

SOUNDNESS IN NEGOTIATIONS*

JAVIER ESPARZA ^a, DENIS KUPERBERG ^b, ANCA MUSCHOLL ^c, AND IGOR WALUKIEWICZ ^c

^a Technical University of Munich
e-mail address: esparza@in.tum.de

^b École Normale Supérieure de Lyon
e-mail address: denis.kuperberg@ens-lyon.fr

^c LaBRI, University of Bordeaux
e-mail address: {anca,igw}@labri.fr

ABSTRACT. Negotiations are a formalism for describing multiparty distributed cooperation. Alternatively, they can be seen as a model of concurrency with synchronized choice as communication primitive.

Well-designed negotiations must be sound, meaning that, whatever its current state, the negotiation can still be completed. In earlier work, Esparza and Desel have shown that deciding soundness of a negotiation is PSPACE-complete, and in PTIME if the negotiation is deterministic. They have also extended their polynomial soundness algorithm to an intermediate class of acyclic, non-deterministic negotiations. However, they did not analyze the runtime of the extended algorithm, and also left open the complexity of the soundness problem for the intermediate class.

In the first part of this paper we revisit the soundness problem for deterministic negotiations, and show that it is NLOGSPACE-complete, improving on the earlier algorithm, which requires linear space.

In the second part we answer the question left open by Esparza and Desel. We prove that the soundness problem can be solved in polynomial time for acyclic, weakly non-deterministic negotiations, a more general class than the one considered by them.

In the third and final part, we show that the techniques developed in the first two parts of the paper can be applied to analysis problems other than soundness, including the problem of detecting race conditions, and several classical static analysis problems. More specifically, we show that, while these problems are intractable for arbitrary acyclic deterministic negotiations, they become tractable in the sound case. So soundness is not only a desirable behavioral property in itself, but also helps to analyze other properties.

Key words and phrases: Negotiations, workflows, soundness, verification, concurrency.

* A preliminary version of this paper appeared in [6].

This work was carried out while Anca Muscholl was on leave at the Institute of Advanced Studies of the Technical University of Munich, supported by a Hans-Fischer Fellowship, and Denis Kuperberg had a postdoctoral position at the same institute.

The work was also supported by the DFG Project “Negotiations: A Model for Tractable Concurrency”.

1. INTRODUCTION

A multiparty atomic negotiation is an event in which several processes (agents) synchronize in order to select one out of a number of possible results. In [3] Esparza and Desel introduced *negotiation diagrams*, or just *negotiations*, a model of concurrency with multiparty atomic negotiation as interaction primitive. A negotiation diagram describes a workflow of “atomic” negotiations. After an atomic negotiation concludes with the selection of a result, the workflow determines the set of atomic negotiations each agent is ready to participate in next.

Negotiation diagrams are closely related to workflow Petri nets, a very successful formalism for the description of business processes, and a back-end for graphical notations like BPMN (Business Process Modeling Notation), EPC (Event-driven Process Chain), or UML Activity Diagrams (see e.g. [17, 20]). In a nutshell, negotiation diagrams are workflow Petri nets that can be decomposed into communicating sequential Petri nets, a feature that makes them more amenable to theoretical study, while the translation into workflow nets (described in [1]) allows to transfer results and algorithms to business process applications.

The most prominent analysis problem for the negotiation model is checking *soundness*, a notion originally introduced for workflow Petri nets. Loosely speaking, a negotiation is sound if from every reachable configuration there is an execution leading to proper termination of the negotiation. In [3] it is shown that the soundness problem is PSPACE-complete for non-deterministic negotiations and CONP-complete for acyclic non-deterministic negotiations. For this reason, and in search of a tractable class, [3] introduces the class of *deterministic negotiations*. In deterministic negotiations all agents are deterministic, meaning that they are never ready to engage in more than one atomic negotiation per result (similarly to a deterministic automaton, that for each action can move to at most one state). In [3] the soundness problem is investigated for acyclic negotiations. The main results are a polynomial time algorithm for checking soundness of deterministic negotiations, and an extension of the algorithm to the more expressive class of weakly deterministic negotiations. However, whether the extended algorithm is polynomial or not was left open. In [4] the polynomial result for acyclic deterministic negotiations is extended to the cyclic case.

In this paper we continue the line of research initiated in [3, 4], and present three contributions.

In the first contribution we revisit the soundness problem for deterministic negotiations. It should be noted that the notion of soundness in [3] has one more requirement (which makes the soundness problem for acyclic negotiations CONP-hard and in DP). We show here that for deterministic, possibly cyclic, negotiations this second requirement is unnecessary, modulo a weak assumption saying that every atomic negotiation is reachable from the initial negotiation by a local path in the graph on the negotiation. We then identify *anti-patterns*, i.e., structures of the graph of a negotiation, and show that a deterministic negotiation is unsound iff it exhibits at least one of them. As an easy consequence of this theorem, we obtain an NLOGSPACE algorithm for checking soundness, whereas the algorithm of [4] requires linear space. Since soundness of deterministic negotiations is easily shown to be NLOGSPACE-hard, our result settles the complexity of the soundness problem for deterministic negotiations.

In the second contribution we answer the question left open in [3]. We prove that the soundness problem can be solved in polynomial time for acyclic, weakly non-deterministic negotiations, a class that is more general than the one considered in [3]¹. The result

¹The weakly deterministic negotiations of [3] are called *very weakly non-deterministic negotiations* in this paper. As the name indicates, every negotiation in this class is also weakly non-deterministic.

is based on a game-theoretic solution to the *omitting problem*, an analysis problem of independent interest. Further, we show that if we leave out one of the two assumptions, acyclicity or weak non-determinism, then the problem becomes coNP -complete². These results set a limit to the class of negotiations with a polynomial soundness problems, but also admit a positive interpretation. Indeed, the soundness problem for arbitrary negotiations is PSPACE -complete [3], and so of higher complexity (under the usual assumption $\text{PTIME} \subset \text{NP} \cup \text{coNP} \subset \text{PSPACE}$).

In the third and final contribution, we show that the techniques developed in the first two parts of the paper, namely anti-patterns and our game-theoretic solution to the omitting problem, can be applied to analysis problems other than soundness. More specifically, we show that, while these problems are intractable for arbitrary deterministic negotiations, they become tractable in the sound case. So soundness is not only a desirable behavioral property in itself, but also helps to analyze other properties. The first problem we consider is the existence of *racers*, i.e., executions in which two given atomic negotiations are concurrently enabled. We show that for acyclic deterministic negotiations the problem is in NLOGSPACE . Then we analyze several classical program analysis problems for negotiations that manipulate data, for example whether every value written into a variable is guaranteed to be read. Such problems have been studied for workflow nets in [16, 13], and exponential algorithms have been proposed. We show that for acyclic sound deterministic negotiations the problems can be solved in polynomial time.

Related formalisms and related work. The connection between negotiations and Petri nets is studied in detail in [1]. The connection is particularly close between deterministic negotiations and free-choice workflow nets. The complexity of the soundness problem for workflow nets has been studied in several papers [18, 19, 11, 15], and in particular in [8] soundness of free-choice workflow nets is also characterized in terms of anti-patterns, which can be used to explain why a given workflow net is unsound. In the conclusions we discuss the consequences of our results for the analysis of soundness in workflow Petri nets in detail.

As a process-based concurrent model, negotiations can be compared with another well-studied model for distributed computation, namely Zielonka automata [22, 2, 12]. Such an automaton is a parallel composition of finite transition systems with synchronization on common actions. The important point is that a synchronization in Zielonka automata involves exchange of information between states of agents: the result of the synchronization depends on the states of all the components taking part in it. Zielonka automata have the same expressive power as arbitrary, possibly nondeterministic negotiations. Deterministic negotiations correspond to a subclass that does not seem to have been studied yet, and for which verification becomes considerably easier. For example, the question whether some local state occurs in some execution is PSPACE -complete for “sound” Zielonka automata, while it can be answered in polynomial time for sound deterministic negotiations.

A somewhat similar graphical formalism are message sequence charts/graphs, used to describe asynchronous communication. Questions like non-emptiness of intersection are in general undecidable for this model, even assuming that communication buffers are bounded. Subclasses of message sequence graphs with decidable model-checking problem were proposed, but the complexity is PSPACE -complete [9].

²We show that coNP -hardness holds even for a very mild relaxation of acyclicity.

Overview. Section 2 introduces definitions and notations. Section 3 revisits the soundness problem for deterministic negotiations and is new compared to the conference version [6]. Section 4 shows that soundness of acyclic weakly non-deterministic negotiations can be decided in polynomial time; the first part of the section solves the omitting problem, and the second part applies the solution to the soundness problem. Section 5 proves that dropping acyclicity or weak non-determinism makes the soundness problem intractable. Section 6 gives polynomial algorithms for the race problem and the static analysis problems of sound deterministic negotiations. Section 7 presents our conclusions.

2. NEGOTIATIONS

A *negotiation* \mathcal{N} is a tuple $\langle Proc, N, dom, R, \delta \rangle$, where $Proc$ is a finite set of *processes* (or agents) that can participate in negotiations, and N is a finite set of *nodes* (or *atomic negotiations*) where the processes can synchronize. The function $dom : N \rightarrow \mathcal{P}(Proc)$ associates to every atomic negotiation $n \in N$ the (non-empty) set $dom(n)$ of processes participating in it (*domain* of n). Negotiations come equipped with two distinguished initial and final atomic negotiations n_{init} and n_{fin} in which *all* processes in $Proc$ participate. Nodes are denoted as m, n, \dots , and processes as p, q, \dots , possibly with indices.

The set of possible results of atomic negotiations is denoted R , and we use a, b, \dots to range over its elements. Every atomic negotiation $n \in N$ has its set of possible results $out(n) \subseteq R$. We assume that every atomic negotiation (except possibly for n_{fin}) has at least one result. (This is a slight change with respect to the definitions of [3, 4], due to the fact that the final result is not relevant for the soundness question.) The control flow in a negotiation is determined by a partial transition function $\delta : N \times R \times Proc \rightarrow \mathcal{P}(N)$, telling that after the result $a \in out(n)$ of an atomic negotiation n , process $p \in dom(n)$ is ready to participate in any of the negotiations from the set $\delta(n, a, p)$. So for every $n' \in \delta(n, a, p)$ we have $p \in dom(n') \cap dom(n)$. For every $n, a \in out(n)$ and $p \in dom(n)$ the result $\delta(n, a, p)$ has to be defined and non-empty. So all processes involved in an atomic negotiation should be ready for all its possible results. Observe that atomic negotiations may have one single participant process, and/or have one single result.

Negotiations admit a graphical representation, cf. Figure 1. A node (atomic negotiation) n is represented as a black bar with a white circles, called *ports*, one for each process in $dom(n)$. So, for example, the negotiation on the left of Figure 1 has six nodes n_0, \dots, n_5 . Nodes n_0, n_4 , and n_5 have two ports each, while nodes n_1, n_2 , and n_3 only have one port. An entry $\delta(n, a, p) = \{n_1, \dots, n_k\}$ is represented by a hyper-arc, labeled by a , that connects the port of process p in n with the ports of process p in n_1, \dots, n_k . In particular, for the negotiation on the left of Figure 1 we have $Proc = \{p_0, p_1\}$, $N = \{n_0, \dots, n_5\}$, $R = \{a, b\}$, and, for example $dom(n_1) = \{p_0\}$, $dom(n_4) = \{p_0, p_1\}$, $\delta(n_4, b, p_0) = \{n_1\}$ and $\delta(n_4, b, p_1) = \{n_2\}$. This negotiation does not contain any proper hyper-arcs, but the second one does; there we have $\delta(n_0, a, p_1) = \{n_2, n_3\}$. In general two nodes may be connected by several arcs, carrying different process/result labels. For instance, node n_4 in the negotiation on the left of Figure 1 has two arcs to n_5 .

Configurations. A *configuration* of a negotiation is a function $C : Proc \rightarrow \mathcal{P}(N)$ mapping each process p to the (non-empty) set of atomic negotiations in which p is ready to engage. The *initial and final configurations* C_{init}, C_{fin} are given by $C_{init}(p) = \{n_{init}\}$ and $C_{fin}(p) = \{n_{fin}\}$ for all $p \in Proc$. An atomic negotiation n is *enabled* in a configuration C if $n \in C(p)$




Figure 1: Three negotiations. In all of them, n_0 is the initial node and the final node is the bottom one.

for every $p \in \text{dom}(n)$, that is, if all processes that participate in n are ready to proceed with it. A configuration is a *deadlock* if no atomic negotiation is enabled in it. If an atomic negotiation n is enabled in C , and a is a result of n , then we say that (n, a) can be executed, and its execution produces a new configuration C' given by $C'(p) = \delta(n, a, p)$ for $p \in \text{dom}(n)$ and $C'(p) = C(p)$ for $p \notin \text{dom}(n)$. We denote this by $C \xrightarrow{(n,a)} C'$. For example, in the negotiation on the right of Figure 1 we have $C \xrightarrow{(n_0,a)} C'$ for $C(p_0) = \{n_0\} = C(p_1)$ and $C'(p_0) = \{n_1\}, C'(p_1) = \{n_2, n_3\}$. Observe that in all three negotiations, the final configuration cannot be executed, since the final node has no result. However, by definition the final node is enabled in the final configuration, and so the final configuration is not a deadlock.³

Runs. A *run* of a negotiation \mathcal{N} from a configuration C_1 is a finite or infinite sequence $w = (n_1, a_1)(n_2, a_2) \dots$ such that there are configurations C_2, C_3, \dots with

$$C_1 \xrightarrow{(n_1, a_1)} C_2 \xrightarrow{(n_2, a_2)} C_3 \dots$$

We denote this by $C_1 \xrightarrow{w}$, or $C_1 \xrightarrow{w} C_k$ if the sequence is finite and finishes with C_k . In the latter case we say that C_k is *reachable from C_1 on w* . We simply call it *reachable* if w is irrelevant, and write $C_1 \xrightarrow{*} C_k$. Consider for example the third negotiation of Figure 1. If we represent a configuration C by the tuple $(C(p_0), C(p_1))$ then

$$(\{n_0\}, \{n_0\}) \xrightarrow{(n_0, a)} (\{n_1\}, \{n_2, n_3\}) \xrightarrow{(n_1, b)} (\{n_2\}, \{n_2, n_3\}) \xrightarrow{(n_2, a)} (\{n_3\}, \{n_3\})$$

is a run.

³While an enabled atomic negotiation that cannot be executed is a bit artificial, we adopt it because it allows to separate deadlock and termination.

A run is called *initial* if it starts in C_{init} . An initial run is *successful* if it starts in C_{init} and ends in C_{fin} . In the three negotiations of Figure 1 we have $n_{init} = n_0$, and $n_{fin} = n_5, n_4, n_3$, respectively. The run shown above is both initial and successful.

Acyclicity. The graph of a negotiation has N , the set of atomic negotiations, as the set of vertices; the edges are $n \xrightarrow{p,a} n'$ if $n' \in \delta(n, a, p)$. Observe that $p \in \text{dom}(n) \cap \text{dom}(n')$.

A negotiation is *acyclic* if its graph is so. For an acyclic negotiation \mathcal{N} we fix a linear order $\preceq_{\mathcal{N}}$ on its nodes that is a topological order on the graph of \mathcal{N} . This means that if there is an edge from m to n in the graph of \mathcal{N} then $m \preceq_{\mathcal{N}} n$. The last two negotiations of Figure 1 are acyclic, while the first one is not. For the negotiation in the middle of the figure there are two options for the topological order, corresponding to fixing $n_2 \preceq_{\mathcal{N}} n_3$ or $n_3 \preceq_{\mathcal{N}} n_2$.

Paths. Fix a negotiation \mathcal{N} . A *local path* of \mathcal{N} is a path $n_0 \xrightarrow{p_0, a_0} n_1 \xrightarrow{p_1, a_1} \dots \xrightarrow{p_{k-1}, a_{k-1}} n_k$ in the graph of \mathcal{N} . A local path is

- a *circuit* if $n_0 = n_k$ and $k \geq 1$;
- a *p-path* if $p_0 = \dots = p_{k-1} = p$;
- *realizable* from a configuration C if there is a run

$$C \xrightarrow{(n_0, a_0)} C'_0 \xrightarrow{w_1} C_1 \xrightarrow{(n_1, a_1)} C'_1 \dots C_{k-1} \xrightarrow{(n_{k-1}, a_{k-1})} C'_{k-1} \xrightarrow{w_k} C_k$$

such that $p_i \notin \text{dom}(w_{i+1})$ for all $i = 0, \dots, k-1$. (Here $\text{dom}(v)$ denotes the set of processes involved in v , that is, $\text{dom}(v) = \bigcup \{ \text{dom}(n) : (n, a) \text{ appears in } v \text{ for some } a \in \text{out}(n) \}$.)

We say that the run *realizes* the path from C .

Example 2.1. For example, $n_0 \xrightarrow{p_1, a} n_2 \xrightarrow{p_1, a} n_4 \xrightarrow{p_0, a} n_5$ is a local path of the first negotiation of Figure 1. The path is realized from the initial configuration by the run

$$\begin{aligned} (\{n_0\}, \{n_0\}) &\xrightarrow{(n_0, a)} (\{n_1\}, \{n_2\}) \xrightarrow{(n_1, a)} (\{n_3\}, \{n_2\}) \\ &\xrightarrow{(n_3, a)} (\{n_4\}, \{n_2\}) \xrightarrow{(n_2, a)} (\{n_4\}, \{n_4\}) \xrightarrow{(n_4, a)} (\{n_5\}, \{n_5\}) \end{aligned}$$

Indeed, we can take $w = (n_0, a) w_1 (n_2, a) w_2 (n_4, a)$, with $w_1 = (n_1, a) (n_3, a)$ and $w_2 = \epsilon$.

Soundness. A negotiation \mathcal{N} is *sound* if every initial run can be completed to a successful run. If a negotiation has no infinite runs (for example, this is the case if the negotiation is acyclic), then it is sound iff it has no reachable deadlock configuration. The three negotiations of Figure 1 are sound. If in the negotiation on the right we change $\delta(n_0, a, p_1)$ from $\{n_2, n_3\}$ to $\{n_3\}$, then the negotiation is no longer sound. Indeed, after the change the negotiation has the run

$$(\{n_0\}, \{n_0\}) \xrightarrow{(n_0, a)} (\{n_1\}, \{n_3\}) \xrightarrow{(n_1, b)} (\{n_2\}, \{n_3\})$$

which leads to a deadlock.

Remark 2.2. Our definition of soundness is slightly different from the one used in [3]. The definition of [3], which follows the definition of soundness for workflow Petri nets introduced in [17], requires an additional property: for every atomic negotiation n there is an initial run that enables n . We use the weaker definition because it leads to cleaner theoretical results,

and because, as we shall see, the two definitions are essentially equivalent for deterministic negotiations (see Remark 3.12).

For deterministic negotiations, being unsound is closely related to deadlocks, as the observation below shows:⁴

Lemma 2.3. *Let \mathcal{N} be a deterministic negotiation such that for every node $n \in N$ that is accessible from n_{init} , and every process $p \in \text{dom}(n)$, there exists some p -path from n to n_{fin} . Then \mathcal{N} is unsound iff some initial run leads to a deadlock configuration.*

Proof. Since the right-to-left implication is obvious, we assume that \mathcal{N} is unsound and we show that a reachable deadlock must exist.

For every atomic negotiation n and $p \in \text{dom}(n)$, let $d(n, p)$ be the length of the shortest p -path leading from n to n_{fin} (or ∞ if the path does not exist), and for every reachable configuration C , define $d(C) = (d(C(p_1), p_1), \dots, d(C(p_n), p_n))$. Since \mathcal{N} is unsound, some initial run $C_{init} \xrightarrow{*} C$ cannot be extended to a successful run. Choose the run so that $d(C)$ is minimal w.r.t. the lexicographic order according to some order on $Proc$. We claim that C is a deadlock. Assume the contrary, so C enables some node n . Let p be the smallest process of $\text{dom}(n)$. By assumption, there is some p -path from n to n_{fin} . Let $a \in \text{out}(n)$ and $n' \in N$ be such that $n' = \delta(n, a, p)$ is the first node on the shortest p -path from n to n_{fin} . In particular, we have $d(n', p) < d(n, p)$. Taking $C \xrightarrow{(n,a)} C'$, we get that $d(C')$ is lexicographically smaller than $d(C)$, contradicting the minimality of C , and so the claim is proved. \square

Determinism. Process p is *deterministic* in a negotiation \mathcal{N} if for every $n \in N$ and every $a \in \text{out}(n)$, the set of possible next negotiations, $\delta(n, a, p)$, is a singleton. A negotiation is *deterministic* if every process $p \in Proc$ is deterministic. Graphically, a negotiation is deterministic if it does not have any proper hyper-arc. The negotiation on the left of Figure 1 is deterministic.

A negotiation is *weakly non-deterministic* if for every $n \in N$ at least one of the processes in $\text{dom}(n)$ is deterministic. A negotiation is *very weakly non-deterministic*⁵ if for every $n \in N$, $a \in \text{out}(n)$, and $p \in \text{dom}(n)$, there is a deterministic process q such that $q \in \text{dom}(n')$ for all $n' \in \delta(n, a, p)$. As the names suggest, a very weakly non-deterministic negotiation is weakly non-deterministic, under the (very weak) assumption that every node, but n_{init} is a target of some transition. Indeed, if a node n in a very weakly non-deterministic negotiation is reachable from some other node then there is a deterministic process in $\text{dom}(n)$. The same is true if $n = n_{init}$, since all processes are in the domain of the initial node. This shows the claim. The intuition behind very weak nondeterminism is that every nondeterminism in the transition function should be resolved by a deterministic process.

The negotiation in the middle of Figure 1 is weakly non-deterministic. Indeed, the processes p_0 and p_2 are deterministic, and every node has p_0 or p_2 (or both) in its domain. However, it is not very weakly non-deterministic. To see this, observe that $\delta(n_0, a, p_1) = \{n_2, n_3\}$, but the intersection $\text{dom}(n_2) \cap \text{dom}(n_3) = \{p_1\}$ does not contain any deterministic process. On the contrary, the negotiation on the right of the figure is very weakly non-deterministic, because the deterministic process p_0 belongs to the domain of all nodes.

⁴We could have used the weaker requirement that every accessible node has some local path, not necessarily p -path, to n_{fin} . However, we use the precise statement of the lemma later.

⁵This class was called *weakly deterministic* in [3].

Weakly non-deterministic negotiations allow to model deterministic negotiations with global resources (see Section 6). The resource (say, a piece of data) can be modeled as an additional process, which participates in the atomic negotiations that use the resource. For example, the negotiation in the middle of Figure 1 models a situation in which processes p_0 and p_2 negotiate in n_1 which of the two will have access to the resource modeled by p_1 . If the result of n_1 is a , then p_0 has access to the resource at node n_2 , and if it is b , then p_2 has access to it at node n_3 .

3. SOUNDNESS OF DETERMINISTIC NEGOTIATIONS

We revisit the soundness problem for deterministic negotiations. We give the first NLOGSPACE algorithm for the problem, in contrast with the polynomial algorithm of [3], which requires linear space. The algorithm is based on analysis of the graph of a negotiation as defined on page 6. More precisely we show a novel characterization of soundness in terms of *anti-patterns* in this graph. The characterization allows not only to check soundness, but also to diagnose why a given negotiation is unsound.

The following lemma shows that every local path of a sound and deterministic negotiation is realizable from some reachable configuration.

Lemma 3.1. *Let π be a local path of a sound deterministic negotiation \mathcal{N} , and let n_0 be the first node of π . Then π is realizable from every reachable configuration that enables n_0 .*

Proof. Let $\pi = n_0 \xrightarrow{p_0, a_0} n_1 \xrightarrow{p_1, a_1} \dots \xrightarrow{p_{k-1}, a_{k-1}} n_k$, and let C be a reachable configuration such that $C(p) = n_0$ for every $p \in \text{dom}(n_0)$. By induction on i we show that there is a run $C \xrightarrow{*} C_i$ realizing $n_0 \xrightarrow{p_0, a_0} n_1 \xrightarrow{p_1, a_1} \dots \xrightarrow{p_{i-1}, a_{i-1}} n_i$ and such that n_i is enabled in C_i .

For $i = 0$, we simply take $C_i = C$. For the induction step we assume the existence of C_i in which n_i is enabled. Let C'_{i+1} be the result of executing (n_i, a_i) from C_i . Observe that $C'_{i+1}(p_i) = n_{i+1}$ (recall that \mathcal{N} is deterministic). Since \mathcal{N} is sound, and C'_{i+1} is reachable, there is a run from C'_{i+1} to C_{fin} . We set then C_{i+1} to be the first configuration on this run when n_{i+1} is enabled. \square

In particular, Lemma 3.1 states that there is an initial run containing the atomic negotiation m iff there is a local path from n_{init} to m . If $\text{dom}(m) \cap \text{dom}(n) \neq \emptyset$ then the lemma also provides an easy test for deciding the existence of a run containing both m, n : it suffices to check the existence of a local path $n_{init} \xrightarrow{*} m \xrightarrow{*} n$, or with m, n interchanged.

Our algorithm for checking soundness of deterministic negotiations checks for certain patterns in the graph of the negotiation. Since negotiations exhibiting the patterns are unsound, we call them *anti-patterns*. In order to define them we need to introduce *forks* and recall the notion of *dominating node* of a local path introduced in [4].

Definition 3.2. Let $\mathcal{N} = (\text{Proc}, N, \text{dom}, R, \delta)$ be a deterministic negotiation. A tuple $(p_1, p_2, n_1, n_2) \in \text{Proc}^2 \times N^2$ is a *fork* of \mathcal{N} if there exists a local path from n_{init} to a node $n \in N$ and a result $a \in \text{out}(n)$ such that

- $p_i \in \text{dom}(n) \cap \text{dom}(n_i)$ for $i = 1, 2$;
- for $i = 1, 2$ there exists a p_i -path π_i leading from $\delta(n, a, p_i)$ to n_i ; and
- π_1 and π_2 are disjoint, i.e., no node appears in both.

Definition 3.3. A node n of a local path π *dominates* π if $\text{dom}(m) \subseteq \text{dom}(n)$ for every node m of π .

Example 3.4. The tuple (p_0, p_1, n_3, n_4) is a fork of the negotiation on the left of Figure 1. We can choose $n = n_0$. The p_0 -path is $n_1 \xrightarrow{p_0, a} n_3$, and the p_1 -path is $n_2 \xrightarrow{p_1, a} n_4$. The tuple (p_0, p_1, n_4, n_5) is not a fork since all p_1 -paths from n_0 to n_5 go through n_4 , but the p_0 -path and the p_1 -path are required to be disjoint. If we change this negotiation by setting $\delta(n_4, b, p_1) = n_5$ then (p_0, p_1, n_4, n_5) becomes a fork with $n = n_4$, result $b \in \text{out}(n)$, $n_1 \xrightarrow{p_0, a} n_3 \xrightarrow{p_0, a} n_4$ as p_0 -path, and the p_1 -path consisting of the single node n_5 .

Consider now the local circuit $n_2 \xrightarrow{p_1, a} n_4 \xrightarrow{p_1, b} n_2$ in the graph of the negotiation on the left of Figure 1. The node n_4 is dominating, since its domain includes all processes; node n_2 is not dominating since $p_0 \notin \text{dom}(n_2)$.

Lemma 3.5 [4, Lemma 2]. *Every reachable local circuit of a sound deterministic negotiation (that is, every local circuit containing a node reachable from n_{init} by a local path) has a dominating node.*⁶

Example 3.6. Lemma 3.5 does not hold for arbitrary sound negotiations. Consider the non-deterministic negotiation of Figure 2. It is easy to see that the negotiation is sound. However, the local circuit $n_1 \xrightarrow{p_1, a} n_2 \xrightarrow{p_1, a} n_1$ has no dominating node, because $\text{dom}(n_1) = \{p_0, p_1\}$ and $\text{dom}(n_2) = \{p_1, p_2\}$.




Figure 2: A local circuit without a dominating node

- Definition 3.7.** (1) An anti-pattern of type \mathcal{B} is a p -path leading from n_{init} to a node n such that no p -path leads from n to n_{fin} .
(2) An anti-pattern of type \mathcal{F} is a fork (p_1, p_2, n_1, n_2) such that $p_2 \in \text{dom}(n_1)$ and $p_1 \in \text{dom}(n_2)$.
(3) An anti-pattern of type \mathcal{C} is a local circuit without a dominating node.

Example 3.8. The last two anti-patterns are illustrated in Figure 3. The tuple (p_0, p_1, n_1, n_2) is a fork of the negotiation on the left satisfying $p_1 \in \text{dom}(n_1)$ and $p_0 \in \text{dom}(n_2)$. The local

⁶In [4] dominating nodes of circuits are called synchronizers.

circuit $n_1 \xrightarrow{p_0, a} n_2 \xrightarrow{p_1, a} n_3 \xrightarrow{p_2, a} n_1$ of the negotiation on the right has no dominating node. Observe that the negotiation has no anti-pattern of type \mathcal{F} .

Another example of anti-pattern of type \mathcal{F} appears in the negotiation on the left of Figure 1 modified by letting $\delta(n_4, b, p_1) = n_5$: the tuple (p_0, p_1, n_4, n_5) is a fork with $p_0 \in \text{dom}(n_5)$ and $p_1 \in \text{dom}(n_4)$.




Figure 3: Anti-patterns

Lemma 3.9. *Any deterministic negotiation containing some anti-pattern is unsound.*

Proof. Assume that a sound, deterministic negotiation \mathcal{N} contains an anti-pattern. If the anti-pattern is of type \mathcal{B} , then since the p -path leading to n is realizable (Lemma 3.1), some reachable configuration C satisfies $C(p) = \{n\}$. Since no p -path leads from n to n_{fn} , we have $C'(p) \neq \{n_{fn}\}$ for every configuration C' reachable from C , and so \mathcal{N} is not sound.

If the anti-pattern is of type \mathcal{C} , then the result follows from Lemma 3.5.

If the anti-pattern is of type \mathcal{F} , then \mathcal{N} contains a fork (p_1, p_2, n_1, n_2) such that $p_2 \in \text{dom}(n_1)$, and $p_1 \in \text{dom}(n_2)$. Let n and a be the node and the result required by the definition of a fork. Since local paths of sound deterministic negotiations are realizable (Lemma 3.1), some reachable configuration C enables n . Let $C \xrightarrow{(n, a)} C'$. By the definition of the anti-pattern, there are disjoint p_1 - and p_2 -paths π_1 and π_2 leading to n_1 and n_2 , respectively. Using again soundness and the fact that π_1 and π_2 are disjoint, we can show by induction on $|\pi_1| + |\pi_2|$ that there is a run $C' \xrightarrow{w} C''$ that realizes both π_1 and π_2 from C' . So we have $C''(p_1) = n_1$ and $C''(p_2) = n_2$, because neither n_1 can be executed before n_2 gets enabled, nor the other way round. Since $p_2 \in \text{dom}(n_1)$ and $p_1 \in \text{dom}(n_2)$, neither n_1 nor n_2 can ever occur from C'' . By the same argument as above, this implies that the initial run leading to C'' cannot be extended to a successful run, and so \mathcal{N} is unsound. \square

Lemma 3.10. *Any unsound deterministic negotiation contains some anti-pattern.*

Proof. Let \mathcal{N} be an unsound deterministic negotiation without anti-patterns of type \mathcal{B} . We prove that \mathcal{N} has some anti-pattern of type \mathcal{F} or \mathcal{C} . Let $Proc = \{p_1, \dots, p_n\}$ be the set of processes of \mathcal{N} .

By Lemma 2.3 we know that \mathcal{N} has some reachable deadlock configuration, let us call it C . This means that there are processes p_0, \dots, p_{k-1} , and distinct nodes $n_0, \dots, n_{k-1}, n_k = n_0$ ($k \geq 2$) such that $C(p_i) = n_i$ and $p_i \in \text{dom}(n_{i+1})$ for every $0 \leq i < k$. (Intuitively, at the configuration C process p_0 waits for p_{k-1} , p_1 waits for p_0 , etc.) A run from C_{init} to C yields a fork $(p_i, p_{i+1}, n_i, n_{i+1})$ for every pair of nodes n_i, n_{i+1} : Indeed, we can choose the pair (n, a) required in the definition of a fork as the last pair in the run satisfying $\{p_i, p_{i+1}\} \subseteq \text{dom}(n)$, and choose the path π (resp. π') as a p_i -path from $\delta(n, a, p_i)$ to n_i (resp. a p_{i+1} -path from $\delta(n, a, p_{i+1})$ to n_{i+1}).

If $k = 2$ then (p_0, p_1, n_0, n_1) is a pattern of type \mathcal{F} . So assume that $k > 2$ and that \mathcal{N} has no pattern of type \mathcal{F} . We prove that \mathcal{N} has a local circuit without a dominating node.

We claim that for every $0 \leq i < k$ there is some p_i -path π_i from n_i to n_{i+1} (setting $n_k = n_0$). Assume the contrary, and let (n, a) be the pair in the definition of the fork $(p_i, p_{i+1}, n_i, n_{i+1})$. Then we can extend the p_i -path π from $\delta(n, a, p_i)$ to n_i with some p_i -path π' from n_i to some node n' with $p_{i+1} \in \text{dom}(n')$, and such that p_{i+1} does not occur in π' except for n' . Such a node n' exists because \mathcal{N} has no anti-pattern of type \mathcal{B} , and the final node n_{fin} contains all processes. The p_i -path $\pi\pi'$ is still disjoint with the p_{i+1} -path leading from $\delta(n, a, p_{i+1})$ to n_{i+1} . Therefore, the fork $(p_i, p_{i+1}, n', n_{i+1})$ is an anti-pattern of type \mathcal{F} , contradicting the hypothesis, and the claim is proven.

Note that $p_{i+1} \notin \text{dom}(n_i)$, since otherwise $(p_i, p_{i+1}, n_i, n_{i+1})$ would be a pattern of type \mathcal{F} . Also, p_{i+1} cannot occur in π_i except for n_{i+1} : otherwise we could take the shortest prefix π of π_i ending in a node $n'_i \neq n_{i+1}$ with $\{p_i, p_{i+1}\} \subseteq \text{dom}(n'_i)$ and find another \mathcal{F} pattern, using the paths $\pi_i\pi$ and π_{i+1} . Let γ be the concatenation of $\pi_0, \pi_1, \dots, \pi_{k-1}$. Then γ is a local circuit of \mathcal{N} , and every p_i belongs to the domain of some node of γ . Since p_{i+1} does not belong to the domain of any node of π_i , except n_{i+1} , no node of γ has all of p_0, \dots, p_{k-1} in its domain. So γ has no dominating node, and therefore \mathcal{N} contains a pattern of type \mathcal{C} . \square

The above two lemmas allow us to prove:

Theorem 3.11. *Soundness of deterministic negotiations is NLOGSPACE-complete.*

Proof. The existence of an anti-pattern of type \mathcal{B} or \mathcal{F} is clearly verifiable in NLOGSPACE.

To non-deterministically check the existence of a path of type \mathcal{C} using logarithmic space, we guess an integer M and a process p . Then, we guess on-the fly a circuit π , such that (1) each node of π has domain of size at most M , (2) p occurs in π , and (3) one node of π has domain of size M and does not contain p . Clearly, this implies that π contains at least $M + 1$ processes, but no dominant node.

We prove NLOGSPACE-hardness by reduction from the reachability problem for directed graphs, which is NLOGSPACE-complete. Given a directed graph G with nodes s and t , where s has no incoming and t no outgoing edges, we reduce the question whether $s \xrightarrow{*} t$ to the soundness of a deterministic negotiation $\mathcal{N}(G)$ with one process p . The atomic negotiations of $\mathcal{N}(G)$ are the vertices of G , with domain $\{p\}$. The initial node is s and the final node is t . If $u \rightarrow v$ is an edge of G , then the atomic negotiation u has a result (u, v) with $\delta(u, (u, v), p) = v$. Moreover, for every vertex u of G except t the atomic negotiation u has

a result *back* with $\delta(u, \text{back}, p) = s$. Clearly, $\mathcal{N}(G)$ is sound iff there is some path $s \xrightarrow{*} t$ in G . \square

Remark 3.12. As announced in Remark 2.2, for deterministic negotiations the notion of soundness used in this paper and the one of [3] essentially coincide. More precisely, we show that the two notions coincide under the very weak assumption that for every atomic negotiation n there is a local path from n_{init} to n . (Atomic negotiations that do not satisfy this condition can be identified and removed in NLOGSPACE, and their removal does not change the behavior of the negotiation.)

Recall that the definition of [3] requires that (a) every run can be extended to a successful run *and* (b) that for every atomic negotiation n some initial run enables n . Now, let \mathcal{N} be a deterministic negotiation that is sound according to the definition of this paper, i.e., satisfies (a), and such that for every atomic negotiation n there is a local path from n_{init} to n . We show that (b) also holds. For $n = n_{fin}$, (b) holds because every run can be extended to a successful run. For $n \neq n_{fin}$, we observe that, by Lemma 3.1, the local path leading from n_{init} to n is realizable. By the definition of realizability, there is a reachable configuration C and a process p such that $C(p) = \{n\}$. Since \mathcal{N} is sound, it has a run leading from C to C_{fin} . Since \mathcal{N} is deterministic and $n \neq n_{fin}$, this run necessarily executes n , and we are done.

4. BEYOND DETERMINISM: TRACTABLE CASES

In this section we investigate how much nondeterminism we can allow, while retaining polynomiality of the soundness problem, a question that was left open in [3, 4]. We prove that soundness of acyclic, weakly non-deterministic negotiations can be decided in polynomial time. In Section 5 we show that the problem becomes intractable for both arbitrary weakly non-deterministic negotiations, and for acyclic non-deterministic negotiations.

The polynomial time algorithm is based on a game-theoretic solution to the *omitting problem* for deterministic negotiations, which is a problem of independent interest. Section 4.1 introduces an extended version of the omitting problem for deterministic negotiations where we not only want to omit some nodes but also to reach some other nodes. We show that for a fixed number of nodes to be reached, the problem can be solved in polynomial time for sound, acyclic and deterministic negotiations. Section 4.2 uses this result to prove that the soundness problem for acyclic weakly non-deterministic negotiations is in PTIME.

4.1. Omitting problem. Let $B \subseteq N$ be a set of nodes of a deterministic negotiation \mathcal{N} . We say that a run $(n_1, a_1)(n_2, a_2) \cdots$ of \mathcal{N} *omits* B if $n_i \notin B$ for all i . Let $P \subseteq N \times R$ be a set of pairs consisting of a node and a result. We say that a run of \mathcal{N} *includes* P *and omits* B if it omits B and contains all the pairs from P .

Definition 4.1. The *omitting problem* consists of deciding, given \mathcal{N} , P , and B , whether there is a successful run of \mathcal{N} including P and omitting B .

Given a constant K , the *K-omitting problem* is the subproblem of the omitting problem in which P has size at most K .

We show that the omitting problem for sound, acyclic, and deterministic negotiations can be reduced to solving a safety game on a finite graph (see e.g. [14] for an introduction to games). As a first step we define a *two-player game* $G(\mathcal{N}, B)$ with players Adam and Eve,

where the goal of Eve is to produce a successful run that omits B . The game arena is a finite graph, with vertex set partitioned between Adam and Eve:

- the positions of Eve are $N \setminus B$, namely the set of all the nodes of the negotiation but those in B ;
- the positions of Adam are $N \times R$, namely the set of pairs consisting of a node and an action

The edges of the game arena are:

- $n \rightarrow (n, a)$, for every $n \in N \setminus B$, $a \in out(n)$;
- $(n, a) \rightarrow \delta(n, a, p)$, for every $n \in N$, $a \in out(n)$, $p \in dom(n)$.

The initial game position is n_{init} . Adam wins a play if it reaches a node in B , Eve wins if the play reaches n_{fin} . Eve has a winning strategy in the game, if all plays (from n_{init}) are won by her, no matter which are the moves of Adam.

The idea behind the game is that Eve chooses results, and builds thus a run of the negotiation. Adam challenges Eve by choosing at each step a process, thus a play is a local path of the negotiation. The winning condition for Eve ensures that the execution is successful. Moreover, the execution avoids B since Eve is not allowed to continue from nodes in B .

Observe that since \mathcal{N} is acyclic, the winning condition for Eve is actually a safety condition: every maximal play avoiding B is winning for Eve. So if Eve can win, then she wins with a positional strategy. A *deterministic positional strategy for Eve* is a function $\sigma : N \rightarrow R$, it indicates that at position n Eve should go to position $(n, \sigma(n))$. Since $G(\mathcal{N}, B)$ is a safety game for Eve, there is a *biggest non-deterministic winning strategy* for Eve, i.e., a strategy of type $\sigma_{max} : N \rightarrow \mathcal{P}(R)$. The strategy σ_{max} is obtained by computing the set W_E of all winning positions for Eve in $G(\mathcal{N}, B)$, and then setting for every $n \in N$:

$$\sigma_{max}(n) = \{a \in out(n) : \text{for all } p \in dom(n), \delta(n, a, p) \in W_E\}.$$

Lemma 4.2. *Let \mathcal{N} be an acyclic negotiation. If \mathcal{N} has a successful run omitting B then Eve has a winning strategy in $G(\mathcal{N}, B)$.*

Proof. Define $\sigma(n) = a$ if (n, a) appears in the run. For other nodes define the strategy arbitrarily. To check that this strategy is winning, it is enough to verify that every play respecting the strategy stays in the nodes appearing in the run. \square

Lemma 4.3. *Let \mathcal{N} be a sound, acyclic, and deterministic negotiation such that Eve has a deterministic winning strategy $\sigma : N \rightarrow R$ in $G(\mathcal{N}, B)$. Consider the set S of nodes that are reachable on some play from n_{init} respecting σ . Then there exists a successful run of \mathcal{N} containing precisely the nodes of S .*

Proof. Recall that for an acyclic \mathcal{N} we can fix a topological order $\preceq_{\mathcal{N}}$. Let n_1, n_2, \dots, n_k be an enumeration of the nodes in $S \subseteq (N \setminus B)$ according to $\preceq_{\mathcal{N}}$. Let $w_i = (n_1, \sigma(n_1)) \cdots (n_i, \sigma(n_i))$. By induction on $i \in \{1, \dots, k\}$ we prove that there is a configuration C_i such that $C_{init} \xrightarrow{w_i} C_i$ is a run of \mathcal{N} . This will show that w_k is a successful run containing precisely the nodes of S .

For $i = 1$, $n_1 = n_{init}$, in C_{init} all processes are ready to do n_1 , so C_1 is the result of performing $(n_1, \sigma(n_1))$.

For the inductive step, we assume that we have a run $C_{init} \xrightarrow{w_i} C_i$, and we want to extend it by $C_i \xrightarrow{(n_{i+1}, \sigma(n_{i+1}))} C_{i+1}$. Consider a play respecting σ and reaching n_{i+1} . The last step in this play is $(n_j, \sigma(n_j)) \rightarrow n_{i+1}$, for some $j \leq i$ and n_j in S . Since \mathcal{N} is deterministic,

we have $\delta(n_j, \sigma(n_j), p) = n_{i+1}$ for some process p . Since $j \leq i$ and $(n_j, \sigma(n_j))$ occurred in w_i (but n_{i+1} has not occurred), we have $C_i(p) = \{n_{i+1}\}$. If we show that $C_i(q) = \{n_{i+1}\}$ for all $q \in \text{dom}(n_{i+1})$ then we obtain that n_{i+1} is enabled in C_i and we get the required C_{i+1} . Suppose by contradiction that $C_i(q) = \{n_l\}$ for some $l \neq i+1$. We must have $l > i+1$, since otherwise n_l already occurred in w_i . By definition of our indexing $n_{i+1} \prec_N n_l$. But then no run from C_i can bring process q to a state where it is ready to participate in negotiation n_{i+1} , and so, since \mathcal{N} is deterministic, we have $C(p) = \{n_{i+1}\}$ for every configuration C reachable from C_i . This contradicts the fact that \mathcal{N} is sound. \square

Lemmas 4.2 and 4.3 show that Eve wins in $G(\mathcal{N}, B)$ iff \mathcal{N} has a successful run omitting B . We apply this result to the omitting problem.

Theorem 4.4. *For every constant K , the K -omitting problem for deterministic, acyclic, and sound negotiations is in PTIME.*

Proof. If for some atomic negotiation m we have $(m, a) \in P$ and $(m, b) \in P$ for $a \neq b$ then the answer is negative as \mathcal{N} is acyclic. So let us suppose that it is not the case. By Lemmas 4.2 and 4.3 our problem is equivalent to determining the existence of a deterministic strategy σ for Eve in the game $G(\mathcal{N}, B)$ such that $\sigma(m) = a$ for all $(m, a) \in P$, and all these (m, a) are reachable on a play respecting σ .

To decide this we calculate σ_{max} , the biggest non-deterministic winning strategy for Eve in $G(\mathcal{N}, B)$. This can be done in PTIME as the size of $G(\mathcal{N}, B)$ is proportional to the size of the negotiation. Strategy σ_{max} defines a graph $G(\sigma_{max})$ whose nodes are atomic negotiations, and edges are (m, a, m') if $(m, a) \in \sigma_{max}$ and $m' = \delta(m, a, p)$ for some process p . The size of this graph is proportional to the size of the negotiation. In this graph we look for a subgraph H such that:

- for every node m in H there is at most one a such that (m, a, m') is an edge of H for some m' ;
- for every $(m, a) \in P$ there is an edge (m, a, m') in H for some m' , and moreover m is reachable from n_{init} in H .

We show that such a graph H exists iff there is a strategy σ with the required properties.

Suppose there is a deterministic winning strategy σ such that $\sigma(m) = a$ for all $(m, a) \in P$, and all these (m, a) are reachable on a play respecting σ . We now define H by putting an edge (m, a, m') in H if $\sigma(m) = a$ and $m' = \delta(m, a, p)$ for some process p . As σ is deterministic and winning, this definition guarantees that H satisfies the first item above. The second item is guaranteed by the reachability property that σ satisfies.

For the other direction, given such a graph H we define a deterministic strategy σ_H . We put $\sigma_H(m) = a$ if (m, a, m') is an edge of H . If m is not a node in H , or has no outgoing edges in H then we put $\sigma_H(m) = b$ for some arbitrary $b \in \sigma_{max}(m)$. It should be clear that σ_H is winning since every play respecting σ_H stays in the winning nodes for Eve. By definition $\sigma_H(m) = a$ for all $(m, a) \in P$, and all these (m, a) are reachable on a play respecting σ_H .

So we have reduced the problem stated in the theorem to finding a subgraph H of $G(\sigma_{max})$ as described above. If there is such a subgraph H then there is one in form of a tree, where the edges leading to leaves are of the form (m, a, m') with $(m, a) \in P$. Moreover, there is such a tree with a few, at most $|P|$, nodes with more than one child. So finding such a tree can be done by guessing the $|P|$ branching nodes and solving $|P| + 1$ reachability problems in $G(\sigma_{max})$. This can be done in PTIME since the size of P is bounded by K . \square

4.2. Soundness of acyclic weakly non-deterministic negotiations. In this section we consider acyclic, weakly non-deterministic negotiations, c.f. page 7. That is, we allow some processes to be non-deterministic, but every atomic negotiation should involve at least one deterministic process. We will often work with a part of a negotiation as defined below.

Definition 4.5. The *restriction* of a negotiation $\mathcal{N} = \langle Proc, N, dom, R, \delta \rangle$ to a subset $Proc' \subseteq Proc$ of its processes is the negotiation $\langle Proc', N', dom', R, \delta' \rangle$ where $N' = \{n \in N : dom(n) \cap Proc' \neq \emptyset\}$, $dom'(n) = dom(n) \cap Proc'$, and $\delta'(n, r, p) = \delta(n, r, p) \cap N'$. The restriction of \mathcal{N} to its deterministic processes is denoted \mathcal{N}_D .

If \mathcal{N} is weakly non-deterministic, every atomic negotiation involves a deterministic process, so $\mathcal{N}_D = \mathcal{N}$ have the same set of nodes; but they do not have the same set of processes unless \mathcal{N} is deterministic.

Recall also that for an acyclic negotiation \mathcal{N} we fixed some linear order $\preceq_{\mathcal{N}}$ on N , that is a topological order of the graph of \mathcal{N} .

We show that deciding soundness for acyclic, weakly non-deterministic negotiations is in PTIME. The proof is divided into three parts. In Section 4.2.1 we prove some preliminary lemmas. In Section 4.2.2 we consider the special case in which the negotiation \mathcal{N} has one single non-deterministic process, and show that \mathcal{N} is sound iff \mathcal{N}_D is sound and Eve wins a certain instance of the omitting problem for \mathcal{N}_D . Finally, Section 4.2.3 shows how to reduce the general case to the case with only one non-deterministic process.

4.2.1. Preliminaries. We first show two auxiliary lemmas on the structure of runs in negotiations. Two runs w, w' of \mathcal{N} are called *equivalent* (and we write $w \equiv w'$) if one can be obtained from the other by repeatedly permuting adjacent pairs $(m, a), (n, b)$, with $dom(m) \cap dom(n) = \emptyset$.

The next lemma shows that for acyclic negotiations we can restrict our considerations to runs respecting the order $\preceq_{\mathcal{N}}$.

Lemma 4.6. *Every run of an acyclic negotiation \mathcal{N} has an equivalent run that respects the topological order $\preceq_{\mathcal{N}}$.*

Proof. Let $C_{init} \xrightarrow{(n_0, a_0) \cdots (n_p, a_p)} C$ be some run of \mathcal{N} . The proof is by induction on the number of pairs $0 \leq i < j \leq p$ such that $n_j \preceq_{\mathcal{N}} n_i$. If there are no such pairs then we are done. Otherwise, assume that a pair i, j as above exists. Note that in particular, $dom(n_i) \cap dom(n_j) = \emptyset$ holds. It is not hard to see that there must exist some $i \leq k < j$ such that $n_{k+1} \preceq_{\mathcal{N}} n_k$. As before, we note that $dom(n_k) \cap dom(n_{k+1}) = \emptyset$. Clearly, $(n_0, a_0) \cdots (n_{k-1}, a_{k-1})(n_{k+1}, a_{k+1})(n_k, a_k)(n_{k+2}, a_{k+2}) \cdots (n_p, a_p)$ is an equivalent run of \mathcal{N} , with one less pair that violates $\preceq_{\mathcal{N}}$. \square

The next lemma states some properties of runs respecting the order $\preceq_{\mathcal{N}}$.

Lemma 4.7. *Let \mathcal{N} be an acyclic, weakly non-deterministic negotiation. Consider some reachable configuration C . Let n be the $\preceq_{\mathcal{N}}$ -smallest atomic negotiation with $C(d) = n$ for some deterministic process d . The following properties hold:*

- (1) *If \mathcal{N} is sound then n is enabled in C .*
- (2) *For all runs $C \xrightarrow{w} C_{fin}$, all atomic negotiations n' in w are $\preceq_{\mathcal{N}}$ -bigger than n .*

Proof. For the first item we use the soundness of \mathcal{N} : if n were not enabled in C , then there would exist some n' that is enabled in C and such that the execution of n' makes

the execution of n eventually possible. So in particular, we would have $n' \prec_{\mathcal{N}} n$, which contradicts the choice of n .

The second item is shown similarly: every atomic negotiation n' occurring in w satisfies $n'' \preceq_{\mathcal{N}} n'$ for some n'' that is enabled in C . By the choice of n we have $n \preceq_{\mathcal{N}} n''$, which shows the claim. \square

It is easy to see that whenever C is a reachable configuration of a weakly non-deterministic negotiation \mathcal{N} , the restriction C^D of C to deterministic processes is a reachable configuration of \mathcal{N}_D . The next lemma considers the reverse construction, lifting runs of \mathcal{N}_D to runs of \mathcal{N} . For this, we need to assume that the runs respect the order $\preceq_{\mathcal{N}}$.

Lemma 4.8. *Suppose \mathcal{N} is a sound, acyclic and weakly non-deterministic negotiation. Let $C_{init}^D \xrightarrow{w} C^D$ be a run of \mathcal{N}_D respecting the order $\preceq_{\mathcal{N}}$. Let also C_v^D denote the configuration reached by some prefix v of w : $C_{init}^D \xrightarrow{v} C_v^D$. Then \mathcal{N} has a run $C_{init} \xrightarrow{v} C_v$ with $C_v(d) = C_v^D(d)$, for every prefix v of w and every deterministic process d .*

Proof. The proof is by induction on the length of v . Consider a prefix $v(m, a)$ of w . Recall that w respects the order $\preceq_{\mathcal{N}}$, thus m is the $\preceq_{\mathcal{N}}$ -smallest atomic negotiation enabled in C_v^D . In particular, Lemma 4.7 (2) applied to \mathcal{N}_D implies that m is the $\preceq_{\mathcal{N}}$ -smallest atomic negotiation with $C_v(d) = m$ for some deterministic process. Applying Lemma 4.7 (1) to \mathcal{N} shows finally that m is also enabled in C_v . \square

The next lemma gives a necessary condition for the soundness of \mathcal{N} that is easy to check. It is proved by showing that \mathcal{N}_D cannot have much more behaviors than \mathcal{N} .

Lemma 4.9. *If \mathcal{N} is a sound, acyclic, weakly non-deterministic negotiation then \mathcal{N}_D is sound.*

Proof. Suppose to the contrary that \mathcal{N}_D has a run reaching a deadlock configuration $C_{init} \xrightarrow{w} C$. By Lemma 4.6 we can assume that w respects the $\preceq_{\mathcal{N}}$ ordering. By Lemma 4.8 we get a run of \mathcal{N} to a deadlock configuration, but this is impossible. \square

4.2.2. *Negotiations with one non-deterministic process.* We consider a special case of a negotiation with only one non-deterministic process. The next lemma establishes the connection to the omitting problem: \mathcal{N} is unsound if the deterministic part \mathcal{N}_D is unsound, or \mathcal{N}_D is sound but has a certain successful omitting run.

Lemma 4.10. *Let \mathcal{N} be an acyclic, weakly non-deterministic negotiation with a single non-deterministic process p . Then \mathcal{N} is not sound if and only if either:*

- \mathcal{N}_D is not sound, or
- \mathcal{N}_D is sound, and it has two nodes $m \preceq_{\mathcal{N}} n$ with results $a \in out(m)$, $b \in out(n)$ such that:
 - $p \in dom(m) \cap dom(n)$, $n \notin \delta(m, a, p)$, and
 - there is a successful run of \mathcal{N}_D containing $P = \{(m, a), (n, b)\}$ and omitting $B = \{n' \in \delta(m, a, p) : m \prec_{\mathcal{N}} n' \prec_{\mathcal{N}} n\}$.

Proof. Consider the right-to-left direction. We abbreviate $S_p := \delta(m, a, p)$. If \mathcal{N}_D is not sound then by Lemma 4.9, \mathcal{N} is not sound.

Suppose then that \mathcal{N}_D satisfies the second item from the statement of the lemma, and take a run w of \mathcal{N}_D as it is assumed there. By Lemma 4.6 we can assume that this run respects $\preceq_{\mathcal{N}}$. Towards a contradiction suppose also that \mathcal{N} is sound. Lemma 4.8 says that w

is also a run of \mathcal{N} . Let C_1 be the configuration of this run just after (m, a) was executed, so we have $C_1(p) = S_p$. Let C_2 be the first configuration after C_1 such that $C_2(d) = n$ for some process d (it may be that $C_2 = C_1$). We have $C_2(p) = S_p$ since the run w omits $\{n' \in S_p : m \prec_{\mathcal{N}_D} n' \prec_{\mathcal{N}_D} n\}$, so p cannot move between C_1 and C_2 . When we continue following w from C_2 we see that either p will never move, or it will move to some $n' \neq n$ with $n \prec_{\mathcal{N}_D} n'$. But then d will not be able to move. So this run leads to a deadlock, contradiction with the soundness of \mathcal{N} .

For the left-to-right direction, assume that \mathcal{N}_D is sound. We need to show the second item of the lemma. Observe that since \mathcal{N} is acyclic and not sound, there is a run $C_{init} \xrightarrow{w} C$ where C is a deadlock. By Lemma 4.6 we can assume that w respects $\preceq_{\mathcal{N}}$. Let n be the $\preceq_{\mathcal{N}}$ -smallest atomic negotiation such that $C(d) = n$ for some deterministic process d . Applying Lemma 4.7 (1) to \mathcal{N}_D yields that n must be deterministically enabled in C : that is $C(d') = n$ for all deterministic processes $d' \in \text{dom}(n)$. This implies $p \in \text{dom}(n)$, and $n \notin C(p)$.

Let us split w as $u(m, a)v$ where $p \in \text{dom}(m)$ and all atomic negotiations in v involving p are $\preceq_{\mathcal{N}}$ -bigger than n (it may be that v is empty). That is to say, we define m as the last atomic negotiation in w involving p that is $\preceq_{\mathcal{N}}$ -smaller than n . Let C_m be the configuration reached after doing (m, a) : $C_{init} \xrightarrow{u(m, a)} C_m$. Take $S_p = C_m(p) = \delta(m, a, p)$. By the choice of m and v , the run from C_m to C does not use any atomic negotiation from the set $\{n' \in S_p : m \prec_{\mathcal{N}} n' \prec_{\mathcal{N}} n\}$.

By soundness of \mathcal{N}_D , from C^D there is a run to the final configuration, and by the choice of n and Lemma 4.7 (2) this run cannot use any atomic negotiation that is $\preceq_{\mathcal{N}}$ -smaller than n . Let b be the result such that (n, b) appears in this run. Putting these pieces together we have a successful run of \mathcal{N}_D containing $\{(m, a), (n, b)\}$ and omitting $\{n' \in S_p : m \prec_{\mathcal{N}} n' \prec_{\mathcal{N}} n\}$. We have already observed that $p \in \text{dom}(m) \cap \text{dom}(n)$ and $n \notin S_p = \delta(m, a, p)$. So all the requirements of the lemma are met. \square

Lemma 4.11. *Soundness of acyclic, weakly non-deterministic negotiations with only one non-deterministic process can be checked in PTIME.*

Proof. For every $m \preceq_{\mathcal{N}} n$, a and b we check the conditions described in Lemma 4.10. The existence of a run of \mathcal{N}_D can be checked in PTIME thanks to Theorem 4.4 and the fact that the size of P is always 2. \square

4.2.3. *General weakly non-deterministic negotiations.* The next lemma deals with the case where there is more than one non-deterministic process. Loosely speaking, in this case \mathcal{N} is unsound iff there is a non-deterministic process such that the restriction of \mathcal{N} to the deterministic processes *and* this process is unsound.

Lemma 4.12. *An acyclic, weakly non-deterministic negotiation \mathcal{N} is unsound iff:*

- (1) *its restriction \mathcal{N}_D to deterministic processes is unsound, or*
- (2) *for some non-deterministic process p , its restriction \mathcal{N}^p to p and the deterministic processes is unsound.*

Proof. For the right-to-left direction the case where \mathcal{N}_D is unsound follows directly from Lemma 4.9. It remains to check the case where \mathcal{N}^p is not sound for some non-deterministic process p . Consider a run $C_{init}^p \xrightarrow{w} C^p$ where C^p is a deadlock. By Lemma 4.6 we can assume that w respects $\preceq_{\mathcal{N}}$. Now let us try to make \mathcal{N} execute the sequence w .

If $C_{init} \xrightarrow{w} C$ is a run of \mathcal{N} then C is a deadlock. Indeed if in \mathcal{N} it would be possible to do $C \xrightarrow{(n,a)}$ for some (n, a) then $C^p \xrightarrow{(n,a)}$ would be possible in \mathcal{N}^p .

The other case is when in \mathcal{N} it is not possible to execute all the sequence w . Then we have $w = v(n, a)v'$, a run $C_{init} \xrightarrow{v} C_1$ and from C_1 action (n, a) is not possible. Since $C_{init}^p \xrightarrow{v} C_1^p \xrightarrow{(n,a)} C_2^p$ is a run of \mathcal{N}^p , we know from Lemma 4.7 (2) that there is a deterministic process d with $C_1(d) = n$, and $n \preceq_{\mathcal{N}} C_1(d')$ for all other deterministic processes d' . Thus C_1 is a deadlock because by Lemma 4.7 (1) if there is an action possible from C_1 then this must be n .

For the left-to-right direction, suppose that \mathcal{N} is not sound and take a run $C_{init} \xrightarrow{w} C$ with C a deadlock configuration. Take the $\preceq_{\mathcal{N}}$ -smallest atomic negotiation n such that $n = C(d)$ for some deterministic process d . Consider an arbitrary non-deterministic process p and a run $C_{init}^p \xrightarrow{w} C^p$ of \mathcal{N}^p ; it is indeed a run since \mathcal{N}^p is a restriction of \mathcal{N} . As \mathcal{N}^p is sound, it is possible to extend this run. By Lemma 4.7 (1), it should be possible to execute n from C^p . Hence for every deterministic process $d \in \text{dom}(n)$, we have $C(d) = n$. Moreover $n \in C(p)$ if $p \in \text{dom}(n)$ is non-deterministic. Since the choice of p was arbitrary, we have $n \in C(p)$ for all $p \in \text{dom}(n)$. Thus it is possible to execute n from C , a contradiction. \square

Theorem 4.13. *Soundness can be decided in PTIME for acyclic, weakly non-deterministic negotiations.*

Proof. By Lemma 4.12 we can restrict to negotiations \mathcal{N} with one non-deterministic process. For every $m \preceq_{\mathcal{N}} n$, a and b we check the conditions described in Lemma 4.10. The existence of a run of \mathcal{N}_D can be checked in PTIME thanks to Theorem 4.4 and the fact that the size of P is 2. \square

5. BEYOND DETERMINISM: INTRACTABLE CASES

We show that if we remove any of the two assumptions of Theorem 4.13 (acyclicity and weak non-determinism) then the soundness problem becomes CONP-complete. In fact, even a very mild relaxation of acyclicity suffices.

It is not very surprising that deciding soundness for acyclic, non-deterministic negotiations is CONP-complete. In the acyclic case the negotiation is unsound iff it can reach a deadlock. Clearly, by acyclicity this would happen after at most $|N|$ steps. So it suffices to guess a run of linear length step by step and check if it leads to a deadlock. CONP-hardness is proved by a simple reduction of SAT to the complement of the soundness problem. The reduction, which strongly relies on non-determinism, is presented in [3]. We get:

Proposition 5.1 ([3]). *Soundness of acyclic non-deterministic negotiations is CONP-complete.*

Now we consider a very mild relaxation of acyclicity: deterministic processes still need to be acyclic, but non-deterministic processes may have cycles.

Recall that \mathcal{N}_D is the restriction of \mathcal{N} to deterministic processes.

Definition 5.2. A negotiation \mathcal{N} is *det-acyclic* if \mathcal{N}_D is acyclic.

It follows easily from this definition that all runs of a weakly non-deterministic, det-acyclic negotiation have length at most $|N|$. However, we show that even in the *very* weakly non-deterministic case (c.f. page 7) the soundness problem is CONP-complete.

Theorem 5.3. *Non-soundness of det-acyclic, very weakly non-deterministic negotiations is NP-complete.*

Proof. We describe a reduction from 3-SAT and fix a 3-CNF formula $\varphi = c_1 \wedge \dots \wedge c_m$, with clauses c_1, \dots, c_m , each of length 3, and k variables x_1, \dots, x_k . Let $c_j = \ell_{j,1} \vee \ell_{j,2} \vee \ell_{j,3}$, with $\ell_{j,d} \in \{x_i, \bar{x}_i : 1 \leq i \leq k\}$ for every $1 \leq j \leq m$, $d \in \{1, 2, 3\}$. We construct a det-acyclic, very weakly non-deterministic negotiation \mathcal{N} such that φ is satisfiable iff \mathcal{N} is not sound.

The atomic negotiations of \mathcal{N} are (apart from n_{init} and n_{fin}):

- An atomic negotiation m_0 , and for each variable x_i three atomic negotiations n_i^+, n_i^-, m_i .
- For every pair clause/literal (j, d) , $j = 1, \dots, m$ and $d = 1, 2, 3$, three atomic negotiations $m_{j,d}$, $n_{j,d}$, and $r_{j,d}$.
- For every clause c_j , two auxiliary atomic negotiations t_j and t'_j .

The processes of \mathcal{N} are:

- A deterministic process E .
- For each clause c_j , a deterministic process V_j .
- For every pair clause/literal (j, d) , where $j = 1, \dots, m$ and $d = 1, 2, 3$: two deterministic processes $T_{j,d}, T'_{j,d}$, and a non-deterministic process $P_{j,d}$.

Now we describe the behavior of each process P by means of a graph. The nodes of the graph for P are the atomic negotiations in which P participates. The graph has an edge $n \rightarrow n'$ if there is a result a of n such that P moves with a from n to n' . If P is nondeterministic, and after a is ready to engage in a set of atomic negotiations $\{n_1, \dots, n_k\}$, then the graph contains a *hyperarc* leading from n to $\{n_1, \dots, n_k\}$.

The graphs of all processes are shown in Figure 4. Intuitively, process E is in charge of producing a valuation of x_1, \dots, x_k : it chooses between n_1^+ and n_1^- , then between n_2^+ and n_2^- , etc. Choosing n_i^+ stands for setting x_i to true, and choosing n_i^- for setting x_i to false.

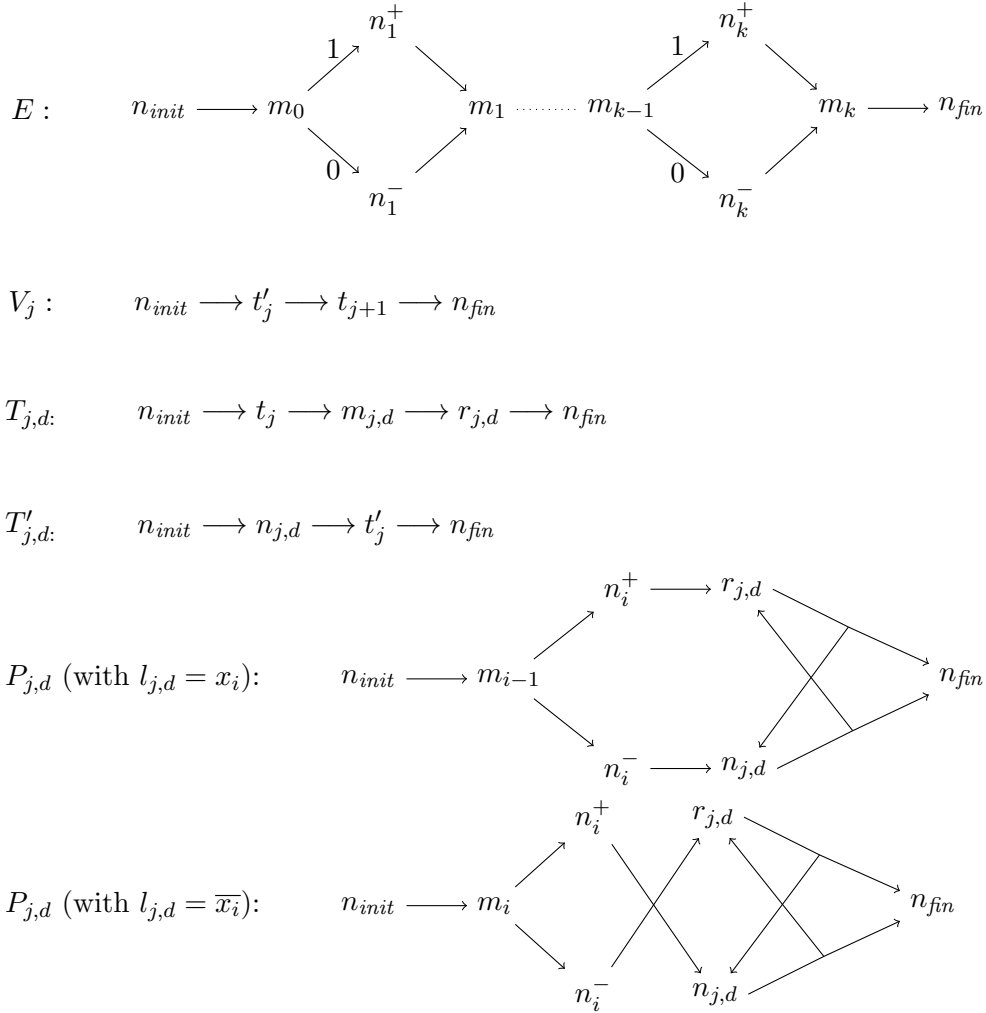
Observe that in the graph for the non-deterministic process $P_{j,d}$ we assume $\ell_{j,d} \in \{x_i, \bar{x}_i\}$. After n_{init} , the process goes to m_i , and then to n_i^+ or n_i^- in a deterministic way, depending on the result chosen at m_i . The rest of its behavior depends on whether $\ell_{j,d} = x_i$ or $\ell_{j,d} = \bar{x}_i$. If $\ell_{j,d} = x_i$, then after n_i^+ the process goes to $r_{j,d}$, and then to one of $\{n_{fin}, n_{j,d}\}$ (nondeterminism!). For the other case, see Figure 4.

Process $P_{j,d}$ is designed with the following purpose. If process E sets literal $\ell_{j,d}$ to true, then $P_{j,d}$ (together with $T_{j,d}$) guarantees that $m_{j,d}$ is executed before $n_{j,d}$; if E sets literal $\ell_{j,d}$ to false, then $m_{j,d}$ and $n_{j,d}$ can occur in any order. In other words, in every successful run containing n_i^+ : if $\ell_{j,d} = x_i$ then node $m_{j,d}$ appears before $n_{j,d}$; if $\ell_{j,d} = \bar{x}_i$ then the nodes $m_{j,d}$ and $n_{j,d}$ can appear in any order. Similarly for runs containing n_i^- , interchanging x_i, \bar{x}_i .

It is easy to see that \mathcal{N} is very weakly non-deterministic. The only nondeterministic processes are the $P_{j,d}$ processes. Moreover, the sets $\{n_{fin}, r_{j,d}\}$ and $\{n_{fin}, n_{j,d}\}$ are the only two sets of atomic negotiations such that (a) there are configurations such that $P_{j,d}$ is ready to engage in them, and (b) contain more than one element. Since n_{fin} contains all processes, the condition for very weak non-determinism is clearly verified.

If the valuation chosen by E makes all c_j true, we claim that the partial run corresponding to this valuation cannot be completed to a successful run. Indeed, in this case for each c_j there is a true literal $\ell_{j,d}$, and so $P_{j,d}$ enforces that $m_{j,d}$ is executed before $n_{j,d}$. We denote this by $m_{j,d} < n_{j,d}$. But then, since process $T_{j,d}$ enforces $t_j < m_{j,d}$, process $T'_{j,d}$ enforces $n_{j,d} < t'_j$, and process V_j enforces $t'_j < t_{j+1 \bmod (m+1)}$, we get the cycle

$$t_1 < t'_1 < t_2 < \dots < t_m < t'_m < t_1$$

Figure 4: Graphs of the processes of \mathcal{N} .

Since, say, t_2 cannot occur before and after t_1 , because the deterministic process are acyclic, the partial run cannot be completed.

Otherwise, if the valuation makes at least one c_j false, then no cycle is created and the partial run can be extended to a successful run.

Here is a more formal version of the proof. By construction, $dom(t_j) = \{V_{j-1}, T_{j,d} \mid d = 1, 2, 3\}$, $dom(t'_j) = \{V_j, T'_{j,d} \mid d = 1, 2, 3\}$, $dom(m_{j,d}) = \{T_{j,d}\}$, $dom(n_{j,d}) = \{T'_{j,d}, P_{j,d}\}$, $dom(r_{j,d}) = \{T_{j,d}, P_{j,d}\}$.

Let $\nu : \{1, \dots, k\} \rightarrow \{0, 1\}$ be a valuation of the variables x_1, \dots, x_k . By C_ν we denote the following configuration (note that only the position of processes $P_{j,d}$ depends on the valuation):

- $C_\nu(E) = n_{fin}$
- If the literal $l_{j,d}$ is true under ν then $C_\nu(P_{j,d}) = r_{j,d}$, and otherwise $C_\nu(P_{j,d}) = n_{j,d}$.
- $C_\nu(T_{j,d}) = t_j$, $C_\nu(T'_{j,d}) = n_{j,d}$
- $C_\nu(V_j) = t'_j$

The following property is easy to check:

Fact 1. Every configuration C_ν is reachable. Moreover, every maximal run has some trace-equivalent prefix that reaches one of the configurations C_ν .

Fact 2. If ν does not satisfy the formula then there is a successful run from C_ν .

Assume that ν does not satisfy clause c_j . By Fact 1, nodes $n_{j,d}$, $d = 1, 2, 3$, are enabled. By executing these nodes we can reach a configuration with $P_{j,d}$ in $\{r_{j,d}, n_{fin}\}$ and $T'_{j,d}$ in t'_j . Now t'_j is executable, and V_j moves to node t_{j+1} . Notice that t_{j+1} is now executable, so $T_{j+1,d}$ can move first to $m_{j+1,d}$, then to $r_{j+1,d}$, $d = 1, 2, 3$ (and V_j goes to n_{fin}). Now either $P_{j+1,d}$ is already in $r_{j+1,d}$ or it can come there after $n_{j+1,d}$ is executed. In the first case nodes $r_{j+1,d}, n_{j+1,d}$ can be executed (in this order), in the second case $n_{j+1,d}, r_{j+1,d}$ can be executed (in this order). In either case we can get to a configuration where processes $T_{j+1,d}$ are in n_{fin} and processes $T'_{j+1,d}$ are in t'_{j+1} . Therefore, t'_{j+1} is also executable. By iterating this argument we obtain a successful run executing $t'_j, t_{j+1}, t'_{j+1}, \dots, t_1, t'_1, \dots, t_j$ in this order.

Fact 3. If ν does satisfy the formula then there is no successful run from C_ν .

Suppose by contradiction that there is a successful run σ from C_ν . By assumption, for each j there is some $d = 1, 2, 3$ such that $C_\nu(P_{j,d}) = r_{j,d}$. Therefore $m_{j,d}$ is necessarily executed before $n_{j,d}$ in σ . By construction this implies that t_j is executed before t'_j in σ . Also by construction, t'_j must be executed before $t_{(j+1) \bmod (m+1)}$, for every j . This means that t_1 should be executed before t'_m , and t'_m before t_1 , a contradiction. \square

6. BEYOND SOUNDNESS

We show that the techniques we have developed for the soundness problem, like forks and the omitting game, can be used to check other functional properties of acyclic deterministic negotiations. Moreover, we show that soundness reduces the complexity of verifying these properties: while checking the property for arbitrary deterministic negotiations—sound or not—is an intractable problem, it is polynomial in the sound case.

In Section 6.1 we study the race problem: Given two atomic negotiations m and n , is there a run in which they occur concurrently? In Section 6.2 we address the static analysis of negotiations with data. We assume that the result of an atomic negotiation corresponds to some operations acting on a set of variables. We study some standard questions of static analysis, for instance whether a variable can be allocated and then never deallocated before the execution ends. These questions can be answered in exponential time by constructing the reachability graph of the negotiation and applying standard model checking algorithms. (This is the approach followed in [17], which studies the questions for workflow Petri nets.) We exhibit polynomial algorithms for the acyclic deterministic case.

Definition 6.1. Let $\mathcal{N} = \langle Proc, N, dom, R, \delta \rangle$ be a negotiation. Two atomic negotiations $m, n \in N$ can be *concurrently enabled*, denoted $m \parallel n$, if $dom(m) \cap dom(n) = \emptyset$ and there is a reachable configuration C of \mathcal{N} where both m and n are enabled.

6.1. Races. For a given pair of atomic negotiations m, n of a deterministic negotiation $\mathcal{N} = \langle Proc, N, dom, R, \delta \rangle$, we want to determine if there is a *race* between m and n , i.e., if there is a reachable configuration that concurrently enables m and n .

This question was answered in [10] for live and safe free-choice nets, where a polynomial fixed point algorithm was given. The algorithm can also be applied to sound deterministic negotiations (cyclic or acyclic), but has cubic complexity in the number of atomic negotiations. We show that in the acyclic case there is a simple anti-pattern characterization of the race pairs m, n , which leads to an algorithm that runs in linear time and logarithmic space.

In the rest of the section we give a syntactic characterization of the pairs m, n such that $m \parallel n$. We proceed in several steps. Lemma 6.2 shows the equivalence of the semantic definition with a more concrete, but still semantic, condition. Proposition 6.3 transforms this condition into the conjunction of a syntactic and a semantic condition. Proposition 6.4 replaces the latter by the existence of a certain fork. The final result, given in Theorem 6.5, just puts the two propositions together.

Recall that two runs $w, w' \in (N \times R)^*$ are equivalent if w' can be obtained from w by repeatedly exchanging adjacent pairs $(m, a)(n, b)$ into $(n, b)(m, a)$ whenever $dom(m) \cap dom(n) = \emptyset$.

Lemma 6.2. *Let \mathcal{N} be an acyclic, deterministic, sound negotiation, and let m, n be two atomic negotiations in \mathcal{N} . Then $m \parallel n$ iff every run w from n_{init} containing both m and n has an equivalent run $w' = w_1 w_2$ such that $w' = C_{init} \xrightarrow{w_1} C \xrightarrow{w_2} C'$ for some configuration C where both m and n are enabled.*

Proof. It suffices to show the implication from left to right. So assume that there exists some reachable configuration C_0 where both m and n are enabled. In particular, $dom(m) \cap dom(n) = \emptyset$. By way of contradiction, let us suppose that there exists some run containing both m and n , but this run cannot be reordered as in the statement of the lemma.

Assuming that m appears before n in this run, we claim that there must be some local path from m to n in \mathcal{N} . To see this, assume the contrary and consider a run of the form $w = w_1(m, a)w_2(n, b)w_3$. The run w defines a partial order (actually a Mazurkiewicz trace) $tr(w)$ with nodes corresponding to positions in w , and edges from (m', c) to (n', d) if $dom(m') \cap dom(n') \neq \emptyset$ and (m', c) precedes (n', d) in w . Since there is no path from m to n in \mathcal{N} , nodes (m, a) and (n, b) are unordered in $tr(w)$. So we can choose a topological order w' of $tr(w)$ of the form $w' = w'_1(m, a)(n, b)w'_2$. This shows the claim.

So let π be a path in \mathcal{N} from m to n , say m, n_1, \dots, n_k, n . Let p be some process such that $n_k \xrightarrow{p, a'} n$ for some result a' . Let us go back to C_0 . Since both m and n are enabled in C_0 , we have a transition $C_0 \xrightarrow{n, b} C_1$, for some $b \in out(n)$. Note that m is still enabled in C_1 , since $dom(m) \cap dom(n) = \emptyset$. So we can apply Lemma 3.1 to C_1 and π (because \mathcal{N} is sound), obtaining a configuration C_2 where $C_2(p) = n$. But since n was executed before C_1 , this violates the acyclicity of \mathcal{N} . \square

The next step is to convert the condition from Lemma 6.2 to a condition on the graph of a negotiation. This condition uses the notion of fork from Definition 3.2.

Proposition 6.3. *Let \mathcal{N} be an acyclic, deterministic, sound negotiation, and let m, n be two atomic negotiations in \mathcal{N} with $dom(m) \cap dom(n) = \emptyset$. Then $m \parallel n$ iff there exists an initial run containing both m, n , and there is neither a local path from m to n nor a local path from n to m .*

Proof. For the left-to-right implication, assume by contradiction that there is some local path π from m to n . Consider some reachable configuration C such that m is enabled in C . By Lemma 3.1 we also find a run $C \xrightarrow{*} C'$ such that n is enabled in C' . But note that the run $C_{init} \xrightarrow{*} C \xrightarrow{*} C'$ cannot be reordered as stated in Lemma 6.2, a contradiction.

For the converse, consider some run w containing both m, n . Since there are no local paths in \mathcal{N} between m, n , we can reorder, as in the proof of Lemma 6.2, the run w into some w' such that we find a configuration C of w' where both m and n are enabled. \square

Proposition 6.4. *Let \mathcal{N} be a sound deterministic negotiation, and let m, n be two atomic negotiations of \mathcal{N} . Then \mathcal{N} has an initial run containing both m and n iff it has a fork (p, q, m, n) for some p, q .*

Proof. Right-to-left implication: the proof is similar to the one of Lemma 3.1, but we need to consider three paths instead of a single one. First we realize the path from n_{init} to the branching node n' (with result a) in the definition of fork, using Lemma 3.1. Suppose that the run from C_{init} to the configuration C that enables n' , contains neither m nor n (otherwise another application of Lemma 3.1 suffices). Let $C \xrightarrow{(n', a)} C_1$. We show by induction on the sum of the lengths of the two local paths of the fork how to construct a run containing both m and n . Let $C_1(p) = m_0, C_1(q) = n_0$ and let

$$m_0 \xrightarrow{p, a_0} \dots \xrightarrow{p, a_{k-1}} m_k = m \quad \text{and} \quad n_0 \xrightarrow{q, b_0} \dots \xrightarrow{q, b_{l-1}} n_l = n.$$

be the local paths of the fork. Since \mathcal{N} is sound, some run leads from C_1 to the final configuration. Since \mathcal{N} is p in C_1 can do only m_0 , and q only n_0 , the run must execute both m_0 and n_0 at some point. Suppose without loss of generality that m_0 is executed first. We have $C_1 \xrightarrow{*} C_2 \xrightarrow{(m_0, a_0)} C_3$ for some C_2, C_3 such that $C_3(p) = m_1$ (since \mathcal{N} is deterministic) and $C_3(q) = n_0$. We can now apply the induction hypothesis to the local paths $m_1 \xrightarrow{*} m$, $n_0 \xrightarrow{*} n$, and we are done.

Left-to-right implication: By assumption \mathcal{N} has a run from n_{init} of the form $w = w_1(m, b)w_2(n, c)$. Choose some $p \in \text{dom}(m), q \in \text{dom}(n)$. Let (n', a) be the rightmost letter of w_1 such that $\{p, q\} \subseteq \text{dom}(n')$ (which exists because $\{p, q\} \subseteq \text{dom}(n_{init})$ by definition), and let $m_0 = \delta(n', a, p), n_0 = \delta(n', a, q)$. Then we can extract from w_1 a p -path leading from m_0 to m , and from w_1w_2 a q -path leading from n_0 to n . By the choice of n' these two paths are disjoint. \square

From Proposition 6.3 and 6.4 we immediately obtain:

Theorem 6.5. *For any acyclic, deterministic, sound negotiation \mathcal{N} we can decide in linear time (resp., in logarithmic space) whether two atomic negotiations m, n of \mathcal{N} satisfy $m \parallel n$. The above problem is NLOGSPACE-complete.*

It is not difficult to show that soundness is essential for this result. Indeed, the race problem is NP-hard for acyclic and deterministic, but not necessarily sound, negotiations. We sketch a reduction from 3CNF-SAT. Given a 3CNF-formula ϕ , we construct a deterministic negotiation \mathcal{N}_ϕ with distinguished atomic negotiations n, m and n_{true} satisfying two properties: (a) \mathcal{N}_ϕ has a run enabling n_{true} iff ϕ is satisfiable, and (b) n and m are concurrently enable iff n_{true} is executed. Clearly, (a) and (b) imply that there is a race between n and m iff ϕ is satisfiable.

\mathcal{N}_ϕ has a process C_j for each clause c_j , and a process $P_{i,j}$ for each clause c_j and variable x_i occurring in c_j (positively or negatively). Besides the initial and final atomic negotiations, the nodes of \mathcal{N}_ϕ are:

- a node $SetVariable_i$ for each variable x_i , with all $P_{i,j}$ as participants and two outcomes **true** and **false**;
- a node $GuessLiteral_j$ for each clause c_j , with C_j as only participant and three outcomes, one for each literal of c_j ;
- a node $TrueLiteral_{i,j}$ for each clause c_j and variable x_i occurring in c_j , with $P_{i,j}$ and C_j as participants and one outcome **true**;
- a node n_{true} , with all C_j as participants, and also one outcome **true**; and
- nodes n and m , with participants C_1 and C_2 , respectively.

The transition function is designed so that \mathcal{N}_ϕ behaves as follows. The initial atomic negotiation concurrently enables all nodes $SetVariable_i$ and $GuessLiteral_j$. At $SetVariable_i$ processes $P_{i,j}$'s jointly select a truth assignment for x_i , and process $P_{i,j}$ gets ready to engage in $TrueLiteral_{i,j}$ if the chosen assignment makes the corresponding literal true. At $GuessLiteral_j$ process C_j guesses a literal of c_j , say with variable i , and gets ready to engage in $TrueLiteral_{i,j}$. If $TrueLiteral_{i,j}$ gets enabled, action **true** is executed and C_j becomes ready to engage in n_{true} . So n_{true} becomes enabled iff the assignment guessed at $SetVariable$ nodes satisfies all clauses and at $GuessLiteral_j$ true literals are guessed correctly. After n_{true} occurs, processes C_1 and C_2 become ready to engage in n and m , respectively, and so n and m become enabled. Therefore, if ϕ is satisfiable, then n and m can be concurrently enabled by selecting a satisfying assignment and correctly guessing true literals. If ϕ is unsatisfiable, then n_{true} can never occur, and so n and m cannot occur either. It is easy to see that the transition function is deterministic.

6.2. Negotiations with data. A negotiation with data is a negotiation over a given set X of variables (over finite domains). Each result $(n, a) \in N \times R$ comes with a set Σ of operations on the shared variables. In our examples this set Σ is composed of $alloc(x)$, $read(x)$, $write(x)$, and $dealloc(x)$.

Formally, a *negotiation with data* is a negotiation with one additional component: $\mathcal{N} = \langle Proc, N, dom, R, \delta, \ell \rangle$ where $\ell: (N \times R) \rightarrow \mathcal{P}(\Sigma \times X)$ maps every result to a (possibly empty) set of data operations on variables from X . We assume that for each $(n, a) \in N \times R$ and for each variable $x \in X$ the label $\ell(n, a)$ contains at most one operation on x , that is, at most one element of $\Sigma \times \{x\}$.

As an example, we enrich the deterministic negotiation on the left of Figure 1 with data (this example is adapted from [16]). The negotiation is shown again in Figure 5. Its results perform operations on two variables x_1 and x_2 . The table on the right of the figure gives for each result and for each operation the (indices of the) variables to which the result applies this operation. Observe that the atomic negotiation n_5 has one result that is not represented graphically.

In [16] some examples of data specifications for workflows are considered.

- (1) *Inconsistent data*: an atomic negotiation reads or writes a variable x while another atomic negotiation is writing, allocating, or deallocating it in parallel. In our example there is a run in which (n_2, a) and (n_3, b) read and write to x_2 in parallel.




Figure 5: A negotiation operating on data.

- (2) *Weakly redundant data*: there is a run in which a variable is written and never read before it is deallocated or the run ends. In the example, there is a run in which x_2 is written by (n_3, b) , and never read again before it is deallocated by (n_5, a) .
- (3) *Never destroyed*: there is an execution in which a variable is allocated and then never deallocated before the execution ends. The example has a run which never takes (n_4, b) . In this run the variable x_1 is never deallocated.

It is easy to give algorithms for these properties that are polynomial *in the size of the reachability graph*. We give the first algorithms that check these properties in polynomial time *in the size of the negotiation*, which can be exponentially smaller than its reachability graph.

For the first property we can directly use the algorithm for the race problem: It suffices to check if the negotiation has two results $(m, a), (n, b)$ such that $m \parallel n$ and there is a variable x such that $\ell(a) \cap \{\text{read}(x), \text{write}(x)\} \neq \emptyset$ and $\ell(b) \cap \{\text{write}(x), \text{alloc}(x), \text{dealloc}(x)\} \neq \emptyset$.

In the rest of the section we present a polynomial algorithm for the following abstract problem, which has the problems (2) and (3) above as special instances.

Given sets $\mathcal{O}_1, \mathcal{O}_2, \mathcal{O} \subseteq N \times R$ we say that the negotiation \mathcal{N} *violates the specification* $(\mathcal{O}_1, \mathcal{O}_2, \mathcal{O})$ if there is a successful run $w = (n_1, a_1) \cdots (n_k, a_k)$ of \mathcal{N} , and indices $0 \leq i < j \leq k$ such that $(n_i, a_i) \in \mathcal{O}_1, (n_j, a_j) \in \mathcal{O}_2$, and $(n_l, a_l) \notin \mathcal{O}$ for all $i < l < j$. In this case we also say that $(\mathcal{O}_1, \mathcal{O}_2, \mathcal{O})$ is violated at $(n_i, a_i), (n_j, a_j)$. Otherwise \mathcal{N} *complies with* $(\mathcal{O}_1, \mathcal{O}_2, \mathcal{O})$.

Example 6.6. In order to simplify notations below, we assume that n_{fin} has some result - if not, we can add a self-loop to n_{fin} .

Observe that a variable x is weakly redundant (specification of type (2)) iff \mathcal{N} violates $(\mathcal{O}_1, \mathcal{O}_2, \mathcal{O})$, where $\mathcal{O}_1 = \{(n, a) \in N \times R : \text{write}(x) \in \ell(n, a)\}$, $\mathcal{O}_2 = \{(n, b) \in N \times R : n = n_{fin} \vee \text{dealloc}(x) \in \ell(n, b)\}$ and $\mathcal{O} = \{(n, c) : \ell(n, c) \cap (\Sigma \times \{x\}) \neq \emptyset\}$.

A variable x is never destroyed (specification of type (3)) iff \mathcal{N} violates $(\mathcal{O}_1, \mathcal{O}_2, \mathcal{O})$, where $\mathcal{O}_1 = \{(n, a) \in N \times R : \text{alloc}(x) \in \ell(n, a)\}$, $\mathcal{O}_2 = \{n_{fin}\}$, $\mathcal{O} = \{(n, c) : n = n_{fin} \vee \ell(n, c) \cap \{\text{alloc}(x), \text{dealloc}(x)\} \neq \emptyset\}$.

For the next proposition it is convenient to use the notation $(m, a) \xrightarrow{\pm} (n, b)$, for $m, n \in N, a \in \text{out}(m), b \in \text{out}(n)$ whenever there is a (non-empty) local path in \mathcal{N} from $\delta(m, a, p)$ to n , for some $p \in \text{dom}(m)$.

Proposition 6.7. *Let \mathcal{N} be an acyclic, deterministic, sound negotiation with data, and $(\mathcal{O}_1, \mathcal{O}_2, \mathcal{O})$ a specification. Let $(m, a) \in \mathcal{O}_1$, $(n, b) \in \mathcal{O}_2$. Then \mathcal{N} violates $(\mathcal{O}_1, \mathcal{O}_2, \mathcal{O})$ at $(m, a), (n, b)$ iff $(n, b) \not\rightarrow (m, a)$ and \mathcal{N} has a successful run containing $P = \{(m, a), (n, b)\}$, and omitting the set $B = \{(m', c) \in \mathcal{O} : (m, a) \rightarrow (m', c) \rightarrow (n, b)\}$.*

Proof. For the right-to-left direction: assume that \mathcal{N} has a run w as claimed. Since $(n, b) \not\rightarrow (m, a)$, w can be assumed to be of the form $w = w_1(m, a)w_2(n, b)w_3$. By reordering w we may suppose that for every (m', c) in w_2 , we have $(m, a) \rightarrow (m', c) \rightarrow (n, b)$. Thus, since w omits B this means that $(m', c) \notin \mathcal{O}$, so the claim follows.

For the left-to-right direction: if \mathcal{N} violates $(\mathcal{O}_1, \mathcal{O}_2, \mathcal{O})$ at $(m, a), (n, b)$ then there is a run $w = (n_1, a_1) \cdots (n_k, a_k)$ with $(n_i, a_i) = (m, a)$, $(n_j, a_j) = (n, b)$ and such that $(n_l, a_l) \notin \mathcal{O}$ for all $i < l < j$. Since $(\{(m', c) : (m, a) \rightarrow (m', c) \rightarrow (n, b)\} \cap \{n_i : 1 \leq i \leq k\}) \subseteq \{n_l : i < l < j\}$, the run w contains $(m, a), (n, b)$ and omits B . \square

Remark 6.8. Note that Proposition 6.7 refers to omitting pairs of atomic negotiation/result, whereas the original omitting problem refers to omitting atomic negotiations. However, it is straightforward to adapt the omitting problem and the corresponding results as to handle pairs.

Putting together Proposition 6.7 and Theorem 4.4 we obtain:

Corollary 6.9. *Given an acyclic, deterministic, sound negotiation with data \mathcal{N} , and a specification $(\mathcal{O}_1, \mathcal{O}_2, \mathcal{O})$, it can be checked in polynomial time whether \mathcal{N} complies with $(\mathcal{O}_1, \mathcal{O}_2, \mathcal{O})$.*

7. CONCLUSIONS

Verification questions for finite-state of concurrent systems are very often PSPACE-hard because of the state explosion problem. One approach to this challenge is to search for restrictions in the communication primitives that permit non-trivial interactions, and yet are algorithmically easier to analyze. We have shown that negotiations are a suitable model for this line of attack. On the one hand, non-deterministic processes can simulate any other communication primitive (up to reasonable equivalences); on the other hand, the definition of non-determinism immediately suggests to investigate the deterministic and weakly non-deterministic classes. Even the deterministic negotiations have enough expressive power for interesting applications in the workflow modeling domain. Indeed, they are equivalent to workflow free-choice nets [1], and acyclicity and free-choiceness are quite common: about 70% of the industrial workflow nets of [21, 7, 5] are free-choice, and about 60% are both acyclic and free-choice (see e.g. the table at the end of [5]). We have shown that a number of verification problems for sound deterministic negotiations can be solved in PTIME or even in NLOGSPACE.

Connection to workflow Petri nets. The connection between negotiations and Petri nets is studied in detail in [1]. We explain the consequences of our results for workflow nets. The following discussion assumes that the reader is familiar with workflow Petri nets and the concept of an S-component of a net.

Let us explain how using the results of this paper we can give a NLOGSPACE-algorithm to check soundness of free-choice workflow nets, assuming that the input is not only the workflow net itself, but also a certain decomposition of the net into S-components.

Before proceeding, we need to examine the definitions of soundness used in this paper and in [1]. There exist many notions of soundness for workflow Petri nets [19]. Soundness as defined in this paper corresponds to weak soundness in [19]. Fortunately, it has been shown in [11] that soundness and weak soundness coincide for free-choice workflow Petri nets, and so we do not need to worry about the differences between the two.

The following discussion refers to results in [1]. There, Desel and Esparza describe a mapping that assigns to every negotiation a so-called in/out net (Proposition 8) satisfying the following properties:

- (1) the in/out net is *decomposable*: it admits a covering by S-components such that any two distinct S-components share the initial and final places, and no other place.
- (2) the in/out net is only linearly larger than the deterministic negotiation, and can be constructed in linear time.

Further, [1] proves the following two results:

- (a) The in/out net of a sound deterministic negotiation is a sound free-choice workflow net (Corollary 4 and Proposition 7).
- (b) Every sound decomposable free-choice workflow net is the in/out net of a sound deterministic negotiation (Proposition 9).

Finally, it follows easily from the proof of Proposition 9 that the sound deterministic negotiation of (b) can be constructed in linear time and logarithmic space from the net *and* its covering by S-components. This shows that the following problem is in NLOGSPACE: Given a decomposable free-choice workflow net W , and a decomposition of W into S-components, is W sound?

Our results on acyclic negotiations also have consequences for workflow nets. It is known that soundness and weak soundness of 1-safe acyclic workflow Petri nets are CONP-complete, and so most likely not solvable in polynomial time [15]. Since the in/out nets of acyclic negotiations are also acyclic, Theorem 4.13 of the present paper identifies a proper superclass of acyclic free-choice workflow nets for which weak soundness remains polynomial: the workflow nets derived from acyclic, weakly nondeterministic negotiations. This class does not coincide with any of the known Petri net classes beyond free-choice, like asymmetric choice nets. We see this as a point in favor of the negotiation model: it can be used to investigate classes of nets beyond the free-choice class for which some important properties can be efficiently decided.

Future work. There are several open problems we are currently working on, or intend to address. We do not know if the soundness problem for acyclic weakly non-deterministic negotiations is PTIME-complete. Also, we conjecture that the polynomial algorithms for acyclic deterministic negotiations of Section 6 can be extended to the cyclic case. More generally, we would like to have a better understanding of which verification problems for sound deterministic negotiations can be solved in PTIME. Since these negotiations are not closed under products with automata, we should not expect to be able to polynomially decide arbitrary safety properties. Finally, the analogous question for sound weakly non-deterministic negotiations and the class NP is also increasingly interesting, due to the important advances in SMT tools.

Acknowledgements. We thank the referees for careful reading of the paper.

REFERENCES

- [1] Jörg Desel and Javier Esparza. Negotiations and petri nets. *T. Petri Nets and Other Models of Concurrency*, 11:203–225, 2016.
- [2] Volker Diekert and Grzegorz Rozenberg, editors. *The Book of Traces*. World Scientific, Singapore, 1995.
- [3] Javier Esparza and Jörg Desel. On negotiation as concurrency primitive. In *CONCUR*, pages 440–454, 2013. Extended version in CoRR abs/1307.2145.
- [4] Javier Esparza and Jörg Desel. On negotiation as concurrency primitive II: Deterministic cyclic negotiations. In *FoSSaCS*, pages 258–273, 2014. Extended version in CoRR abs/1403.4958.
- [5] Javier Esparza and Philipp Hoffmann. Reduction rules for colored workflow nets. In *FASE*, volume 9633 of *LNCS*, pages 342–358. Springer, 2016.
- [6] Javier Esparza, Denis Kuperberg, Anca Muscholl, and Igor Walukiewicz. Soundness in negotiations. In *27th International Conference on Concurrency Theory, CONCUR 2016, August 23–26, 2016, Québec City, Canada*, volume 59 of *LIPICs*, pages 12:1–12:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016.
- [7] Dirk Fahland, Cédric Favre, Barbara Jobstmann, Jana Koehler, Niels Lohmann, Hagen Völzer, and Karsten Wolf. Instantaneous soundness checking of industrial business process models. In *Business Process Management*, pages 278–293. Springer, 2009.
- [8] Cédric Favre, Hagen Völzer, and Peter Müller. Diagnostic information for control-flow analysis of workflow graphs (a.k.a. free-choice workflow nets). In *TACAS 2016*, volume 9636 of *LNCS*, pages 463–479. Springer, 2016.
- [9] Blaise Genest, Dietrich Kuske, and Anca Muscholl. A Kleene theorem and model checking algorithms for existentially bounded communicating automata. *Inf. & Comput.*, 204(6):920–956, 2006.
- [10] Andrei Kovalyov and Javier Esparza. A polynomial algorithm to compute the concurrency relation of free-choice signal transition graphs. In *Workshop on Discrete Event Systems, WODES’96*, pages 1–7. Institute of Electrical Engineers, 1996.
- [11] GuanJun Liu. Some complexity results for the soundness problem of workflow nets. *IEEE Trans. Services Computing*, 7(2):322–328, 2014.
- [12] Anca Muscholl. Automated synthesis of distributed controllers. In *ICALP 2015*, volume 9135 of *LNCS*, pages 11–27. Springer, 2015.
- [13] Natalia Sidorova, Christian Stahl, and Nikola Trcka. Soundness verification for conceptual workflow nets with data: Early detection of errors with the most precision possible. *Inf. Syst.*, 36(7):1026–1043, 2011.
- [14] Wolfgang Thomas. Church’s problem and a tour through automata theory. In *Pillars of Computer Science, Essays Dedicated to Boris (Boaz) Trakhtenbrot on the Occasion of His 85th Birthday*, volume 4800 of *LNCS*, pages 635–655. Springer, 2008.
- [15] Ferucio Laurentiu Tiplea, Corina Bocaneala, and Raluca Chiroasca. On the complexity of deciding soundness of acyclic workflow nets. *IEEE Trans. Systems, Man, and Cybernetics: Systems*, 45(9):1292–1298, 2015.
- [16] Nikola Trcka, Wil M. P. van der Aalst, and Natalia Sidorova. Data-flow anti-patterns: Discovering data-flow errors in workflows. In *Advanced Information Systems Engineering (CAiSE)*, volume 5565 of *LNCS*, pages 425–439. Springer, 2009.
- [17] W. M. P. van der Aalst. The application of Petri nets to workflow management. *J. Circuits, Syst. and Comput.*, 08(01):21–66, 1998.
- [18] Wil M. P. van der Aalst. Workflow verification: Finding control-flow errors using petri-net-based techniques. In Wil M. P. van der Aalst, Jörg Desel, and Andreas Oberweis, editors, *Business Process Management, Models, Techniques, and Empirical Studies*, volume 1806 of *Lecture Notes in Computer Science*, pages 161–183. Springer, 2000.
- [19] Wil M. P. van der Aalst, Kees M. van Hee, Arthur H. M. ter Hofstede, Natalia Sidorova, H. M. W. Verbeek, Marc Voorhoeve, and Moe Thandar Wynn. Soundness of workflow nets: classification, decidability, and analysis. *Formal Asp. Comput.*, 23(3):333–363, 2011.
- [20] Wil M. P. van der Aalst and Kees Max van Hee. *Workflow management: models, methods, and systems*. MIT press, 2004.

- [21] B. van Dongen, M. Jansen-Vullers, HMW. Verbeek, and Wil MP. van der Aalst. Verification of the SAP reference models using EPC reduction, state-space analysis, and invariants. *Computers in Industry*, 58(6):578–601, 2007.
- [22] W. Zielonka. Notes on finite asynchronous automata. *RAIRO–Theoretical Informatics and Applications*, 21:99–135, 1987.