finkel who established the decidability in polynomial time of several properties. These authors also identified some problems in generalising half-duplex systems to multi-party communications. We introduce greedy systems, as a candidate to generalise binary half-duplex systems. We show that greedy systems retain the same good properties as binary half-duplex systems, and that, in the setting of mailbox communications, greedy systems are quite closely related to a multiparty generalisation of half-duplex systems.

1 Introduction

FIFO automata, also known as asynchronous communicating automata (i.e., finite state automata that exchange messages via FIFO queues) are an interesting formalism for modeling distributed protocols. In their most general formulation, these automata are Turing powerful, and in order to be able to model check them it is necessary to reduce their expressiveness.

Binary half-duplex systems, introduced by Cece and Finkel [5], are systems with two participants and a bidirectional channel formed of two FIFO queues, such that communication happens only in one direction at a time. The stereotypical half-duplex device is the walkie-talkie (or the CB). In several applications, in particular when FIFO buffers are bounded and sends may be blocking, half-duplex communications are considered a good practice to avoid send-send deadlocks. Language support for enforcing this discipline of communication includes, for instance, binary session types [14, 15] or Sing# channel contracts [10, 20].

In [5], Cece and Finkel show that (1) whether a system is half-duplex is decidable in polynomial time, (2) the set of reachable configurations is regular, and (3) properties like progress and boundedness are decidable in polynomial time. Cece and Finkel also present two possible notions ‘multiparty half-duplex’ systems generalizing their class to systems of any number of machines for p2p communications (one FIFO queue per pair of machine).

The first generalisation involves assuming that at most one queue over all queues is non-empty at any time. This generalisation preserves decidability but is very restrictive. The second generalisation restricts the communications between each pair of participants to half-duplex communications, that is only one buffer per bidirectional channel can be used simultaneously. This generalisation however does

*This work has been supported by the French government, through the EUR DS4H Investments in the Future project managed by the National Research Agency (ANR) with the reference number ANR-17-EURE- 0004.

not preserve decidability: the systems captured by this definition form a Turing powerful class. In fact, with just three machines it is possible to mimic the tape of a Turing machine.

It could be believed that these results end the discussion about multi-party half-duplex systems. In this work, we claim conversely that there is another natural and relevant notion of multi-party half-duplex communications that allows us to generalise the results of Cece and Finkel. We introduce *greedy systems*, which are systems for which all executions can be rescheduled in such a way that all receptions are immediately preceded by their corresponding send. This notion is quite natural, and closely related to other notions like synchronisability [3], 1-synchronous systems [4], or existentially 1-bounded system [12] (see Section 6 for a detailed discussion).

In this work, we establish the following results:

1. whether a system is greedy is decidable in polynomial time (when the number of processes is fixed);
2. for greedy systems, all regular safety properties, which includes reachability, absence of unspecified receptions, progress, and boundedness are decidable in polynomial time.
3. we generalize binary half-duplex systems to multiparty mailbox half-duplex systems and we show that (1) mailbox half-duplex systems are greedy, and (2) greedy systems without orphan messages, at least in the binary case, are half-duplex.

The first result follows from techniques developed by Bouajjani *et al* [4] for k -synchronous systems. The main challenge here is that we address a more general model of communicating systems that encompasses both mailbox and p2p communications, but also allows any form of sharing of buffers among processes. The second result is based on an approach that, to the best of our knowledge, is new, although it borrows from some general principles from regular model-checking. The challenge is that, unlike for binary half-duplex systems, the reachability set of greedy systems is not regular, which complicates how automata-based techniques can be used to solve regular safety. The third contribution aims at answering, although incompletely, the question we would like to address with this work: what is a relevant notion of multi-party half-duplex systems?

Outline. The paper is organised as follows: Section 2 introduces communicating automata and systems. Section 3 defines greedy systems and establishes the decidability of the greediness of a system. Section 4 discusses regular safety for greedy systems. Section 5 compares greedy systems and half-duplex systems, first in the binary setting, then in the multi-party setting, by introducing the notion of mailbox half-duplex systems. Finally, Section 6 concludes with some final remarks and discusses related works.

2 Preliminaries

For a finite set S , S^* denotes the set of finite words over S , $w_1 \cdot w_2$ denotes the concatenation of two words, $|w|$ denotes the length of w , and ε denotes the empty word. We assume some familiarity with non-deterministic finite state automata, and we write $\mathcal{L}(\mathcal{A})$ for the language accepted by the automaton \mathcal{A} . For two sets S and I , we write \mathbf{b} (in bold) for an element of S^I , and b_i for the i -th component of \mathbf{b} , so that $\mathbf{b} = (b_i)_{i \in I}$.

A *FIFO automaton* is basically a finite state machine equipped with FIFO queues where transitions are labelled with either queuing or dequeuing actions. More precisely:

Definition 1 (FIFO automaton). A FIFO automaton is a tuple¹ $\mathcal{A} = (L, \mathbb{V}, I, \text{Act}, \delta, l_0)$ where (1) L is a finite set of control states, (2) \mathbb{V} is a finite set of messages, (3) I is a finite set of buffer identifiers, (4) $\text{Act} \subseteq I \times \{!, ?\} \times \mathbb{V}$ is a finite set of actions, (5) $\delta \subseteq L \times \text{Act} \times L$ is the transition relation, and (6) $l_0 \in L$ is the initial control state. The size $|\mathcal{A}|$ of \mathcal{A} is $|L| + |\delta|$.

Actions $a = (i, \dagger, v)$, for $\dagger \in \{!, ?\}$ are also denoted by $i\dagger v$. Given a FIFO automaton $\mathcal{A} = (L, \mathbb{V}, I, \delta, l_0)$, a configuration of \mathcal{A} is a tuple $\gamma = (l, \mathbf{b}) \in L \times \mathbb{V}^I$. The initial configuration is $\gamma_0 = (l_0, \mathbf{b}^0)$ where for all $i \in I$, $b_i^0 = \varepsilon$. A step is a tuple (γ, a, γ') (often written $\gamma \xrightarrow[\mathcal{A}]{a} \gamma'$) where $\gamma = (l, \mathbf{b})$ and $\gamma' = (l', \mathbf{b}')$ are configurations and a is an action, such that the following holds:

- $(l, a, l') \in \delta$,
- if $a = i!v$, then $b'_i = b_i \cdot v$ and $b'_j = b_j$ for all $j \in I \setminus \{i\}$.
- if $a = i?v$, then $b_i = v \cdot b'_i$ and $b'_j = b_j$ for all $j \in I \setminus \{i\}$.

Next, we define systems of FIFO automata. We pick a very general definition, where each FIFO queue might be queued (resp. dequeued) by more than one automaton, and where an automaton might ‘send a message to itself’. Most of the theory can be done in this general setting without extra cost, but we merely have in mind either mailbox systems or p2p systems (defined below).

A FIFO system, later called simply a *system*, is a family $\mathfrak{S} = (\mathcal{A}_p)_{p \in \mathbb{P}}$ of FIFO automata such that all FIFO automata have disjoint sets of actions: for all $p \neq q \in \mathbb{P}$, $\text{Act}_p \cap \text{Act}_q = \emptyset$. Each $p \in \mathbb{P}$ is referred to as a *process*. The condition on the disjointness of the sets of actions helps to identify the process that is responsible for a given action: for an action a , we write $\text{process}(a)$ to denote the unique p (when it exists) such that $a \in \text{Act}_p$.

Let us now define mailbox and p2p systems. Informally, a p2p system is a system in which each pair of processes has a dedicated buffer for exchanging messages. Instead, for mailbox communication, each process receives messages from all other processes in a single buffer. Let us now formally define these notions. To this aim, it will be useful to identify what are the buffers an automaton queues in (resp. dequeues from). Hence, for a given FIFO automaton $\mathcal{A}_p = (L_p, \mathbb{V}_p, I_p, \text{Act}_p, \delta_p, l_{0,p})$, and for $\dagger \in \{!, ?\}$, we write I_p^\dagger for the set of buffer identifiers i such that there exists $v \in \mathbb{V}_p$ such that $i\dagger v \in \text{Act}_p$. A system $\mathfrak{S} = (\mathcal{A}_p)_{p \in \mathbb{P}}$ is *p2p* if for all $p \in \mathbb{P}$, $I_p \subseteq \mathbb{P}^2$, $I_p^! = \{p\} \times (\mathbb{P} \setminus \{p\})$, and $I_p^? = (\mathbb{P} \setminus \{p\}) \times \{p\}$. A system $\mathfrak{S} = (\mathcal{A}_p)_{p \in \mathbb{P}}$ is a *mailbox system* if for all $p \in \mathbb{P}$, $I_p \subseteq \mathbb{P}$, $I_p^! = \mathbb{P} \setminus \{p\}$ and $I_p^? = \{p\}$. Thus in a p2p system with n processes, there are at most $n(n-1)$ buffers, and in a mailbox system with n processes there are at most n buffers. A *binary system* is a system $\mathfrak{S} = (\mathcal{A}_p)_{p \in \mathbb{P}}$ such that $\mathbb{P} = \{1, 2\}$ and for all $p \in \mathbb{P}$, $I_p^! = \{3-p\} = I_{3-p}^?$; note that a binary system is both p2p and mailbox. We sometimes use a more handy notation for actions of a mailbox (resp. p2p) system: if $\text{process}(i!v) = p$ and $\text{process}(i?v) = q$, we sometimes write $!v^{p \rightarrow q}$ instead of $i!v$ and $?v^{p \rightarrow q}$ instead of $i?v$.

Example 1 (FIFO Automata). Figure 1 shows a graphical representation of a FIFO system, borrowed from [2]. This system represents a protocol between a client, a server and a database logging requests from the client and the server. In this protocol, a client can log something on the database or send requests to the server, when those requests are satisfied the server logs them in a database. Each automaton is equipped with a buffer in which it receives messages from all other participants: this system is an example of a mailbox system. To improve readability of the graphical representation, we refer to the buffers with the initial of the automaton to which they are associated. For example, s is the buffer into which the server can receive messages. This simple system will be used as a running example throughout the paper. \square

¹Note that FIFO automata do not have accepting states, therefore they are not a special case of non-deterministic finite state automaton, and there is not such a thing as “the language of a FIFO automaton”.




Figure 1: Client/Server/Database protocol

The size $|\mathfrak{S}|$ of \mathfrak{S} is the sum of the $|\mathcal{A}_p|$. Note that $|\mathbb{P}|$ and $|I|$ are independent from the size of \mathfrak{S} . In particular, when we say that an algorithm is in polynomial time, we mean in time $\mathcal{O}(|\mathfrak{S}|^k)$ for some k that may depend on $|\mathbb{P}|$ and $|I|$.

The FIFO automaton $\text{product}(\mathfrak{S})$ associated with \mathfrak{S} is the standard asynchronous product automaton: $\text{product}(\mathfrak{S}) = (\prod_{p \in \mathbb{P}} L_p, \cup_{p \in \mathbb{P}} \mathbb{V}_p, \cup_{p \in \mathbb{P}} I_p, \delta, \mathbf{l}_0)$ where $\mathbf{l}_0 = (l_{0,p})_{p \in \mathbb{P}}$ and δ is the set of triples $(\mathbf{l}, a, \mathbf{l}')$ for which there exists $p \in \mathbb{P}$ such that $(l_p, a, l'_p) \in \delta_p$ and $l_q = l'_q$ for all $q \in \mathbb{P} \setminus \{p\}$. We often identify \mathfrak{S} and $\text{product}(\mathfrak{S})$ and write for instance $\xrightarrow[\mathfrak{S}]{a}$ instead of $\xrightarrow[\text{product}(\mathfrak{S})]{a}$. Similarly, we say that $\gamma = (\mathbf{l}, \mathbf{b})$ is a configuration of \mathfrak{S} while we mean that γ is a configuration of $\text{product}(\mathfrak{S})$. An execution $e = a_1 \cdots a_n \in \text{Act}^*$ is a sequence of actions. As usual \xRightarrow{e} stands for $\xrightarrow{a_1} \cdots \xrightarrow{a_n}$. We write $\text{executions}(\mathfrak{S})$ for $\{e \in \text{Act}^* \mid \gamma_0 \xRightarrow[\mathfrak{S}]{e} \gamma \text{ for some } \gamma\}$.

Next, we introduce the definition of reachable configuration:

Definition 2 (Reachable configuration). *Let \mathfrak{S} be a system. A configuration γ is reachable if there exists $e \in \text{Act}^*$ such that $\gamma_0 \xRightarrow[\mathfrak{S}]{e} \gamma$. The set of all reachable configurations of \mathfrak{S} is denoted $RS(\mathfrak{S})$.*

Given an execution $e = a_1 \cdots a_n$, we say that $\{j, j'\} \subseteq \{1, \dots, n\}^2$ is a *matching pair* if there exists a buffer identifier i , a message v and natural number k such that (1) $a_j = i!v$, (2) $a_{j'} = i?v$, (3) a_j is the k -th send action on i in e , and (4) $a_{j'}$ is the k -th receive action on i in e . A *communication* of e is either a matching pair $\{j, j'\}$, or a singleton $\{j\}$ such that j does not belong to any matching pair (such a communication is called unmatched). We write $\text{com}(e)$ to denote the set of communications of e .

An execution imposes a total order on the actions. Sometimes, however, it is useful to visualise only the causal dependencies between actions. Message sequence charts [17] are usually used to this aim, as they only depict an order between matched pairs of actions and between actions of the same process. However, message sequence charts do not represent graphically the causal dependencies due to shared buffers, like the ones found in mailbox systems. Here we define *action graphs* that depict all causal dependencies. When considering p2p communications, action graphs and message sequence charts coincide. We say that two actions a_1, a_2 commute if $\text{process}(a_1) \neq \text{process}(a_2)$ and it is not the case that a_1 and a_2 are two actions of the same type on a same buffer: there is no $\dagger \in \{!, ?\}$, $i \in I$ and $v_1, v_2 \in \mathbb{V}$ such that $a_1 = i\dagger v_1$ and $a_2 = i\dagger v_2$.

Definition 3 (Action graph). *Given an execution $e = a_1 \cdots a_n$, the action graph $\text{agraph}(e)$ is the vertex-labeled directed graph $(\{1, \dots, n\}, \prec_e, \lambda_e)$ where $\lambda_e(j) = a_j$ and $j \prec_e j'$ if (1) $j < j'$ and (2) either $a_j, a_{j'}$ do not commute, or $\{j, j'\}$ is a matching pair.*

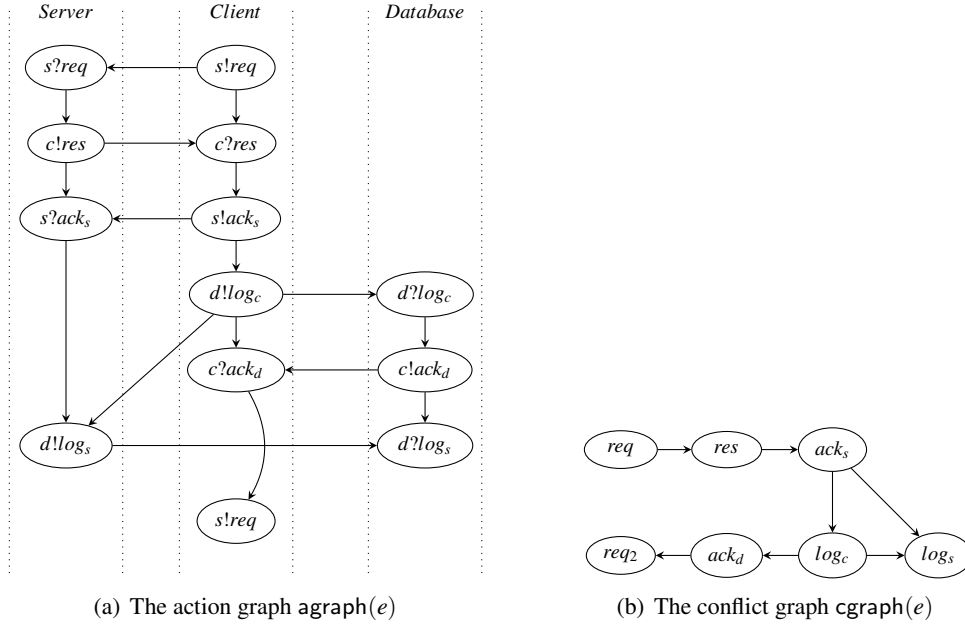


Figure 2: Causal dependencies of an execution of Client/Server/Database protocol

When $j \prec_e j'$, we say that a_j happens before $a_{j'}$. Note that for an execution e with $e \in \text{executions}(\mathfrak{G})$, \prec_e^* is a partial order. Two executions e, e' are causally equivalent, denoted by $e \stackrel{\mathfrak{G}}{\sim} e'$, if their action graphs are isomorphic. Said differently, e, e' are causally equivalent if they are two linearisations of a same happens before partial order. Note that if $\gamma_0 \stackrel{e}{\xrightarrow{\mathfrak{G}}} \gamma$ and $e \stackrel{\mathfrak{G}}{\sim} e'$, then $\gamma_0 \stackrel{e'}{\xrightarrow{\mathfrak{G}}} \gamma$.

Another graphical tool that we will use to talk about equivalent executions is the *conflict graph*, which is intuitively obtained from the action graph by merging matching pairs of vertices.

Definition 4 (Conflict graph). *Given an execution $e = a_1 \cdots a_n$, the conflict graph $\text{cgraph}(e)$ of the execution e is the directed graph $(\text{com}(e), \rightarrow_e)$ where for all communications $c_1, c_2 \in \text{com}(e)$, $c_1 \rightarrow_e c_2$ if there is $j_1 \in c_1$ and $j_2 \in c_2$ such that $j_1 \prec_e j_2$.*

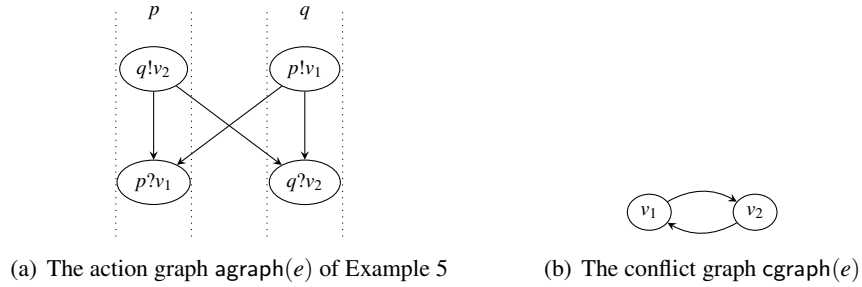
Example 2 (Action and Conflict Graphs). We go back to the system depicted in Figure 1. One of its executions is

$$e = s!req \cdot s?req \cdot c!res \cdot c?res \cdot s!ack_s \cdot s?ack_s \cdot d!log_c \cdot d!log_s \cdot d?log_c \cdot c!ack_d \cdot c?ack_d \cdot s!req \cdot d?log_s$$

Figure 2(a) shows $\text{agrap}(e)$. Actions of the same process are represented vertically between the same dotted lines. As formally explained in Definition 3, an arc from an action a and another a' means that a happens before a' . To ease readability, the arcs that follow from transitivity are omitted. For example, in a given column, there should be an arc between every pair of actions.

Figure 2(b) shows $\text{cgraph}(e)$. To simplify the graph, instead of marking the matching pairs we simply identify them with the message exchanged. Message req_2 represent the second send of req in the execution above.

□

Figure 3: Visual representations of a non-greedy execution e

3 Greedy systems

In this section we introduce greedy systems. Those systems aim at mimicking rendez-vous or synchronous communications by checking whether each execution can be rescheduled to an equivalent one where all receptions immediately follow their corresponding send.

Definition 5 (Greedy system). *An execution e is greedy if all matching pairs are of the form $\{j, j + 1\}$. A system \mathfrak{S} is greedy if for all execution $e \in \text{executions}(\mathfrak{S})$, there exists a greedy execution e' such that $e \stackrel{\mathfrak{S}}{\sim} e'$.*

Example 3. The execution $s!req \cdot s?req \cdot c!res \cdot c?res \cdot s!ack_s \cdot d!log_c \cdot d?log_c$ is greedy, but the execution $s!req \cdot s?req \cdot c!res \cdot c?res \cdot s!ack_s \cdot d!log_c \cdot s?ack_s$ is not greedy, although causally equivalent to a greedy execution. \square

Example 4 (Greedy system). The system in Figure 1 is greedy. Take for instance execution e of the example 2. This execution is not greedy as messages log_c , log_s , and ack_d are not received right after their send. Still since those action can commute and by observing that the conflict graph in Figure 2(b) does not present any cycle (see Lemma 6) below, there exists an equivalent greedy execution e' :

$$e' = s!req \cdot s?req \cdot c!res \cdot c?res \cdot s!ack_s \cdot s?ack_s \cdot d!log_c \cdot d?log_c \cdot c!ack_d \cdot c?ack_d \cdot s!req \cdot d!log_s \cdot d?log_s.$$

\square

Example 5. Consider a system with two processes p and q each sending a message to the other, and whose corresponding receptions only happen after the send (see Figure 3). Then the execution $e = p!v_1 \cdot q!v_2 \cdot p?v_1 \cdot q?v_2$ is not causally equivalent to a greedy execution, therefore the whole system is not greedy. \square

In the remainder of this section, we show that deciding whether a system is greedy is feasible in polynomial time. The proof is in three steps: first, we give a graphical characterisation of the executions that are causally equivalent to greedy executions; second, we show that the non-greediness of a system is revealed by the existence of ‘bad’ executions of a certain shape, called borderline executions. Finally, we show that the graphical characterisation can be exploited to show the regularity of the language of borderline violations, from which we get the decidability of greediness.

Lemma 6. *e is causally equivalent to a greedy execution if and only if $\text{cgraph}(e)$ is acyclic.*

Proof. The left to right implication follows from two observations: first, two causally equivalent executions have isomorphic conflict graphs (because they have isomorphic action graphs), and second, the conflict graph of a greedy execution is acyclic, because for a greedy execution $c_1 \rightarrow_e c_2$ induces

$\min c_1 < \min c_2$. Conversely, let $e = a_1 \cdots a_n$, and assume that $\text{cgraph}(e)$ is acyclic. Let $c_1 \ll \dots \ll c_n$ be a topological order on $\text{com}(e)$, $e' = c_1 \cdots c_n$ be the corresponding greedy execution, and σ be the permutation such that $e' = a_{\sigma(1)} \cdots a_{\sigma(n)}$. The claim is that $e \stackrel{\mathfrak{S}}{\sim} e'$. Let j, j' be two indices of e , and let us show that $j \prec_e j'$ iff $\sigma(j) \prec_{e'} \sigma(j')$. First, $\{j, j'\}$ is a matching pair of e if and only if $\{\sigma(j), \sigma(j')\}$ is a matching pair of e' , which shows the equivalence in that case. Assume now that $\{j, j'\}$ is not a matching pair of e , and let c, c' be the communications of e containing j and j' respectively. Assume also that $j \prec_e j'$, we show that $\sigma(j) \prec_{e'} \sigma(j')$ (the other implication, similar, is omitted). First, $j \prec_e j'$ implies $c \rightarrow_e c'$, which entails $\sigma(c) \rightarrow_{e'} \sigma(c')$ because e and e' have the same conflict graph. Moreover, $j \prec_e j'$ implies that a_j and $a_{j'}$ do not commute, therefore either $\sigma(j) \prec_{e'} \sigma(j')$ or $\sigma(j') \prec_{e'} \sigma(j)$. By contradiction let $\sigma(j') \prec_{e'} \sigma(j)$; then $\sigma(c') \rightarrow_{e'} \sigma(c)$, contradicting the acyclicity of the conflict graph of e' , which ends the proof. \square

Definition 7 (Borderline violations). *An execution $e \in \text{executions}(\mathfrak{S})$ is a borderline violation if (1) e is not causally equivalent to a greedy execution, (2) $e = e_1 \cdot i?v$ for some greedy execution e_1 and receive action $i?v$.*

Example 6 (Borderline violation). An example of a borderline violation for the system whose unique maximal execution is the one of Figure 3 is the execution

$$e = !v_2^{p \rightarrow q} . !v_1^{q \rightarrow p} . ?v_1^{q \rightarrow p} . ?v_2^{p \rightarrow q} .$$

Figure 3 shows its action and conflict graph. The action graph makes it easy to see that any execution e' equivalent to e will require both the send actions to be done before the first reception, therefore at least one reception will not follow its matching send action. \square

Lemma 8. *\mathfrak{S} is greedy if and only if $\text{executions}(\mathfrak{S})$ contains no borderline violation.*

Proof. Obviously, if $\text{executions}(\mathfrak{S})$ contains a borderline violation, \mathfrak{S} is not greedy. Conversely, assume that \mathfrak{S} is not greedy, and let us show that $\text{executions}(\mathfrak{S})$ contains a borderline violation. Let $e \in \text{executions}(\mathfrak{S})$ be an execution that is not causally equivalent to a greedy execution and of minimal length among all such executions. Then $e = e_1 \cdot a$ with e_1 causally equivalent to a greedy execution. Let e'_1 be a greedy execution causally equivalent to e_1 . Then $e' = e'_1 \cdot a \in \text{executions}(\mathfrak{S})$. Moreover, if a is a send action, then e' is greedy, contradicting the fact that e is not causally equivalent to a greedy execution. Therefore, e' is a borderline violation. \square

Let $\Sigma = I \times \{!, !?\} \times \mathbb{V}$ denote the set of communications, and let $\Sigma_? = I \times \{?\} \times \mathbb{V}$ be the set of receive actions. Then a greedy execution can be represented by a word in Σ^* and a borderline violation is represented by a word in $\Sigma^* \cdot \Sigma_?$. So now, we define two non-deterministic finite state automata over $\Sigma \cup \Sigma_?$: the first one accepts all greedy executions of a system, and the second one all borderline violations. We later explain how these automata are used to decide greediness.

Lemma 9. *Let $\mathfrak{S} = (\mathcal{A}_p)_{p \in \mathbb{P}}$ of size n and $\mathbb{V} = \bigcup_{p \in \mathbb{P}} \mathbb{V}_p$, $I = \bigcup_{p \in \mathbb{P}} I_p$ be fixed. There is a non-deterministic finite state automaton \mathcal{A}_{gr} computable in time $\mathcal{O}(|\mathbb{V}| |I|^2 2^{|I|} n^{|\mathbb{P}|+2})$ such that $\mathcal{L}(\mathcal{A}_{gr}) = \{e \cdot i?v \in \Sigma^* \cdot \Sigma_? \mid e \cdot i?v \in \text{executions}(\mathfrak{S}) \text{ and } e \text{ is greedy}\}$.*

Proof. Let $\text{product}(\mathfrak{S}) = (L_{\mathfrak{S}}, \mathbb{V}, I, \text{Act}, \delta_{\mathfrak{S}}, \mathbf{l}_0)$ and let $\mathcal{A}_{gr} = (L_{gr}, \delta_{gr}, l_{gr,0}, F_{gr})$ be the non-deterministic finite state automaton over Σ with $L_{gr} = L_{\mathfrak{S}} \times (\{\mathcal{E}\} \cup I \times \mathbb{V}) \times 2^I \cup \{l_F\}$, $l_{gr,0} = (\mathbf{l}_0, \mathcal{E}, \emptyset)$, $F_{gr} = \{l_F\}$, and the transitions defined as follows. First, while reading a letter $c \in \Sigma$, $((\mathbf{l}, x, S), c, (\mathbf{l}', x', S')) \in \delta_{gr}$ if

- $(\mathbf{l}, \mathbf{b}) \xrightarrow{c}_{\mathfrak{S}} (\mathbf{l}', \mathbf{b}')$ for some \mathbf{b}, \mathbf{b}' such that for all $i \in I$, $b_i \neq \emptyset$ iff $i \in S$, and $b'_i \neq \emptyset$ iff $i \in S'$, and

- one of the two following holds
 - either $x = x'$,
 - or $x = \varepsilon$ and $x' = (i, v)$ and $c = i!v$ and $i \notin S$.

Second, while reading the letter $i?v \in \Sigma_\gamma$, $((\mathbf{l}, x, S), i?v, l_F) \in \delta_{gr}$ if $x = (i, v)$ and $(\mathbf{l}, i?v, \mathbf{l}') \in \delta_{\mathfrak{G}}$ for some $\mathbf{l}' \in L_{\mathfrak{G}}$. Then $\mathcal{L}(\mathcal{A}_{gr})$ is as announced. Moreover, each transition of \mathcal{A}_{gr} can be constructed in constant time, so \mathcal{A}_{gr} can be constructed in time as announced. \square

Lemma 10. *There is a non-deterministic finite state automaton \mathcal{A}_{bv} computable in time $\mathcal{O}(|I|^3|\mathbb{V}|^3)$ such that $\mathcal{L}(\mathcal{A}_{bv}) = \{e \in \Sigma^* \cdot \Sigma_\gamma \mid \text{cgraph}(e) \text{ contains a cycle}\}$.*

Proof. Let $\mathcal{A}_{bv} = (L_{bv}, \delta_{bv}, l_{bv,0}, \{l_{bv,1}\})$ be the non-deterministic finite state automaton over $\Sigma \cup \Sigma_\gamma$ such that $L_{bv} = \{l_{bv,0}, l_{bv,1}\} \cup \Sigma_\gamma \times \Sigma$, and for all $c, c' \in \Sigma$, for all $a \in \Sigma_\gamma$, for all $i \in I, v \in \mathbb{V}$, (1) $(l_{bv,0}, c, l_{bv,0}) \in \delta_{bv}$ (2) $(l_{bv,0}, i!v, (i?v, i!v)) \in \delta_{bv}$, (3) $((a, c), c', (a, c)) \in \delta_{bv}$ (4) $((a, c), c', (a, c')) \in \delta_{bv}$ if $\text{process}(c) \cap \text{process}(c') \neq \emptyset$, and (5) $((i?v, c), i?v, l_{bv,1}) \in \delta_{bv}$ if $\text{process}(c) \cap \text{process}(i?v) \neq \emptyset$. Then $\mathcal{L}(\mathcal{A}_{bv}) = \{e \in \Sigma^* \Sigma_\gamma \mid \text{cgraph}(e) \text{ contains a cycle}\}$. Moreover, each transition of \mathcal{A}_{bv} can be constructed in constant time, so \mathcal{A}_{bv} can be constructed in time as announced. \square

From the computability of the two previous automata, we deduce the decidability of system greediness.

Theorem 11. *Whether a system $\mathfrak{G} = (\mathcal{A}_p)_{p \in \mathbb{P}}$ of size n is greedy is decidable in time $\mathcal{O}(|I|^5|\mathbb{V}|^4 2^{|I|} n^{|\mathbb{P}|+2})$.*

Proof. Let \mathcal{A}_{gr} and \mathcal{A}_{bv} be the two automata defined in Lemmas 9 and 10. By Lemma 6 and by definition of a borderline violation, the set of borderline violations of \mathfrak{G} is $\mathcal{L}(\mathcal{A}_{gr}) \cdot \Sigma_\gamma \cap \mathcal{L}(\mathcal{A}_{bv})$. So, by Lemma 8, \mathfrak{G} is greedy if and only if $\mathcal{L}(\mathcal{A}_{gr}) \cdot \Sigma_\gamma \cap \mathcal{L}(\mathcal{A}_{bv}) = \emptyset$. The claim then directly follows from the fact that emptiness testing for a non-deterministic finite state automaton of size n is in time $\mathcal{O}(n)$. \square

4 Model-Checking Greedy Systems

In this section we explore how to verify various safety properties of greedy systems in polynomial time. Since the reachability set of a greedy system is not regular, it is not obvious that regular safety properties are always decidable. We show that this problem actually is decidable, with a polynomial time complexity under mild assumptions. Then, we list a few regular safety properties that were also considered in other works, in particular for approaches based on session types [16, 8, 22].

4.1 Checking Regular Safety Properties

Let $\mathfrak{G} = (\mathcal{A}_p)_{p \in \mathbb{P}}$ with $I = \bigcup_{p \in \mathbb{P}} I_p = \{1, \dots, |I|\}$ be a system. We identify a word $w = \mathbf{l} \cdot \# \cdot b_1 \cdot \# \cdots \# \cdot b_{|I|} \in L_{\mathfrak{G}} \cdot (\# \cdot \mathbb{V}^*)^{|I|}$ with the configuration $(\mathbf{l}, b_1, \dots, b_{|I|})$. We say that a set of configurations $P(\mathfrak{G})$ is *regular*² if the corresponding set of words coding these configurations is regular. A *property* is a function P that associates to every system \mathfrak{G} a set of configurations $P(\mathfrak{G})$. We say that P is regular if $P(\mathfrak{G})$ is regular for all \mathfrak{G} , and *computable* in time $\mathcal{O}(f(n))$ if a non-deterministic finite state automaton \mathcal{A} accepting $P(\mathfrak{G})$ can be computed in time $\mathcal{O}(f(|\mathfrak{G}|))$. The *P safety problem* is whether a system \mathfrak{G} is such that $RS(\mathfrak{G}) \cap P(\mathfrak{G}) = \emptyset$. Examples of safety problems are discussed below in Section 4.2.

²also called *channel recognizable* by Cece and Finkel [5, Definition 10]

Cece and Finkel showed that, for a binary half-duplex system \mathfrak{G} , $RS(\mathfrak{G})$ is regular and computable in polynomial time [5, Theorem 26]. Since the emptiness of the intersection of two polynomial time computable regular languages is decidable in polynomial time, for any regular polynomial time property P , the P safety problem is decidable in polynomial time for binary half-duplex systems.

For greedy systems, however, the situation is a bit different. Indeed, observe that $RS(\mathfrak{G})$ may be non-regular and even context sensitive (it is easy to define a system with one machine and 3 buffers, such that for some control state l , $RS(\mathfrak{G}) \cap l(\#\mathbb{V}^*)^3 = \{l\#a^n\#b^n\#c^n \mid n \geq 0\}$). So it is not obvious how to decide the emptiness of $RS(\mathfrak{G}) \cap P(\mathfrak{G})$.

Still, the P safety problem remains decidable for a computable regular property P .

Theorem 12. *Let P be a computable regular property. Then, for a given greedy system $\mathfrak{G} = (\mathcal{A}_p)_{p \in \mathbb{P}}$, the P safety problem is decidable. Moreover, if P is computable in time $\mathcal{O}(|\mathfrak{G}|^k)$ for some $k \geq 0$, then the problem is decidable in time $\mathcal{O}(|\mathfrak{G}|^{k+|\mathbb{P}|+2})$.*

Proof. Let \mathfrak{G} be fixed, with $\text{product}(\mathfrak{G}) = (L_{\mathfrak{G}}, \mathbb{V}, I, \text{Act}, \delta_{\mathfrak{G}}, \mathbf{I}_{\mathfrak{G},0})$. Let $\mathcal{A} = (L_{\mathcal{A}}, \delta_{\mathcal{A}}, l_{\mathcal{A},0}, F_{\mathcal{A}})$ be the polynomial time computable non-deterministic finite state automaton over alphabet $L_{\mathfrak{G}} \cup \{\#\} \cup \mathbb{V}$ such that $\mathcal{L}(\mathcal{A}) = P(\mathfrak{G})$. We define an automaton $\mathcal{A}_p = (L_p, \delta_p, L_{p,0}, F_p)$ over the alphabet Σ of communications such that for all greedy execution $e \in \text{executions}(\mathfrak{G})$, $e \in \mathcal{L}(\mathcal{A}_p)$ iff there is $\gamma \in P(\mathfrak{G})$ such that $\gamma \xrightarrow[e]{\mathfrak{G}} \gamma$.

Before defining \mathcal{A}_p formally, let us give some intuitions about how it works on an example. Assume that \mathcal{A}_p reads $e = 1!a \cdot 2!b \cdot 2!c \cdot 1!d$, and that the final configuration γ is $\mathbf{I}\#ad\#c$. While reading e , \mathcal{A}_p should check the existence of an accepting run of \mathcal{A} on γ . When \mathcal{A}_p reads a communication, there are two cases. Either the communication is a matched send (like $2!b$) and therefore it does not contribute to the final configuration, so \mathcal{A}_p merely ignores it. Or the communication is an unmatched send, and it contributes to a piece of the accepting run of γ on \mathcal{A} . However, these pieces of the accepting run of \mathcal{A} are not necessarily consecutive. For instance, in the above execution, a and d are consecutive in the run of \mathcal{A} , but \mathcal{A}_p reads c in between, although c contributes only later to the run of \mathcal{A} . To correctly check the existence of a run of \mathcal{A} on γ , \mathcal{A}_p uses for each buffer a distinguished "pebble" placed on a state of \mathcal{A} . Every time \mathcal{A}_p reads an unmatched send $i!v$, it moves the i -th pebble along an v transition of \mathcal{A} . So each pebble checks for a piece of the accepting run of the whole word coding the final configuration. \mathcal{A}_p therefore also needs to make sure that all of these pieces of runs can be concatenated to form a run of \mathcal{A} . Therefore, at the beginning, \mathcal{A}_p guesses an initial control state l_i and a final control state l'_i for each pebble, and ensures that $(l'_i, \#, l_{i+1}) \in \delta_{\mathcal{A}}$. While doing so, going back to our example, \mathcal{A}_p ensures that $ad\#c$ can be read by \mathcal{A} . It remains also to deal with the control state: indeed, \mathcal{A} should accept $\mathbf{I}\#ad\#c$. So \mathcal{A}_p guesses before reading e that the control state will be \mathbf{I} after executing e , and while reading e , it computes the current control state of \mathfrak{G} . In the end, it checks that this control state actually is \mathbf{I} .

Now that we presented the intuitions about \mathcal{A}_p , let us define it formally. Let us start with the set of control states. Let $L_p = L_{\mathfrak{G}} \times L_{\mathfrak{G}} \times L_{\mathcal{A}}^{|I|} \times L_{\mathcal{A}}^{|I|}$. Intuitively, the control state $(\mathbf{I}_{\mathfrak{G}}, \mathbf{I}_F, \mathbf{I}_{\mathcal{A}}, \mathbf{I}_I)$ of \mathcal{A}_p corresponds to a situation where: (1) the current control state of \mathfrak{G} is $\mathbf{I}_{\mathfrak{G}}$, (2) the guessed final control state of \mathfrak{G} is \mathbf{I}_F , (3) the i -th pebble currently is on state $l_{\mathcal{A},i}$ of \mathcal{A} , and (4) \mathbf{I}_I is a copy of the initial positions of the pebbles and will be checked against their final positions in the end to ensure that all pieces of runs can be concatenated.

Let us now define the set $L_{p,0}$ of initial control states of \mathcal{A}_p . Let us set that $(\mathbf{I}_{\mathfrak{G}}, \mathbf{I}_F, \mathbf{I}_{\mathcal{A}}, \mathbf{I}_I \in L_{p,0}$ if (1) $\mathbf{I}_{\mathcal{A}} = \mathbf{I}_I$, (2) $l_{\mathcal{A},1} \in \delta_{\mathcal{A}}^*(l_{\mathcal{A},0}, \mathbf{I}_F \cdot \#)$, and (3) $\mathbf{I}_{\mathfrak{G}} = \mathbf{I}_{\mathfrak{G},0}$. Intuitively, a control state is initial if (1) \mathbf{I}_I is a copy of $\mathbf{I}_{\mathcal{A}}$, (2) the position of the pebble of buffer 1 is on a state that is reachable in \mathcal{A} after reading $\mathbf{I}_F\#$ and (3) $\mathbf{I}_{\mathfrak{G}}$ is the initial control state of \mathfrak{G} .

Similarly, let us now define the set F_P of final control states of \mathcal{A}_P . Let us set $(\mathbf{l}_G, \mathbf{l}_F, \mathbf{l}_A, \mathbf{l}_I) \in F_P$ if (1) $\mathbf{l}_G = \mathbf{l}_F$, (2) for all $i = 1, \dots, |I| - 1$, $(l_{A,i}, \sharp, l_{A,i+1}) \in \delta_A$, and (3) $l_{A,|I|} \in F_A$. Intuitively, a control state is final if (1) the current control state of \mathfrak{S} corresponds to the guessed final one, (2) the i -th pebble can be moved along a \sharp transition so as to reach the initial position of pebble $i + 1$, and (3) the pebble of the last buffer reached an accepting state of \mathcal{A} .

Let us now define the set δ_P of transitions of \mathcal{A}_P . Let us set that

$$\left((\mathbf{l}_G, \mathbf{l}_F, \mathbf{l}_A, \mathbf{l}_I), c, (\mathbf{l}'_G, \mathbf{l}'_F, \mathbf{l}'_A, \mathbf{l}'_I) \right) \in \delta_P$$

if (1) $\mathbf{l}_F = \mathbf{l}'_F$, (2) $\mathbf{l}_I = \mathbf{l}'_I$, (3) $\mathbf{l}'_G \in \delta_G^*(\mathbf{l}_G, c)$, (4.1) if $c = i!v$, then $\mathbf{l}_A = \mathbf{l}'_A$, and (4.2) if $c = i!v$, then $(l_{A,i}, v, l'_{A,i}) \in \delta_A$ and $l_{A,j} = l'_{A,j}$ for all $j \neq i$. Intuitively, when it reads a matched send $i!v$, \mathcal{A}_P only updates \mathbf{l}_G according to the sequence of actions $i!v$, while when it reads an unmatched send $i!v$ it also updates the position of the i -th token.

Now that \mathcal{A}_P is defined, observe that $RS(\mathfrak{S}) \cap P(\mathfrak{S}) = \emptyset$ iff $\mathcal{L}(\mathcal{A}_P) \cap \mathcal{L}(\mathcal{A}_{gr}) = \emptyset$, where \mathcal{A}_{gr} is the automaton defined in Lemma 9. The emptiness of this intersection is decidable in time $\mathcal{O}(|\mathcal{A}_P| \cdot |\mathcal{A}_{gr}|)$, which shows the claim. \square

4.2 Examples of Regular Safety Problems

In this section we review a few properties of systems that are polynomial-time computable regular properties and showcase some applications of Theorem 12.

Reachability. The control state reachability problem is to decide, given a system \mathfrak{S} and a control state $\mathbf{l} \in L_G$, whether there exists $\mathbf{b} \in (\mathbb{V}^*)^I$ and $e \in \text{Act}^*$ such that $\gamma_0 \xrightarrow[e]{\mathfrak{S}} (\mathbf{l}, \mathbf{b})$. The configuration reachability problem, on the other hand, is to decide, given a system \mathfrak{S} and a configuration γ , whether $\gamma \in RS(\mathfrak{S})$. Both problems are safety problems for a regular property P computable in polynomial time (and even constant time): $P(\mathfrak{S}) = \mathbf{l} \cdot (\sharp \cdot \mathbb{V}^*)^{|I|}$ for the control state reachability problem, and $P(\mathfrak{S}) = \{\gamma\}$ for the configuration reachability problem.

Unspecified reception. Unspecified receptions is one of the errors that session types usually forbid. This error makes more sense for mailbox systems, so let us assume for now that \mathfrak{S} is a mailbox system. A configuration is an *unspecified reception* if one of the participants is in a receiving state, and none of its outgoing transitions can receive the first message in its buffer.

Let us define these notions more formally. A control state l_p of process p is a receiving state if for all a, l' such that $(l_p, a, l') \in \delta_p$, a is a receive action. The set $\{v \mid (l_p, p?v, l') \in \delta_p \text{ for some } l'\}$ is called the ready set of l_p .

A configuration (\mathbf{l}, \mathbf{b}) is said an *unspecified reception configuration* if there is $p \in \mathbb{P}$ such that (1) l_p is a receiving state, (2) $b_p = vb'$ for some $v \in \mathbb{V}$ and $b' \in \mathbb{V}^*$, and (3) v is not in the ready set of l_p . It can be observed that the set $UR(\mathfrak{S})$ of unspecified receptions of \mathfrak{S} defines a regular property that is computable in polynomial time.

Progress. Another property that is central in session types is progress. A global control state \mathbf{l} of \mathfrak{S} is *final* if there is no action a and global control state \mathbf{l}' such that $(\mathbf{l}, a, \mathbf{l}') \in \delta_{\mathfrak{S}}$. A configuration $\gamma = (\mathbf{l}, \mathbf{b})$ of \mathfrak{S} *satisfies progress* if either \mathbf{l} is final or there is a configuration γ' and an action a such that $\gamma \xrightarrow[a]{\mathfrak{S}} \gamma'$. A system satisfies progress if all reachable configurations satisfy progress. It can be observed that the set $NP(\mathfrak{S})$ of configurations that do not satisfy progress is regular and polynomial time computable.

4.3 Boundedness

There are other examples of properties that are regular and polynomial time, but some interesting ones are not safety properties. To conclude this section, we consider one of these properties.

Definition 13 (Boundedness). *Let \mathfrak{S} be a system and $k \geq 0$. A channel $i \in I$ is k -bounded if for all $(\mathbf{l}, \mathbf{b}) \in RS(\mathfrak{S})$ $|b_i| \leq k$. \mathfrak{S} is k -bounded if for all $i \in I$, i is k -bounded.*

Theorem 14. *Whether there exists $k \geq 0$ such that a greedy system \mathfrak{S} is k -bounded is decidable in polynomial time. Moreover, k is computable in polynomial time.*

Proof. Let $\Sigma_! \subseteq \Sigma$ be the set of unmatched communications, and $\sigma : \Sigma^* \rightarrow \Sigma_!^*$ the morphism that erase all matched communications. By Lemma 9, and by the closure of regular languages under morphisms, there is a polynomial time computable non-deterministic finite state automaton \mathcal{A} such that $\mathcal{L}(\mathcal{A}) = \{\sigma(e) \mid e \in \text{executions}(\mathfrak{S}) \text{ is greedy}\}$. Then \mathfrak{S} is k -bounded for some k if and only if $\mathcal{L}(\mathcal{A})$ is finite, or equivalently if and only if \mathcal{A} , once pruned (removing states that are not reachable from the initial state and co-reachable from a final state), is acyclic. This is decidable in time $\mathcal{O}(|\mathcal{A}|)$, and the maximal length k of a word of $\mathcal{L}(\mathcal{A})$ also is computable in time $\mathcal{O}(|\mathcal{A}|)$. \square

5 Mailbox Half-Duplex Systems

Binary half-duplex systems, (called simply half-duplex by Cece and Finkel [5]) are binary systems such that all reachable configurations (l_1, l_2, b_1, b_2) are such that either $b_1 = \varepsilon$ or $b_2 = \varepsilon$. In the previous sections, we established that greedy systems enjoy the same decidability and complexity results as binary half-duplex systems. In this section, we defend the claim that greedy systems could also merit the name of multiparty half-duplex systems.

First, observe that binary half-duplex systems are greedy (see [5, Lemma 20]). The converse does not hold in general: some binary greedy systems are not half-duplex. However, under an extra hypothesis, both are equivalent. A system is called *without orphan messages* if for all reachable configuration containing a message in a buffer, it is possible to reach a configuration where this message has been received. This property is also enforced by session types and is very natural in communicating systems. Then observe that, for a given a binary system \mathfrak{S} without orphan messages, the following two are equivalent: (1) \mathfrak{S} is binary half-duplex, and (2) \mathfrak{S} is greedy.

Let us now consider multiparty systems. Cece and Finkel proposed two notions of multiparty half-duplex systems, but conclude that they were not well behaved (one being too restrictive, and the other Turing powerful). Both of these generalisations relied on peer-to-peer communication. We propose to consider mailbox communication instead.

Definition 15 (Half-duplex execution). *Fix a mailbox system \mathfrak{S} . An execution $e = (\mathbf{l}_0, \mathbf{b}_0) \xrightarrow{a_1} (\mathbf{l}_1, \mathbf{b}_1) \rightarrow \dots \xrightarrow{a_n} (\mathbf{l}_n, \mathbf{b}_n)$ is half-duplex if for all $j = 1, \dots, n$, if a_j is a send action, then $b_p^{i-1} = \varepsilon$, where $p = \text{process}(a_j)$.*

Intuitively, an execution is half-duplex if every process empties its queue of messages before sending.

Definition 16 (Mailbox half-duplex system). *A mailbox system \mathfrak{S} is mailbox half-duplex if for all execution $e \in \text{executions}(\mathfrak{S})$, there is a half-duplex execution e' such that $e \stackrel{\mathfrak{S}}{\sim} e'$.*

Example 7. The system in Figure 1 is half-duplex. Indeed even if execution e of Example 2 is not half-duplex, by considering one of its greedy equivalents e'' :

$$e'' = s!req \cdot s?req \cdot c!res \cdot c?res \cdot s!ack_s \cdot s?ack_s \cdot d!log_c \cdot d?log_c \cdot c!ack_d \cdot c?ack_d \cdot d!log_s \cdot d?log_s \cdot s!req$$

we obtain a half-duplex execution. Notice that execution e' of Example 4 is greedy but not half-duplex as the send of message log_s is done when the buffer of the Server is filled with message req . \square


Note that a binary system is mailbox half-duplex if and only if it is binary half-duplex. In the remainder, we therefore sometimes say simply half-duplex instead of binary half-duplex or mailbox half-duplex.

Theorem 17. *Mailbox half-duplex systems are greedy.*

Proof. We reason by contradiction. Assume that \mathfrak{S} is not greedy, we show that \mathfrak{S} is not half-duplex. Let $e \in \text{executions}(\mathfrak{S})$ be any execution that is not causally equivalent to a greedy execution (for instance, take for e a borderline violation). We claim that for all e' such that $e \stackrel{\mathfrak{S}}{\sim} e'$, e' is not half-duplex. Since e' is not causally equivalent to a greedy execution, we get, by Lemma 6, that $\text{cgraph}(e')$ contains a cycle of communications $c_1 \rightarrow_{e'} c_2 \rightarrow_{e'} \dots \rightarrow_{e'} c_n \rightarrow_{e'} c_1$ where for all $i = 1, \dots, n$, either $c_i = \{j_i, k_i\}$ is a matching pair, or $c_i = \{j_i\}$ is an unmatched send. We assume that $j_i < k_i$, i.e. j_i is the index of the send action and k_i the index of the receive action. Up to a circular permutation, we can also assume, without loss of generality, that j_1 is the first send among them in e' , i.e. $j_1 < j_\ell$ for all $\ell = 2, \dots, n$. Now, let us reason by case analysis on the nature of the conflict edge $c_n \rightarrow_{e'} c_1$.

- case “ $c_n \xrightarrow{SS} c_1$ ”: $j_n \prec_{e'} j_1$. Then $j_n < j_1$, contradicts the minimality of j_1 . Impossible.
- case “ $c_n \xrightarrow{RS} c_1$ ”: $k_n \prec_{e'} k_1$. Then $j_n < k_n < j_1$, impossible.
- case “ $c_n \xrightarrow{RR} c_1$ ”: $k_n \prec_{e'} k_1$. Then $k_n < k_1$ and either (1) $\text{process}(a_{k_n}) = \text{process}(a_{k_1})$ or (2) there is $i \in I, v, v' \in \mathbb{V}$ such that $a_{k_n} = i?v$ and $a_{k_1} = i?v'$. Because of the mailbox semantics, (1) and (2) are equivalent, so (2) is granted. But then $a_{j_n} = i!v$ and $a_{j_1} = i!v'$. Since e' is a FIFO execution, and $k_n < k_1$, we get that $j_n < j_1$, and again the contradiction.
- case “ $c_n \xrightarrow{SR} c_1$ ”: $j_n \prec_{e'} k_1$. Then $j_n < k_1$, and $\text{process}(a_{j_n}) = \text{process}(a_{k_1})$. Moreover, $j_1 < j_n$ by the minimality of j_1 . To sum up, let p, q, r, v_1, v_2 be such that $a_{j_1} = !v_1^{p \rightarrow q}$, $a_{k_1} = ?v_1^{p \rightarrow q}$, and $a_{j_n} = !v_2^{q \rightarrow r}$. Then we just showed that $e' = \dots !v_1^{p \rightarrow q} \dots !v_2^{q \rightarrow r} \dots ?v_1^{p \rightarrow q} \dots$, so e' is not a half-duplex execution. \square

Notice that the converse of Theorem 17 does not hold: being greedy is not a sufficient condition to be half-duplex. Indeed, an unmatched send can fill the buffer of a process willing to send. More precisely, consider the action graph on the right. It depicts a greedy system that is not half-duplex: in fact the buffer for q is not empty when v_2 is sent. We conjecture that this is the only pathological situation, and that, like in the binary setting, if a system is greedy and has no orphan messages, then it is half-duplex.



Example 8. To finish this section we present another small example of a half-duplex system: the classic Client/Seller/Bank protocol. This system is shown in Figure 4. In this protocol a client can ask the price of an item to the seller (ask_price message), and receive the answer. Whenever the client agrees on a price it can place an order (via the message buy). Receiving this message the seller initiates a transaction with the bank. The bank asks the client for its credentials, and when it receives them it either authorizes or refuses the transaction, and notifies the seller accordingly. The seller then confirms or cancels the transaction, sending a message to the client. \square

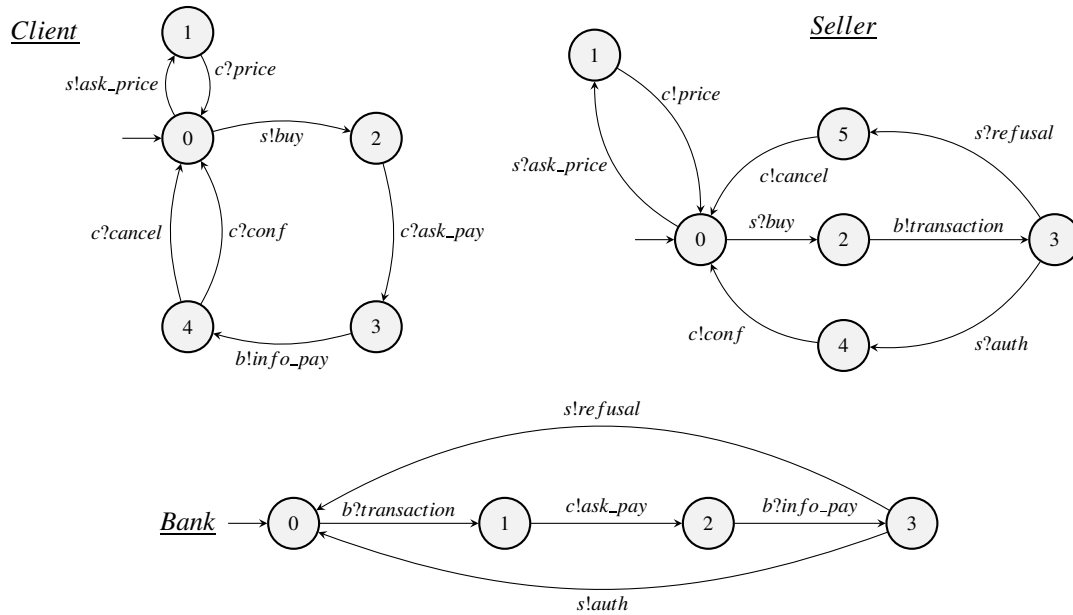


Figure 4: Client/Seller/Bank

6 Conclusion

We have introduced greedy systems, a new class of communicating systems, generalising the notion of half-duplex systems to any number of processes, and to an arbitrary model of FIFO communication (encompassing both p2p and mailbox communications). We have shown that the greediness of a system is decidable in polynomial time, and that for greedy systems regular safety properties, such as reachability, progress, and boundedness are decidable in polynomial time. Finally, we defined mailbox half-duplex systems and showed that greedy systems are intimately related to mailbox half-duplex systems.

Still, the picture is a bit incomplete: we did not address liveness properties nor more general temporal properties, and we also did not propose a notion of p2p half-duplex systems that would enjoy all desirable properties. Also, we did not report on experimental evaluation. We leave these questions for future work.

Pachl gave a general decidability result for systems of communicating finite state machines whose reachability set is regular [21]. Although the algorithm is rather brute-force, the MCSCM tool illustrates that it is amenable to an efficient CEGAR optimization [13]. Beyond regularity, and relying on visibly pushdown languages, La Torre *et al* [26] established that bounded context-switch reachability is also decidable in time exponential in the number of states and doubly exponential in the number of switch. We conjecture that bounded context-switch reachability is complete for greedy systems with a number of switch polynomial in the size of the system.

Several authors considered communicating systems where queues are not plain FIFO but may over-approximate a FIFO behaviour. A representative example of this approach is lossy channels, introduced independently in [6] and [1]. Another example is the bag semantics of buffers where messages are received without loss but out of order. Examples of uses of bag buffers can be found in [25]. Both for lossy and bag systems the reachability problem is decidable but with a high complexity: non primitive recursive for lossy [23, 24]. The exact complexity seems unknown for bag systems, but by reduction to Petri nets it is non-elementary [7].

Aside bounded context-switch model-checking, another form of bounded model-checking has been promoted by Muscholl *et al.*: the notions of universally and existentially bounded message sequence charts [19]. This leads to two notions of universally/existentially bounded systems (both having unfortunately the same name), depending whether only the MSCs leading to final stable configurations are considered [12] or all MSCs [18]. For the latter definition, reachability and membership are decidable in PSPACE. Greedy systems are existentially 1-bounded, but the converse does not hold (for instance, the system of Example 5) is existentially 1-bounded). A system is k -synchronous [4] if the message sequence charts of the system are formed of blocks of k messages. In particular, a system is 1-synchronous if the message lines never cross. For systems with p2p communications, greediness is the same as 1-synchronizability. However, for mailbox communications, some subtle examples are 1-synchronous but not greedy (see *e.g.* [9, Example 1.2]). For k -synchronous systems, reachability is decidable in PSPACE. Finally, greedy systems are synchronizable in the sense of Basu and Bultan [3], but synchronizability is not decidable [11]. We believe that greediness is the notion that Basu and Bultan were aiming at with the notion of synchronizability. It might be wondered if greedy systems were not implicit in Cece and Finkel's work. Actually, some of their arguments rely on the fact that for half-duplex systems, every execution is 'reachability equivalent' to a synchronous execution. This is not exactly the notion of greedy systems we introduced, and our notion of greedy systems is closer to Bouajjani et al notion of 1-synchronous systems, although, as we just explained, they are not the same in some corner cases.

Session types are intimately related to half-duplex systems in the binary setting [20]. Several multi-party extensions of session types have been proposed, the last proposal being [22]. It seems there are also some similarities between multiparty session types and greedy systems. For instance, the notion of greedy execution shares some similarities with the notion of alternation in [8]. The study of the exact relationship between greedy systems and multi-party session typed systems is left for future work.

Acknowledgements. We would like to thank all the ICE reviewers for their comments that greatly improved the present paper.

References

- [1] Parosh Aziz Abdulla & Bengt Jonsson (1996): *Verifying Programs with Unreliable Channels*. *Inf. Comput.* 127(2), pp. 91–101, doi:10.1006/inco.1996.0053.
- [2] Lakhdar Akroun & Gwen Salaün (2018): *Automated verification of automata communicating via FIFO and bag buffers*. *Formal Methods Syst. Des.* 52(3), pp. 260–276, doi:10.1007/s10703-017-0285-8.
- [3] Samik Basu & Tevfik Bultan (2016): *On deciding synchronizability for asynchronously communicating systems*. *Theor. Comput. Sci.* 656, pp. 60–75, doi:10.1016/j.tcs.2016.09.023.
- [4] Ahmed Bouajjani, Constantin Enea, Kailiang Ji & Shaz Qadeer (2018): *On the Completeness of Verifying Message Passing Programs Under Bounded Asynchrony*. In Hana Chockler & Georg Weissenbacher, editors: *Computer Aided Verification - 30th International Conference, CAV 2018, Held as Part of the Federated Logic Conference, FloC 2018, Oxford, UK, July 14-17, 2018, Proceedings, Part II, Lecture Notes in Computer Science* 10982, Springer, pp. 372–391, doi:10.1007/978-3-319-96142-2_23.
- [5] Gérard Cécé & Alain Finkel (2005): *Verification of programs with half-duplex communication*. *Inf. Comput.* 202(2), pp. 166–190, doi:10.1016/j.ic.2005.05.006.
- [6] Gérard Cécé, Alain Finkel & S. Purushothaman Iyer (1996): *Unreliable Channels are Easier to Verify Than Perfect Channels*. *Inf. Comput.* 124(1), pp. 20–31, doi:10.1006/inco.1996.0003.
- [7] Wojciech Czerwinski, Slawomir Lasota, Ranko Lazic, Jérôme Leroux & Filip Mazowiecki (2021): *The Reachability Problem for Petri Nets Is Not Elementary*. *J. ACM* 68(1), pp. 7:1–7:28, doi:10.1145/3422822.

- [8] Pierre-Malo Deniérou & Nobuko Yoshida (2012): *Multiparty Session Types Meet Communicating Automata*. In Helmut Seidl, editor: *Programming Languages and Systems - 21st European Symposium on Programming, ESOP 2012, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2012, Tallinn, Estonia, March 24 - April 1, 2012. Proceedings, Lecture Notes in Computer Science 7211*, Springer, pp. 194–213, doi:10.1007/978-3-642-28869-2_10.
- [9] Cinzia Di Giusto, Laetitia Laversa & Étienne Lozes (2020): *On the k-synchronizability of Systems*. In: *FOSSACS 2020, LNCS 12077*, pp. 157–176, doi:10.1007/978-3-030-45231-5_9.
- [10] Manuel Fähndrich, Mark Aiken, Chris Hawblitzel, Orion Hodson, Galen C. Hunt, James R. Larus & Steven Levi (2006): *Language support for fast and reliable message-based communication in singularity OS*. In Yolande Berbers & Willy Zwaenepoel, editors: *Proceedings of the 2006 EuroSys Conference, Leuven, Belgium, April 18-21, 2006*, ACM, pp. 177–190, doi:10.1145/1217935.1217953.
- [11] Alain Finkel & Étienne Lozes (2017): *Synchronizability of Communicating Finite State Machines is not Decidable*. In Ioannis Chatzigiannakis, Piotr Indyk, Anca Muscholl & Fabian Kuhn, editors: *Proceedings of the 44th International Colloquium on Automata, Languages and Programming (ICALP'17), Leibniz International Proceedings in Informatics 80*, Leibniz-Zentrum für Informatik, Warsaw, Poland, pp. 122:1–122:14, doi:10.4230/LIPIcs.ICALP.2017.122.
- [12] Blaise Genest, Dietrich Kuske & Anca Muscholl (2007): *On Communicating Automata with Bounded Channels*. *Fundam. Inform.* 80(1-3), pp. 147–167. Available at <http://content.iospress.com/articles/fundamenta-informaticae/fi80-1-3-09>.
- [13] Alexander Heußner, Tristan Le Gall & Grégoire Sutre (2012): *McScM: A General Framework for the Verification of Communicating Machines*. In Cormac Flanagan & Barbara König, editors: *Tools and Algorithms for the Construction and Analysis of Systems - 18th International Conference, TACAS 2012, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2012, Tallinn, Estonia, March 24 - April 1, 2012. Proceedings, Lecture Notes in Computer Science 7214*, Springer, pp. 478–484, doi:10.1007/978-3-642-28756-5_34.
- [14] Kohei Honda (1993): *Types for Dyadic Interaction*. In Eike Best, editor: *CONCUR '93, 4th International Conference on Concurrency Theory, Hildesheim, Germany, August 23-26, 1993, Proceedings, Lecture Notes in Computer Science 715*, Springer, pp. 509–523, doi:10.1007/3-540-57208-2_35.
- [15] Kohei Honda, Vasco Thudichum Vasconcelos & Makoto Kubo (1998): *Language Primitives and Type Discipline for Structured Communication-Based Programming*. In Chris Hankin, editor: *Programming Languages and Systems - ESOP'98, 7th European Symposium on Programming, Held as Part of the European Joint Conferences on the Theory and Practice of Software, ETAPS'98, Lisbon, Portugal, March 28 - April 4, 1998, Proceedings, Lecture Notes in Computer Science 1381*, Springer, pp. 122–138, doi:10.1007/BFb0053567.
- [16] Kohei Honda, Nobuko Yoshida & Marco Carbone (2016): *Multiparty Asynchronous Session Types*. *J. ACM* 63(1), pp. 9:1–9:67, doi:10.1145/2827695.
- [17] *ITU-TS: ITU-TS Recommendation Z.120: Message Sequence Chart (MSC). ITU-TS (1999)*.
- [18] Dietrich Kuske & Anca Muscholl (2019): *Communicating Automata*. In J.-E. Pin, editor: *Handbook of Automata: from Mathematics to Applications (AutoMathA)*, European Science Foundation. Available at <https://eiche.theoinf.tu-ilmeneau.de/kuske/Submitted/cfm-final.pdf>.
- [19] Markus Lohrey & Anca Muscholl (2004): *Bounded MSC communication*. *Inf. Comput.* 189(2), pp. 160–181, doi:10.1016/j.ic.2003.10.002.
- [20] Étienne Lozes & Jules Villard (2011): *Reliable Contracts for Unreliable Half-Duplex Communications*. In: *WS-FM 2011, LNCS 7176*, pp. 2–16, doi:10.1007/978-3-642-29834-9_2.
- [21] Jan K. Pahl (2003): *Reachability problems for communicating finite state machines*. CoRR cs.LO/0306121. Available at <http://arxiv.org/abs/cs/0306121>.
- [22] Alceste Scalas & Nobuko Yoshida (2019): *Less is more: multiparty session types revisited*. *Proc. ACM Program. Lang.* 3(POPL), pp. 30:1–30:29, doi:10.1145/3290343.

- [23] Philippe Schnoebelen (2002): *Verifying lossy channel systems has nonprimitive recursive complexity*. *Inf. Process. Lett.* 83(5), pp. 251–261, doi:10.1016/S0020-0190(01)00337-4.
- [24] Philippe Schnoebelen (2021): *On Flat Lossy Channel Machines*. In Christel Baier & Jean Goubault-Larrecq, editors: *29th EACSL Annual Conference on Computer Science Logic, CSL 2021, January 25-28, 2021, Ljubljana, Slovenia (Virtual Conference)*, LIPIcs 183, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 37:1–37:22, doi:10.4230/LIPIcs.CSL.2021.37.
- [25] Koushik Sen & Mahesh Viswanathan (2006): *Model Checking Multithreaded Programs with Asynchronous Atomic Methods*. In Thomas Ball & Robert B. Jones, editors: *Computer Aided Verification*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 300–314, doi:10.1007/11817963_29.
- [26] Salvatore La Torre, P. Madhusudan & Gennaro Parlato (2008): *Context-Bounded Analysis of Concurrent Queue Systems*. In C. R. Ramakrishnan & Jakob Rehof, editors: *Tools and Algorithms for the Construction and Analysis of Systems, 14th International Conference, TACAS 2008, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2008, Budapest, Hungary, March 29-April 6, 2008. Proceedings, Lecture Notes in Computer Science* 4963, Springer, pp. 299–314, doi:10.1007/978-3-540-78800-3_21.