

# Static Analysis of Deterministic Negotiations

Javier Esparza  
Technical University of Munich

Anca Muscholl  
University of Bordeaux, LaBRI

Igor Walukiewicz  
CNRS, LaBRI, University of Bordeaux

**Abstract**—Negotiation diagrams are a model of concurrent computation akin to workflow Petri nets. Deterministic negotiation diagrams, equivalent to the much studied and used free-choice workflow Petri nets, are surprisingly amenable to verification. Soundness (a property close to deadlock-freedom) can be decided in PTIME. Further, other fundamental questions like computing summaries or the expected cost, can also be solved in PTIME for sound deterministic negotiation diagrams, while they are PSPACE-complete in the general case.

In this paper we generalize and explain these results. We extend the classical “meet-over-all-paths” (MOP) formulation of static analysis problems to our concurrent setting, and introduce Mazurkiewicz-invariant analysis problems, which encompass the questions above and new ones. We show that any Mazurkiewicz-invariant analysis problem can be solved in PTIME for sound deterministic negotiations whenever it is in PTIME for sequential flow-graphs—even though the flow-graph of a deterministic negotiation diagram can be exponentially larger than the diagram itself. This gives a common explanation to the low-complexity of all the analysis questions studied so far. Finally, we show that classical gen/kill analyses are also an instance of our framework, and obtain a PTIME algorithm for detecting anti-patterns in free-choice workflow Petri nets.

Our result is based on a novel decomposition theorem, of independent interest, showing that sound deterministic negotiation diagrams can be hierarchically decomposed into (possibly overlapping) smaller sound diagrams.

## 1. Introduction

Concurrent systems are difficult to analyze due to the state explosion problem. Unlike for sequential systems, the flow graph of a concurrent system is often exponential in the size of the system, so that analysis techniques for sequential systems cannot be directly applied. One approach to analyze concurrent systems is to take a general model and design heuristics that work for relevant examples. Another, that we pursue in this paper, is to find a restricted class of concurrent systems and design provably efficient algorithms for particular analysis problems for this class.

In [6] Esparza and Desel introduced *negotiation diagrams*, a model of concurrency closely related to workflow

Petri nets. Workflow nets are 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. [29], [31]). 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 [3]) allows to transfer results and algorithms to business process applications.

A negotiation diagram describes a distributed system with a fixed set of sequential processes. The diagram is composed of “atomic negotiations”, each one involving a (possibly different) subset of processes. An atomic negotiation starts when all its participants are ready to engage in it, and concludes with the selection of one out of a fixed set of possible outcomes; for each participant process, the outcome determines which atomic negotiations the process is willing to engage in at the next step. As workflow Petri nets, negotiations can simulate linearly bounded automata, and so all interesting analysis problems are PSPACE-hard for them.

A negotiation is *deterministic* if for every process the outcome of an atomic negotiation completely determines the next atomic negotiation the process should participate in. As shown in [3], the connection between negotiations diagrams and workflow Petri nets is particularly tight in the deterministic case: Deterministic negotiation diagrams are essentially isomorphic to the class of *free-choice* workflow nets, a class important in practice<sup>1</sup> and extensively studied, see e.g. [8], [9], [10], [14], [17], [18], [30]). The state space of deterministic negotiations/free-choice workflow nets can grow exponentially in their size, and so they are subject to the state explosion problem. However, theoretical research has shown that, remarkably, several fundamental problems can be solved in polynomial time by means of algorithms that avoid direct exploration of the state space (contrary to other techniques, like partial-order reduction, that only reduce the number of states to be explored, and still have exponential worst-case complexity). First, it can be checked in PTIME if a deterministic negotiation diagram is *sound* [7], a variant of deadlock-freedom property [18]<sup>2</sup>. Then, for

1. For example, 70% of the almost 2000 workflow nets from the suite of industrial models studied in [8], [13], [32] are free-choice.

2. About 50% of the free-choice workflow nets from the suite mentioned above are sound.

Partially supported by the DFG Project Negotiations: A Model for Tractable Concurrency, and the Institute of Advance Studies of the Technische Universität München.

sound deterministic negotiation diagrams PTIME algorithms have been proposed for: the *summarization problem* [8], the problem of computing the *expected cost* of a probabilistic free-choice workflow net [9], and the identification of some *anti-patterns* [10].

In this paper we develop a generic approach to the static analysis of sound deterministic negotiation diagrams. It covers all the problems above as particular instances, and new ones, like the computation of the best-case/worst-case execution time. The approach is a generalization to the concurrent setting of the classical lattice-based approach to static analysis of sequential flow-graphs [24]. A flow-graph consists of a set of nodes, modeling program points, and a set of edges, modeling program instructions, like assignments or guards<sup>3</sup>. In the lattice-based approach one (i) defines a lattice  $\mathcal{D}$  of dataflow informations capturing the analysis at hand, (ii) assigns semantic transformers  $\llbracket a \rrbracket: \mathcal{D} \rightarrow \mathcal{D}$  to each action  $a$  of the flow-graph, (iii) assigns to a path  $a_1 \cdots a_n$  of the flow graph the functional composition  $\llbracket a_n \rrbracket \circ \cdots \circ \llbracket a_1 \rrbracket$  of the transformers, and (iv) defines the result of the analysis as the “Merge Over all Paths”, i.e, the join of the transformers of all execution paths, usually called the MOP-solution or just the MOP of the dataflow problem. So performing an analysis amounts to computing the MOP of the flow-graph for the corresponding lattice and transformers.

Katoen *et al.* have recently shown in [22], [5] that in order to adequately deal with quantitative analyses of concurrent systems, like expected costs, one needs a semantics that distinguishes between the inherent nondeterminism of each sequential process, and the nondeterminism introduced by concurrency (the choice of the process that should perform the next step). Following these ideas, we introduce a semantics in which the latter is resolved by an external scheduler, and define the MOP for a given scheduler. The result of a dataflow analysis is then given by the infimum or supremum, depending on the application, of the MOPs for all possible schedulers.

The contributions of the paper are the following:

(1) We present an extension of a static analysis framework to deterministic negotiation diagrams. In particular, we identify the class of *Mazurkiewicz-invariant frameworks* that respect the concurrency relation in negotiations. We prove a theorem showing a first important property of sound deterministic negotiations, namely that the MOP is independent of the scheduler for Mazurkiewicz invariant frameworks. This allows to compute the result of the analysis by fixing a scheduler, and computing the MOP for it. As another motivation for Mazurkiewicz-invariant frameworks we observe that there are static analysis frameworks for which analysis is NP-hard, even for sound deterministic negotiation diagrams.

(2) The main contribution of the paper is a method to compute MOP problems for sound deterministic negotiation

diagrams. The method does not require the computation of the reachable configurations. We prove a novel *decomposition theorem* showing that a deterministic negotiation diagram is composed of smaller subnegotiations involving only a subset of the processes, and that these subnegotiations are themselves sound. This allows us to define a generic PTIME algorithm for computing the MOP for Mazurkiewicz-invariant static analysis frameworks.

(3) Finally, we show that the problems studied in [8], [9], and others, are Mazurkiewicz-invariant. Further, we show that the MOP of an important class of analyses – all four flavors of gen/kill problems, well known in the static analysis community – can be reformulated as invariant frameworks, and computed in PTIME.

*Organization of the paper:* Section 2 introduces the negotiation model and static analysis frameworks. Section 3 proves the decomposition theorem. Section 4 presents the algorithm to compute the MOP of an arbitrary Mazurkiewicz-invariant analysis framework. Section 5 deals with gen/kill analyses. Omitted proofs can be found in the extended version [11].

**Related work.** As we have mentioned, deterministic negotiations are very close to free-choice workflow Petri nets, also called workflow graphs. Algorithms for the analysis of specific properties of these nets have been studied in [8], [9], [10], [14], [17], [18], [28], [30]. We have already described above the relation to these works.

We discuss the connection to work on static analysis for (abstract models of) programming languages. The synchronization-sensitive analysis of concurrent programs has been intensively studied (see e.g [1], [2], [12], [15], [16], [19], [23], [25], [26], [27]). A fundamental fact is that interprocedural synchronization-sensitive analysis is undecidable [25], and intraprocedural synchronization-sensitive analysis has high complexity (ranging from PSPACE-completeness to EXPSpace-completeness, depending on the communication primitive), see e.g. [21]). This is in sharp contrast to the linear complexity of static analysis in the size of the flow graph for sequential programs, and causes work on the subject to roughly split into two research directions. The first one aims at obtaining decidability or low complexity of analyses by restricting the possible synchronization patterns. Many different restrictions have been considered: parbegin-parend constructs [12], [27], generalizations thereof (see e.g. [23]), synchronization by nested locks (see e.g. [15]), and asynchronous programming (see e.g. [19]). The other direction does not restrict the synchronization patterns, at the price of worst-case exponential analysis algorithms (see e.g. [2], [16], where control-flow of parallel programs is modelled by Petri nets, and a notion similar to Mazurkiewicz invariance is also used).

Compared with these papers, the original feature of our work is that we obtain polynomial analysis algorithms without restricting the possible synchronization patterns; instead, deterministic negotiation diagrams restrict the *interplay* between synchronization and choice. This distinction can be best appreciated when we compare these formalisms, but

3. In some papers the roles of nodes and edges are reversed: Nodes are program instructions, and edges are program points. The version with program points as nodes is more convenient for our purposes.

excluding choice. In the programming languages of [12], [15], [19], [23], [27], excluding choice means excluding if-then-else or alternative constructs, while for deterministic negotiations it means considering the special case in which every node has exactly one outcome. Sound deterministic negotiation diagrams can model all synchronization patterns given in terms of Mazurkiewicz traces, which is not the case for the formalisms of [12], [15], [19], [23], [27]. For example, the languages of [12], [27] cannot model a synchronization pattern with three processes  $A, B, C$  in which first  $A$  synchronizes with  $B$ , then  $A$  synchronizes with  $C$ , and finally  $B$  synchronizes with  $C$ . Observe that on the other hand, negotiations are finite state, whereas the other formalisms we have mentioned have non-determinism, recursion, and possibly, thread creation.

## 2. Negotiations

A *negotiation diagram*  $\mathcal{N}$  is a tuple  $\langle Proc, N, dom, R, \delta \rangle$ , where  $Proc$  is a finite set of *processes* (or agents) and  $N$  is a finite set of *nodes* where the processes can synchronize to choose an *outcome*. The function  $dom : N \rightarrow \mathcal{P}(Proc)$  associates to every node  $n \in N$  the (non-empty) set  $dom(n)$  of processes participating in it. Nodes are denoted as  $m$  or  $n$ , and processes as  $p$  or  $q$ ; possibly with indices. The set of possible outcomes<sup>4</sup> of nodes is denoted  $R$ , and we use  $a, b, \dots$  to range over its elements. Every node  $n \in N$  has its set of possible outcomes  $out(n) \subseteq R$ .

The control flow in a negotiation diagram is determined by a partial transition function  $\delta : N \times R \times P \rightarrow \mathcal{P}(N)$ , telling that after the outcome  $a$  of node  $n$ , process  $p \in dom(n)$  is ready to participate in any of the nodes in the set  $\delta(n, a, p)$ . So for every  $n' \in \delta(n, a, p)$  we have  $p \in dom(n') \cap dom(n)$ , and for every  $n, a \in out(n)$  and  $p \in dom(n)$  the result  $\delta(n, a, p)$  is defined. Observe that nodes may have one single participant process, and/or have one single outcome. A *location* is a pair  $(n, a)$  such that  $a \in out(n)$ , and we define its domain as  $dom(n)$ .

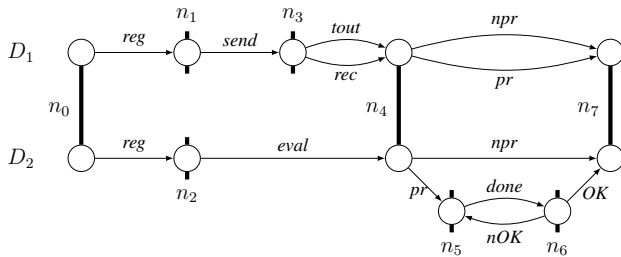


Figure 1. A negotiation diagram with two processes.

**Example:** Figure 1 shows a negotiation diagram for (a slight modification of) the insurance claim example of [29] (see also Fig. 2 of [8] for a workflow Petri net model). The diagram describes a workflow for handling insurance claims by an insurance company with two departments  $D_1$

4.  $R$  stands for *result*; we prefer to avoid the confusing symbol  $O$ .

and  $D_2$ . The processes of the negotiation are  $D_1$  and  $D_2$ . The nodes  $n_0, n_4, n_7$  have domain  $\{D_1, D_2\}$ ;  $n_1$  and  $n_3$  have domain  $D_1$ , and  $n_2, n_5, n_6$  have domain  $D_2$ . After the claim is *registered*, outcome *reg* involves both processes,  $D_1$  *sends* a questionnaire to the client, and concurrently  $D_2$  makes a first *evaluation* of the claim. After the client's answer is *received* or a time-out occurs (outcome *tout*), both departments decide together at node  $n_4$  whether to *process* the claim or not. In both cases  $D_1$  has nothing further to do, and moves to the final node  $n_7$ . If the decision is to process the claim, then  $D_2$  moves to  $n_5$ , and the claim is processed, possibly several times, until a satisfactory result is achieved (outcome *OK*), after which  $D_2$  also moves to  $n_7$ .  $\square$

A *configuration* of a negotiation diagram is a function  $C : Proc \rightarrow \mathcal{P}(N)$  mapping each process  $p$  to the set of nodes in which  $p$  is ready to engage. A node  $n$  is *enabled* in a configuration  $C$  if  $n \in C(p)$  for every  $p \in dom(n)$ , that is, if all processes that participate in  $n$  are ready to proceed with it. A configuration is a *deadlock* if it has no enabled node. If node  $n$  is enabled in  $C$ , and  $a$  is an outcome of  $n$ , then we say that location  $(n, a)$  can be *executed*, and its execution produces a new configuration  $C'$  given by  $C'(p) = \delta(n, a, p)$  for  $p \in dom(n)$  and  $C'(p) = C(p)$  for  $p \notin dom(n)$ . We denote this by  $C \xrightarrow{(n,a)} C'$ . For example, in Figure 1 we have  $C \xrightarrow{(n_0, reg)} C'$  for  $C(D_1) = \{n_0\} = C(D_2)$  and  $C'(D_1) = \{n_1\}, C'(D_2) = \{n_2\}$ .

A *run* of a negotiation diagram  $\mathcal{N}$  from a configuration  $C_1$  is a finite or infinite sequence of locations  $w = (n_1, a_1)(n_2, a_2) \dots$  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  $C_1 \xrightarrow{*} C_k$ .

Negotiation diagrams come equipped with two distinguished *initial* and *final* nodes  $n_{init}$  and  $n_{fin}$ , in which all processes in  $Proc$  participate. 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  $p \in Proc$ . A run is *successful* if it starts in  $C_{init}$  and ends in  $C_{fin}$ . We assume that every node (except for  $n_{fin}$ ) has at least one outcome. In Figure 1,  $n_{init} = n_0$  and  $n_{fin} = n_7$ .

A negotiation diagram  $\mathcal{N}$  is *sound* if every partial run starting at  $C_{init}$  can be completed to a successful run. If a negotiation diagram has no infinite runs, then it is sound iff it has no reachable deadlock configuration. The negotiation diagram of Figure 1 is sound.

Process  $p$  is *deterministic* in a negotiation diagram  $\mathcal{N}$  if for every  $n \in N$ , and  $a \in R$ , the set  $\delta(n, a, p)$  of possible successor nodes is either a singleton or the empty set. A negotiation diagram is *deterministic* if every process  $p \in Proc$  is deterministic. The negotiation diagram of Figure 1 is deterministic.

The *graph of a negotiation diagram* has  $N$  as set of vertices, and there is an edge  $n \xrightarrow{p,a} n'$  iff  $n' \in \delta(n, a, p)$ . Observe that  $p \in \text{dom}(n) \cap \text{dom}(n')$ .

A negotiation diagram is *acyclic* if its graph is acyclic. Acyclic negotiation diagrams cannot have infinite runs, so as mentioned above, soundness is equivalent to deadlock-freedom.

## 2.1. Static analysis frameworks

Let  $(D, \sqcup, \sqcap, \sqsubseteq, \perp, \top)$  be a complete lattice. A function  $f: D \rightarrow D$  is *monotonic* if  $d \sqsubseteq d'$  implies  $f(d) \sqsubseteq f(d')$ , and *distributive* if  $f(\sqcap D') = \sqcap \{f(d) \mid d \in D'\}$ .

An *analysis framework*<sup>5</sup> of a negotiation diagram is a lattice together with a mapping  $\llbracket \_ \rrbracket$  that assigns to each outcome  $\ell$  a monotonic and distributive function  $\llbracket \ell \rrbracket$  in the lattice. Abusing language, we use  $\llbracket \_ \rrbracket$  to denote a framework.

A negotiation diagram has two kinds of nondeterminism, one that picks a node among the ones enabled at a configuration, and a second kind which picks an outcome. We distinguish the two by letting a scheduler to decide the first kind. This is an important design choice, motivated by modeling issues: In distributed systems, one often has information about how the outcomes are picked, but not about the way nondeterminism due to concurrency is resolved. In particular, one may have probabilistic information about the former, but not about the latter. This point has been discussed in detail by Katoen *et al.* [22], [5], who also advocate the separation of the two kinds of nondeterminism.

A *scheduler* of  $\mathcal{N}$  is a partial function  $S$  that assigns to every run  $C_{\text{init}} \xrightarrow{w} C$  a node  $S(w)$  enabled at  $C$ , if it exists. A finite initial run  $w = \ell_1 \cdots \ell_k$ , where  $\ell_i = (n_i, a_i)$ , is *compatible* with  $S$  if  $S(\ell_1 \cdots \ell_i) = n_{i+1}$  for every  $1 \leq i \leq k-1$ .

For example, a scheduler for the negotiation diagram in Figure 1 can give preference to  $n_1$  and  $n_3$  over  $n_2$ . The successful runs compatible with this scheduler are given by the regular expression (omitting the nodes of the locations)  $\text{reg send}(\text{tout}|\text{rec}) \text{eval}(\text{npr}|\text{pr}(\text{done } nOK)^* \text{done } OK)$ .

The abstract semantics of a finite run  $w = \ell_1 \cdots \ell_k$  is the function  $\llbracket w \rrbracket := \llbracket \ell_k \rrbracket \circ \llbracket \ell_{k-1} \rrbracket \circ \cdots \circ \llbracket \ell_1 \rrbracket$ . The abstract semantics of  $\mathcal{N}$  with respect to a scheduler  $S$  is the function  $\llbracket \mathcal{N}, S \rrbracket$  defined by

$$\llbracket \mathcal{N}, S \rrbracket = \bigsqcup \{ \llbracket w \rrbracket \mid w \text{ is a successful run of } \mathcal{N} \text{ compatible with } S \}$$

where the extension of  $\sqcup$  to functions is defined pointwise.

The *abstract semantics*  $\llbracket \mathcal{N} \rrbracket$  of  $\mathcal{N}$  is defined as either  $\bigsqcup \{ \llbracket \mathcal{N}, S \rrbracket \mid S \text{ is a scheduler of } \mathcal{N} \}$ , or as  $\bigsqcup \{ \llbracket \mathcal{N}, S \rrbracket \mid S \text{ is a scheduler of } \mathcal{N} \}$ , depending on the application.

In classical static analysis, analysis frameworks are over flow-graphs, instead of negotiation diagrams [24]. Flow-graphs describe sequential programs. Loosely speaking, a flow-graph is a graph whose nodes are labeled with program

points, and whose edges are labeled with program instructions (assignments or guards). The mapping  $\llbracket \_ \rrbracket$  assigns to an edge the relation describing the effect of the assignment or guard on the program variables. We can see a flow-graph as a degenerate negotiation diagram in which all nodes have one single process. In this case every reachable configuration enables at most one node, and so there is a unique scheduler. So, in this case, the abstract semantics of a flow-graph is the standard ‘‘Merge Over all Paths’’ (MOP), defined by  $\llbracket \mathcal{N} \rrbracket = \bigsqcup \{ \llbracket w \rrbracket \mid w \text{ is a successful run} \}$ .<sup>6</sup>

Several interesting analyses are instances of our framework.

**2.1.1. Input/output semantics.** Let  $V$  be a set of variables and  $Z$  the set of values. A *valuation* is a function  $V \rightarrow Z$ , and  $Val$  denotes the set of all valuations. An element of  $D$  is a set  $T \subseteq Val$ . The join and meet lattice operations are set union and intersection. For each location  $\ell = (n, a)$ , the function  $\llbracket \ell \rrbracket$  describes for each input valuation  $v \in Val$  the set of output valuations  $\llbracket \ell \rrbracket(\{v\})$  that are possible if  $n$  ends with result  $a$ . For any set  $T$  of valuations the function is defined by  $\llbracket \ell \rrbracket(T) = \bigcup_{v \in T} \llbracket \ell \rrbracket(\{v\})$ . The semantics  $\llbracket \mathcal{N} \rrbracket$  is the relation that assigns to every initial valuation the possible final valuations after a successful run.

**2.1.2. Detection of anti-patterns.** Actions in business processes generate, use, modify, and delete resources (for example, a document can be created by a first department, read and used by a second, and classified as confidential by a third). Anti-patterns are used to describe runs that do not correctly access resources; for example, a resource is used before it is created, or a resource is created and then never used. Examples of anti-patterns can be found in [28]. They can be easily formalized as analyses frameworks. Consider for example two locations  $\ell_1$  and  $\ell_2$  that generate a resource, and a set  $K$  of locations that delete it. We wish to know if a given deterministic sound negotiation diagram has a successful run that belongs to

$$L = \mathcal{L}^* \ell_1 (\overline{K})^* \ell_2 \mathcal{L}^*$$

where  $\overline{K}$  denotes the set of locations not in  $K$ . In other words, is there a scenario where a resource is generated twice without deleting it in between.

To encode this problem in our static analysis framework, we take  $D = \{0, 1, 2\}$  with the natural order together with  $\min, \max$  as  $\sqcap$  and  $\sqcup$ , respectively. Intuitively, 0 says that the sequence does not have a suffix of the form  $\ell_1 (\overline{K})^*$ , 1 says that it has such a suffix, and 2 that it has a subword  $\ell_1 (\overline{K})^* \ell_2$ . The semantics of a location is a monotone and distributive function from  $D$  to  $D$  reflecting this intuition:

$$\llbracket \ell_1 \rrbracket(x) = \begin{cases} 2 & \text{if } x = 2 \\ 1 & \text{otherwise} \end{cases} \quad \llbracket \ell_2 \rrbracket(x) = \begin{cases} 0 & \text{if } x = 0 \\ 2 & \text{otherwise} \end{cases}$$

6. Some classical literature uses  $\sqcap$  instead of  $\sqcup$  and speaks of the ‘‘Meet Over all Paths’’, but other standard texts, e.g. [24], use  $\sqcup$ .

5. In [24] this is called a monotone and distributive framework.

$$\llbracket \ell \rrbracket(x) = \begin{cases} x & \text{if } \ell \in \overline{K} \\ 2 & \text{if } \ell \in K \text{ and } x = 2 \\ 0 & \text{if } \ell \in K \text{ and } x = 0, 1 \end{cases}$$

**2.1.3. Minimal/maximal expected cost.** We let  $D = \{(p, c) \mid p \in \mathbb{R}_0^+, c \in \mathbb{R}\}$ , where we interpret  $p$  as a probability and  $c$  as a cost. We take  $(p_1, c_1) \sqcup (p_2, c_2) = (p_1 + p_2, c_1 p_1 + c_2 p_2)$  and  $(p_1, c_1) \sqsubseteq (p_2, c_2)$  if  $p_1 \leq p_2$  and  $c_1 \leq c_2$ .

We define a function  $Prob: N \times R \rightarrow [0, 1]$  such that  $Prob(n, a) = 0$  if  $a \notin out(n)$ , and  $\sum_{a \in R} Prob(n, a) = 1$  for every  $n \in N$ . Intuitively,  $Prob(n, a)$  is the probability that node  $n$  yields the outcome  $a$ . We also define a cost function  $Cost: N \times R \rightarrow \mathbb{R}$  that assigns to each result a cost.

Let  $\llbracket \ell \rrbracket((p, c)) = (p \cdot Prob(\ell), c + Cost(\ell))$ . Then  $\llbracket \mathcal{N}, S \rrbracket(1, 0)$  gives the expected cost of  $\mathcal{N}$  under the scheduler  $S$  (which may be infinite) and  $\llbracket \mathcal{N} \rrbracket(1, 0)$  is the minimal/maximal expected cost.

**2.1.4. Best/worst-case execution time.** Let  $\mathbb{R}_0^+$  denote the nonnegative reals. A *time valuation* is a function  $v: Proc \rightarrow \mathbb{R}_0^+ \cup \{\infty\}$  that assigns to each process  $p$  a time  $v(p)$ , intuitively corresponding to the time that the process has needed so far. The elements of  $D$  are time valuations, with  $(v \sqcup v')(p) = \max\{v(p), v'(p)\}$  for every process  $p$ , and  $v \sqsubseteq v'$  if  $v(p) \leq v'(p)$  for every process  $p$ .

We assign to each outcome  $\ell = (n, a)$  and to each process  $p \in dom(n)$  the time  $t_{\ell, p}$  that  $p$  needs to execute  $a$ . The semantic function  $\llbracket \ell \rrbracket$  is given by  $\llbracket \ell \rrbracket(v) = v'$ , where

$$v'(p) = \begin{cases} v(p) & \text{if } p \notin dom(n) \\ \max_{p' \in dom(n)} v(p') + t_{\ell, p} & \text{if } p \in dom(n) \end{cases}$$

This definition reflects that all processes in  $dom(n)$  must wait until all of them are ready, and then we add to them the time they need to execute  $\ell$ . Since the initial and final atoms involve all processes, the abstract semantics  $\llbracket w \rrbracket$  of a successful run has the form  $\llbracket w \rrbracket(v) = (\max_{p \in Proc} v(p) + t_w(p))_{p \in Proc}$ , where  $t_w(p)$  is the time process  $p$  needs to execute  $w$ . In particular, we have  $\llbracket w \rrbracket(0) = t_w$ . Then  $\llbracket \mathcal{N}, S \rrbracket(0)$  gives the best-case execution time for a scheduler  $S$ , and  $\llbracket \mathcal{N} \rrbracket(0)$  the infimum/ supremum time over the times for each scheduler.

## 2.2. Maximal fixed point of an analysis framework

It is well-known that for sequential flow-graphs the MOP of an analysis framework coincides with the *Maximal Fixed Point* of the framework, or *MFP*. The MOP is the least fixed point of a set of linear equations over the lattice, having one equation for each node of the flow-graph<sup>7</sup>. The least fixed point can be approximated by means of Kleene's theorem, and computed exactly in a number of cases, including the case of lattices satisfying the ascending chain condition, but also others. For example, the lattice for the expected

cost of a flow-graph does not satisfy the ascending chain condition, but yields a set of linear fixed point equations over the rational numbers, which can be solved using standard techniques.

In the concurrent case, the correspondence between MOP and MFP is more delicate. Given a scheduler  $S$ , we can construct the reachability graph of the negotiation diagram, corresponding to the runs compatible with  $S$ . If the graph is finite (for instance, this is always the case if the scheduler is memoryless, i.e., the node selected by the scheduler to extend a run depends only on the configuration reached by the run), then  $\llbracket \mathcal{N}, S \rrbracket$  can be computed as the MFP of this graph, seen as a sequential flow-graph. The corresponding set of linear fixed point equations has one equation for each configuration of the graph. However, this approach has two problems:

- (a) The number of schedulers is infinite, and non-memoryless schedulers may generate an infinite reachability graph; so we do not obtain an algorithm for computing  $\llbracket \mathcal{N} \rrbracket$ .
- (b) Even for memoryless schedulers, the size of the reachability graph may grow exponentially in the size of the negotiation diagram. So the algorithm for computing  $\llbracket \mathcal{N}, S \rrbracket$  needs exponential time, also for lattices with only two elements.

In the remaining of this section we introduce a *Mazurkiewicz-invariant analysis framework*, and show that for this framework and for the class of sound deterministic negotiation diagrams we can overcome these two obstacles. In Section 2.3 we solve problem (a): We show that  $\llbracket \mathcal{N} \rrbracket = \llbracket \mathcal{N}, S \rrbracket$  for every scheduler  $S$  (Theorem 1 below), and so that it suffices to compute  $\llbracket \mathcal{N}, S \rrbracket$  for a scheduler  $S$  of our choice. In the rest of the paper we solve problem (b): We give a procedure that computes  $\llbracket \mathcal{N} \rrbracket$  without ever constructing the reachability graph of the negotiation diagram. The procedure reduces the problem of computing the MOP to computing the MFP of a polynomial number of (sequential) flow-graphs, each of them of size at most linear in the size of the negotiation diagram. This shows that the MOP can be computed in polynomial time for a sound deterministic negotiation diagram iff it can be computed in polynomial time for a sequential flow-graph.

If we remove any of the three conditions of our setting (Mazurkiewicz-invariance, soundness, determinism), then there exist frameworks with the following property: deciding if the MOP has a given value is polynomial in the sequential case (i.e., for flow graphs of sequential programs), but at least NP-hard for negotiations.

We sketch the NP-hardness proof for deterministic, sound negotiations where the framework is not Mazurkiewicz-invariant. Consider the NP-hard problem 1-in-3-SAT, where it is asked if for a CNF formula with  $k$  variables and  $m$  clauses there is an assignment that sets exactly one literal true in each clause. We have  $k$  processes  $p_1, \dots, p_k$ , one for each variable  $x_i$ . We describe the (acyclic) deterministic, sound negotiation  $\mathcal{N}$ . The initial node of  $\mathcal{N}$  has a single outcome, that leads process  $p_i$  to

7. Again, the name ‘‘maximal’’ has historical reasons.

a node with domain  $\{p_i\}$ . From there  $p_i$  branches for the two possible values for  $x_i$ . The “true” branch is a line with outcomes corresponding to clauses that become “true” when  $x_i$  is true, and analogously for the “false” branch - in both cases respecting the order of clauses. Let us denote by  $C_j$  an outcome corresponding to the  $j$ -th clause. The lattice  $\bar{D}$  has elements  $\perp < 1, \dots, (m+1) < \top$ ; so there are  $m+1$  pairwise incomparable elements together with  $\perp$  and  $\top$ . For every node  $n$  and clause  $C_j$  we set  $\llbracket (n, C_j) \rrbracket(j) = j+1$ ,  $\llbracket (n, C_j) \rrbracket(\top) = \top$  and  $\llbracket (n, C_j) \rrbracket(d) = \perp$  otherwise. Moreover  $\llbracket \ell \rrbracket$  is the identity function for all other locations  $\ell$ . This framework is monotonic and distributive. For a run  $w$  of  $\mathcal{N}$  we have  $\llbracket w \rrbracket(1) = m+1$  if the subsequence of clauses appearing in  $w$  is exactly  $C_1 \dots C_m$ ; otherwise  $\llbracket w \rrbracket(1) = \perp$ . Since  $\llbracket \mathcal{N} \rrbracket$  is the  $\sqcup$  over all runs, we get that the 1-in-3-SAT instance is positive iff  $\llbracket \mathcal{N} \rrbracket(1) = m+1$ . In the sequential case, the analysis can be done in polynomial time, since the lattice  $D \rightarrow \bar{D}$  has the height  $\mathcal{O}(m)$ .

The proofs of the other two cases (where determinism or soundness are removed) follow easily from a simple construction shown in Theorem 1 of [6]: Given a deterministic linearly bounded automaton  $A$  and a word  $w$ , one can construct in polynomial time a negotiation  $\mathcal{N}_A$  having one single run that simulates the execution of  $A$  on the input  $w$ . This gives PSPACE-hardness for essentially all non-trivial frameworks, Mazurkiewicz invariant or not.

### 2.3. Mazurkiewicz-invariant analysis frameworks

We introduce the notion of Mazurkiewicz equivalence between runs (also called trace equivalence in the literature [4]). Two equivalent runs started in the same configuration will end up in the same configuration. We call an analysis framework Mazurkiewicz-invariant if the values of equivalent runs are the same. We then show that the MOP of a Mazurkiewicz-invariant analysis is independent of the scheduler.

**Definition 1.** *Two nodes  $n, m$  of a negotiation diagram are independent if  $\text{dom}(n) \cap \text{dom}(m) = \emptyset$ . Two locations are independent if their nodes are independent. Given two finite sequences of locations  $w_1, w_2$ , we write  $w_1 \sim w_2$  if  $w_1 = w\ell_1\ell_2w'$  and  $w_2 = w\ell_2\ell_1w'$  for independent locations  $\ell_1, \ell_2$ . Mazurkiewicz equivalence, denoted by  $\equiv$ , is the reflexive-transitive closure of  $\sim$ .*

The next lemma says that Mazurkiewicz equivalent runs have the same behaviors.

**Lemma 1.** *If  $C_1 \xrightarrow{w} C_2$  and  $v \equiv w$ , then  $C_1 \xrightarrow{v} C_2$ . In particular, if  $w$  is a (successful) run, then  $v$  is.*

Interestingly Mazurkiewicz equivalence behaves very well with respect to schedulers.

**Lemma 2.** *Let  $\mathcal{N}$  be a deterministic negotiation diagram and let  $S$  be a scheduler of  $\mathcal{N}$ . For every successful run  $w$  there is exactly one successful run  $v \equiv w$  that is compatible with  $S$ .*

We observe that Lemma 2 may not hold for runs that are not successful nor for non-deterministic negotiation diagrams.

We can now define Mazurkiewicz-invariant analysis frameworks, and prove that they are independent of schedulers.

**Definition 2.** *An analysis framework is Mazurkiewicz-invariant if  $\llbracket \ell_1 \rrbracket \circ \llbracket \ell_2 \rrbracket = \llbracket \ell_2 \rrbracket \circ \llbracket \ell_1 \rrbracket$  for every two independent outcomes  $\ell_1, \ell_2$ .*

**Theorem 1.** *Let  $\mathcal{N}$  be a negotiation diagram, and let  $\llbracket \_ \rrbracket$  be an analysis framework for  $\mathcal{N}$ . If  $\mathcal{N}$  is deterministic and  $\llbracket \_ \rrbracket$  is Mazurkiewicz-invariant, then  $\llbracket \mathcal{N}, S \rrbracket = \llbracket \mathcal{N}, S' \rrbracket$  for every two schedulers  $S, S'$ , and so  $\llbracket \mathcal{N} \rrbracket = \llbracket \mathcal{N}, S \rrbracket$  for every scheduler  $S$ .*

It turns out that many interesting analysis frameworks are Mazurkiewicz-invariant, or Mazurkiewicz-invariant under natural conditions. Let us look at the examples from Section 2.1.

**The input/output framework** is Mazurkiewicz-invariant if  $\llbracket \ell_1 \rrbracket(\llbracket \ell_2 \rrbracket(\{v\})) = \llbracket \ell_2 \rrbracket(\llbracket \ell_1 \rrbracket(\{v\}))$  holds whenever  $\ell_1$  and  $\ell_2$  are independent. This is not always the case, but holds e.g. when all variables of  $V$  are local variables. Formally, the set  $V$  of variables is partitioned into sets  $V_p$  of local variables for each process  $p$ . Further,  $\llbracket \ell \rrbracket$  involves only the local variables of the processes involved in  $\ell$ : Letting  $(v_\ell, v)$  denote a valuation of  $V$ , split into a valuation  $v_\ell$  of the variables of the processes of  $\ell$  and a valuation  $v$  of the rest, we have  $\llbracket \ell \rrbracket(v_\ell, v) = (v'_\ell, v)$ , and  $\llbracket \ell \rrbracket(v_\ell, v) = \llbracket \ell \rrbracket(v_\ell, v')$  for every  $v_\ell, v, v'$ .

**The anti-pattern framework** is not Mazurkiewicz-invariant. For example if we take some  $\ell \in K$  independent of  $\ell_1$  then  $\llbracket \ell_1 \ell \ell_2 \rrbracket \neq \llbracket \ell \ell_1 \ell_2 \rrbracket$ . However, in Section 5 we will show that there is a Mazurkiewicz-invariant framework for anti-patterns.

**The minimal/maximal expected cost framework** is Mazurkiewicz-invariant. Indeed, it satisfies  $\llbracket \ell_1 \rrbracket \circ \llbracket \ell_2 \rrbracket = \llbracket \ell_2 \rrbracket \circ \llbracket \ell_1 \rrbracket$  for all outcomes  $\ell_1, \ell_2$ , independent or not. Further, by Theorem 1 the expected cost is the same for every scheduler, and so the result of the analysis is the expected cost of the negotiation diagram.

**The best/worse case execution framework** is Mazurkiewicz-invariant. Intuitively, the scheduler introduces an artificial linearization of the nodes, which are however being executed in parallel. As in the previous case, the result of the analysis is the best-case/worst-case execution time of the negotiation diagram (if the negotiation diagram is cyclic and the cycle non-zero time, then the worst-case execution time is infinite).

Our next goal is a generic algorithm for computing the MOP of Mazurkiewicz-invariant frameworks for sound deterministic negotiation diagrams. This will be done in Section 4, but before we will need some results on decomposing negotiation diagrams.

### 3. Decomposing Sound Negotiation Diagrams

We associate with every node  $n$  and every location  $\ell$  of a sound deterministic negotiation diagram  $\mathcal{N}$  a “subnegotiation”  $\mathcal{N}|_n$  and  $\mathcal{N}|_\ell$ , and prove that it is also sound. In Section 4 we use these subnegotiations to define an analysis algorithm for sound deterministic negotiation diagrams. We illustrate the results of this section on the example of Figure 2.

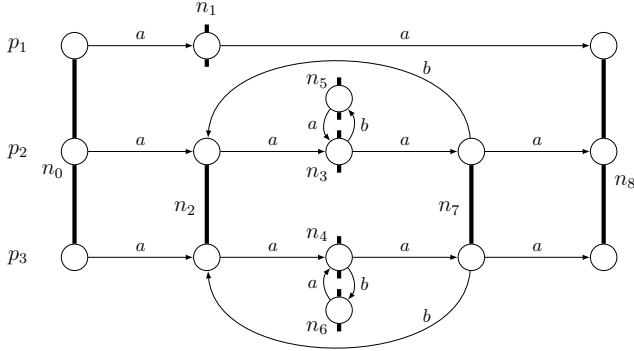


Figure 2. A negotiation diagram with three processes.

Intuitively,  $\mathcal{N}|_n$  contains the nodes  $n'$  such that  $\text{dom}(n') \subseteq \text{dom}(n)$ , with transitions inherited from  $\mathcal{N}$ , and  $n$  as initial node. The non-trivial part is to define the final node and show that  $\mathcal{N}|_n$  is sound. Given a location  $\ell = (n, a)$ , the negotiation  $\mathcal{N}|_\ell$  contains the part of  $\mathcal{N}|_n$  reachable by executions that start with  $\ell$  and afterwards only use nodes with domains *strictly* included in  $\text{dom}(n)$ . Figure 3 shows some of the subnegotiations we will obtain for some nodes and locations of Figure 2.

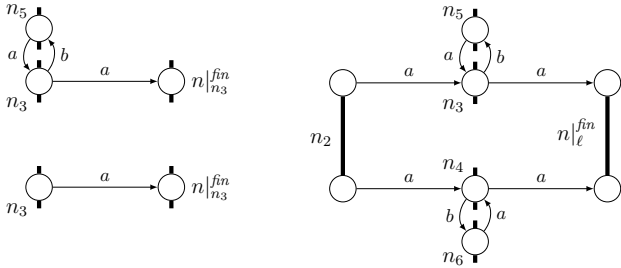


Figure 3. Subnegotiations  $\mathcal{N}|_{n_3}$  (top left),  $\mathcal{N}|_{(n_3,a)}$  (bottom left), and  $\mathcal{N}|_{(n_2,a)}$  (right) of the negotiation diagram of Figure 2. Nodes unreachable from the initial node are not shown.

The rest of the section first presents a theorem showing the existence and uniqueness of some special configurations (Section 3.1), and then uses it to define  $\mathcal{N}|_n$  and  $\mathcal{N}|_\ell$ , and prove their soundness (Section 3.2).

#### 3.1. Unique maximal configurations

Given a node  $n$  of a sound and deterministic negotiation, we prove the existence of a unique reachable configuration

$I(n)$  enabling  $n$  and only  $n$ . Then we show the following: if we start from  $I(n)$  as initial configuration, “freeze” the processes of  $\text{Proc} \setminus \text{dom}(n)$ , and let the processes of  $\text{dom}(n)$  execute maximally (i.e., until they cannot execute any node without the help of processes of  $\text{Proc} \setminus \text{dom}(n)$ ), then we *always* reach the same “final” configuration  $F(n)$ . Additionally, given a location  $\ell = (n, a)$  we show: If from  $I(n)$  we execute  $\ell$  and let the processes of  $\text{dom}(n)$  proceed until no enabled node  $m$  satisfies  $\text{dom}(m) \subset \text{dom}(n)$ , then again we always reach the same “final” configuration  $F(\ell)$ .

Let  $X \subseteq \text{Proc}$  be a set of processes. A sequence of locations  $\ell_1, \dots, \ell_k$  is an  $X$ -sequence if the domains of all  $\ell_i$  are included in  $X$ ; it is a *strict*  $X$ -sequence if moreover the domains of all  $\ell_i$ , but possibly  $\ell_1$ , are strictly included in  $X$ . We write (strict)  $n$ -sequence for (strict)  $\text{dom}(n)$ -sequence.

We write  $n' \preceq n$  if  $\text{dom}(n') \subseteq \text{dom}(n)$ , and  $n' \prec n$  if  $\text{dom}(n') \subsetneq \text{dom}(n)$  (and similarly for  $\ell' \preceq \ell, \ell' \prec \ell, \ell \preceq n$  etc). Our goal is to prove:

**Theorem 2.** *Let  $m$  be a reachable node of a sound deterministic negotiation diagram  $\mathcal{N}$ .*

- (i) *There is a unique reachable configuration  $I(m)$  of  $\mathcal{N}$  that enables  $m$ , and no other node.*
- (ii) *There is a unique configuration  $F(m)$  such that*
  - *$F(m)$  is reachable from  $I(m)$  by means of an  $m$ -sequence, and*
  - *for every node  $n$  enabled at  $F(m)$ ,  $\text{dom}(n)$  is not included in  $\text{dom}(m)$ .*
- (iii) *For every location  $\ell$  of  $m$  there is a unique configuration  $F(\ell)$  such that*
  - *$F(\ell)$  is reachable from  $I(m)$  by means of a strict  $m$ -sequence starting with  $\ell$ , and*
  - *for every node  $n$  enabled at  $F(\ell)$ ,  $\text{dom}(n)$  is not strictly included in  $\text{dom}(m)$ .*

E.g., in Figure 2 we have  $I(n_1) = (n_1, n_8, n_8)$  (an abbreviation for  $I(n_1)(p_1) = \{n_1\}, I(n_1)(p_2) = \{n_8\}, I(n_1)(p_3) = \{n_8\}$ );  $I(n_2) = (n_8, n_2, n_2)$ ; and  $I(n_3) = (n_8, n_3, n_7)$ . Moreover,  $F(n_1) = F(n_2) = (n_8, n_8, n_8)$ ; and  $F(n_3) = (n_8, n_7, n_7)$ . Further, we get  $F(n_7, a) = (n_8, n_7, n_7)$  and  $F(n_7, b) = (n_8, n_2, n_2)$ .

The proof of Theorem 2 is quite involved. The theorem is a consequence of the Unique Configuration lemma (Lemma 4 below), which relies on the Domination lemma (Lemma 3 below), which in turn is based on results of [7], [10].

A *local path* of a negotiation diagram  $\mathcal{N}$  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 *local circuit* if  $k > 0$  and  $n_0 = n_k$ . A local path is *reachable* if some node in the path is reachable. The domination lemma says that every local circuit has some node containing all processes of the circuit.

**Lemma 3. (Domination Lemma)** *Let  $\mathcal{N}$  be a deterministic sound negotiation diagram. Every reachable local circuit of  $\mathcal{N}$  contains a dominant node, i.e. a node  $n$  such that  $m \preceq n$ , for every node  $m$  of the circuit.*

The unique configuration lemma says that if two enabled configurations agree on a set of processes  $X$  and every enabled action in one of the two configurations needs a process from  $X$ , then the two configurations are actually the same.

**Lemma 4. (Unique Configuration Lemma)** *Let  $X \subseteq \text{Proc}$  be a set of processes. Let  $C_1, C_2$  be reachable configurations such that (1)  $C_1(p) = C_2(p)$  for every  $p \in X$ , and (2) every node  $n$  enabled at  $C_1$  or  $C_2$  satisfies  $\text{dom}(n) \cap X \neq \emptyset$ . Then  $C_1 = C_2$ .*

The above lemma gives Theorem 2 rather directly. For (i), take  $X = \text{dom}(m)$ . Suppose that there are two configurations  $I_1$  and  $I_2$  as in (i). The hypotheses of Lemma 4 are satisfied, and so  $I_1 = I_2$ . The case (ii) is equally easy, while (iii) is only a bit more involved.

### 3.2. Subnegotiations for nodes and locations

We use Theorem 2 to define the subnegotiations  $\mathcal{N}|_n$  and  $\mathcal{N}|_\ell$  for each node  $n$  and location  $\ell$  of a sound deterministic negotiation diagram  $\mathcal{N}$ , and prove that they are sound.

**Definition 3.** *Let  $\mathcal{N}$  be a sound deterministic negotiation diagram and let  $n$  be a reachable node of  $\mathcal{N}$ . The negotiation diagram  $\mathcal{N}|_n$  contains all the nodes and locations that appear in the  $n$ -sequences  $u$  such that  $I(n) \xrightarrow{u} F(n)$ , plus a new final node  $n|_n^{\text{fin}}$  with  $\text{dom}(n|_n^{\text{fin}}) = \text{dom}(n)$ . The initial node is  $n$ , and the transition function  $\delta|_n$  is defined as follows. For given  $m, a, p$ , we set:*

$$\delta|_n(m, a, p) = \begin{cases} \delta(m, a, p) & \text{if } \delta(m, a, p) \neq F(n)(p) \\ n|_n^{\text{fin}} & \text{if } \delta(m, a, p) = F(n)(p) \end{cases}$$

An example of  $\mathcal{N}|_n$  is given on the left of Figure 3.

**Lemma 5.** *If  $\mathcal{N}$  is a sound deterministic negotiation diagram then so is  $\mathcal{N}|_n$ .*

The subnegotiation  $\mathcal{N}|_\ell$  induced by a location  $\ell = (n, a)$  is defined analogously to  $\mathcal{N}|_n$ , with two differences. First, in  $\mathcal{N}|_\ell$  the node  $n$  has  $a$  as the unique outcome. Second, the domain of every node of  $\mathcal{N}|_\ell$ , except the node  $n$  itself, is strictly included in  $\text{dom}(n)$ .

**Definition 4.** *Let  $\mathcal{N}$  be a deterministic sound negotiation diagram, let  $n$  be a reachable node of  $\mathcal{N}$ , and let  $\ell = (n, a)$  for some outcome  $a$  of  $n$ . The negotiation diagram  $\mathcal{N}|_\ell$  contains all the nodes and locations that appear in the strict  $n$ -sequences  $u$  such that  $I(n) \xrightarrow{\ell u} F(\ell)$ , plus a new final node  $n|_\ell^{\text{fin}}$ . The initial node is  $n$ , it has the unique outcome  $a$ , and the transition function  $\delta|_\ell$  is defined as follows. For given  $m, b, p$  we set:*

$$\delta|_\ell(m, b, p) = \begin{cases} \delta(m, b, p) & \text{if } \delta(m, b, p) \neq F(\ell)(p) \\ n|_\ell^{\text{fin}} & \text{otherwise} \end{cases}$$

Figure 3 shows (on the right)  $\mathcal{N}|_\ell$  for the location  $\ell = (n_2, a)$  of the negotiation diagram of Figure 2.

**Lemma 6.** *If  $\mathcal{N}$  is a sound deterministic negotiation diagram, then so is  $\mathcal{N}|_\ell$ .*

## 4. Computing the MOP

We use the decomposition of Section 3 to define an algorithm computing the MOP for Mazurkiewicz-invariant frameworks and arbitrary sound deterministic negotiation diagram. The idea is to repeatedly reduce parts of the negotiation diagram without changing the meaning of the whole negotiation. When reduction will be no longer possible, the negotiation diagram will have only one location, whose value will be the value of the negotiation. The goal of this section is to present Algorithm 1 and Theorem 3 that is our main result.

As we have seen in the previous section, for every node  $n$ , the negotiation diagram  $\mathcal{N}|_n$  is sound and the final configuration  $F(n)$  is unique. Thus we can safely replace all the transitions from  $n$  by one transition going directly to  $F(n)$  and assign to this transition the value of  $\mathcal{N}|_n$ . This requires to be able to compute  $F(n)$  as well as the value of  $\mathcal{N}|_n$ . For this we proceed by induction on the domain of  $n$  starting from nodes with the smallest domain. As we will see, this will require us to compute MOP only for negotiations of two (very) special forms

**One-trace negotiations.** These are acyclic negotiation diagrams in which every node has one single outcome. In this degenerate case, all the executions of the negotiation diagram are Mazurkiewicz equivalent; moreover, by acyclicity, the trace contains every location at most once. Since the analysis framework is Mazurkiewicz-invariant, we have  $\llbracket \mathcal{N} \rrbracket = \llbracket w \rrbracket$  for any successful run  $w$ . A successful run can be computed by just executing the negotiation diagram with some arbitrary scheduler. Once a successful run  $w$  is computed, we extract from it a flow-graph with  $|w|$  nodes, that is actually a sequence, and compute MFP.

**Replications.** Intuitively, a replication is a negotiation diagram in which all processes are involved in every node, and all processes move uniformly, that is, after they agree on an outcome they all move to the same node. Formally, a negotiation diagram is a *replication* if for every reachable node  $n$  and every outcome  $(n, a)$  there is a node  $m$  such that  $\delta(n, a, p) = m$  for every process  $p$ . Observe that, in particular, all nodes of a replication have the same domain. It follows that (the reachable part of) a replication is a flow-graph “in disguise”. More precisely, we can assign to it a flow-graph having one node for every reachable node, and an edge for every location  $(n, a)$ , leading from  $n$  to  $\delta(n, a, p)$ , where  $p$  can be chosen arbitrarily out of  $\text{dom}(n)$ . It follows immediately that the MOP for the negotiation diagram is equal to the MFP of this flow-graph.

**Definition 5.** *A node  $n$  is reduced if it has one single outcome, and for this single outcome  $a$  we have  $\delta(n, a, p) = F(n)(p)$  for every  $p \in \text{dom}(n)$ .*

*A location  $\ell = (n, a)$  is reduced if  $\delta(n, a, p) = F(\ell)(p)$  for every  $p \in \text{dom}(n)$ .*

Now we will define an operation of reducing nodes and locations in a negotiation diagram; this is the core operation



---

**Algorithm 1** Algorithm computing MOP for a sound deterministic negotiation diagram  $\mathcal{N}$ .

---

```

1: while  $\mathcal{N}$  has non-reduced nodes do
2:    $m := \prec$ -minimal, non-reduced node
3:    $X = \text{dom}(m)$ 
4:   for every  $\ell = (n, a)$  with  $\text{dom}(n) = X$  do
5:      $\#\# \mathcal{N}|_\ell$  is a one-trace negotiation  $\#\#$ 
6:      $\llbracket n, a_\ell \rrbracket := \text{MOP}(\mathcal{N}|_\ell)$ 
7:      $\mathcal{N} := \text{Red}_\ell(\mathcal{N})$ 
8:   end for
9:   for every node  $n$  such that  $\text{dom}(n) = X$  do
10:     $\#\# \mathcal{N}|_n$  is a replication  $\#\#$ 
11:     $\llbracket n, a_n \rrbracket = \text{MOP}(\mathcal{N}|_n)$ 
12:   end for
13:   for every node  $n$  such that  $\text{dom}(n) = X$  do
14:     $\mathcal{N} := \text{Red}_n(\mathcal{N})$ 
15:   end for
16: end while
17: return  $\llbracket \ell \rrbracket$ , where  $\ell$  is the unique outcome of the initial
    node of  $\mathcal{N}$ 

```

---

of Algorithm 1. For a location  $\ell = (n, a)$ , the operation  $\text{Red}_\ell(\mathcal{N})$  removes the transition of  $n$  on  $\ell$  and adds a new transition on  $(n, a_\ell)$ . Similarly  $\text{Red}_n(\mathcal{N})$ , but this time it removes all transitions from  $n$  and adds a single new transition on  $(n, a_n)$ .

**Definition 6.** Let  $\mathcal{N} = \langle \text{Proc}, N, \text{dom}, R, \delta \rangle$  be a sound deterministic negotiation diagram and let  $\ell = (n, a)$  be a non-reduced outcome of  $\mathcal{N}$ . The negotiation diagram  $\text{Red}_\ell(\mathcal{N})$  has the same components as  $\mathcal{N}$  but for  $\text{out}$  and  $\delta$  that are subject to the following changes:

- $\text{out}(n) := (\text{out}(n) \setminus \{a\}) \cup \{a_\ell\}$ ;
- $\delta(n, a_\ell, p) := F(\ell)(p)$  for every process  $p \in \text{dom}(n)$ .

The negotiation diagram  $\text{Red}_n(\mathcal{N})$  is defined similarly but now:

- $\text{out}(n) := \{a_n\}$ ;
- $\delta(n, a_n, p) := F(n)(p)$  for every process  $p \in \text{dom}(n)$ .

The next lemma states that these reduction operations preserve the meaning of a negotiation diagram.

**Lemma 7.** Let  $\mathcal{N}$  and  $\ell = (n, a)$  be as in Definition 6. Assign to the new location  $\ell' = (n, a_\ell)$  the mapping  $\llbracket \ell' \rrbracket := \llbracket \mathcal{N}|_\ell \rrbracket$ . Then  $\llbracket \mathcal{N} \rrbracket = \llbracket \text{Red}_\ell(\mathcal{N}) \rrbracket$ . Analogously,  $\llbracket \mathcal{N} \rrbracket = \llbracket \text{Red}_n(\mathcal{N}) \rrbracket$  when we assign  $\llbracket (n, a_n) \rrbracket = \llbracket \mathcal{N}|_n \rrbracket$ .

At this point we can examine Algorithm 1. The algorithm repeatedly applies reduction operations to a given negotiation diagram. Thanks to Lemma 7 these reductions preserve the meaning of the negotiation diagram. At every reduction, the number of reachable locations in the negotiation diagram decreases. So the algorithm stops, and when it stops the negotiation diagram has only one reachable location. The abstract semantics of this location is equal to the abstract semantics of the original negotiation diagram.

This argument works if indeed we can compute  $\text{MOP}(\mathcal{N}|_\ell)$  and  $\text{MOP}(\mathcal{N}|_n)$  in lines 6 and 11 of the algorithm, respectively. For this it is enough to show that the invariants immediately preceding these lines hold, as this would mean that we deal with special cases we have discussed at the beginning of this section. The following lemma implies that the invariants indeed hold.

**Lemma 8.** Let  $\mathcal{N}$  be a sound deterministic negotiation diagram, and  $n$  a node such that all nodes  $m \prec n$  are reduced in  $\mathcal{N}$ .

- (1) if  $a$  is an outcome of  $n$ , then for  $\ell = (n, a)$  the negotiation diagram  $\mathcal{N}|_\ell$  is a one-trace negotiation.
- (2) if all locations  $\ell' \prec n$  are reduced, then  $\mathcal{N}|_n$  is a replication.

**Example:** Consider the negotiation diagram of Figure 2. Assume that all locations have cost 1, and that the probability of a location  $\ell = (n, a)$  is  $1/|\text{out}(n)|$  (so, for example, the locations  $(n_3, a)$  and  $(n_3, b)$  have probability  $1/2$ , while  $(n_0, a)$  has probability 1). We compute the expected cost of the diagram using Algorithm 1.

The minimal non-reduced nodes w.r.t.  $\prec$  are  $n_3, n_4, n_5, n_6$ . All their locations satisfy  $\llbracket \mathcal{N}|_\ell \rrbracket = \llbracket \ell \rrbracket$ . The algorithm computes  $\text{MOP}(\mathcal{N}|_{n_i})$  for  $i = 3, 4, 5, 6$ . The subnegotiations  $\mathcal{N}|_{n_3}$  and  $\mathcal{N}|_{n_5}$  are shown in Figure 3; the other two are similar.  $\text{MOP}(\mathcal{N}|_{n_3})$  is the expected cost of reaching  $n_{fin}$  from  $n_3$  in  $\mathcal{N}|_{n_3}$ . Since  $\mathcal{N}|_{n_3}$  is a replication (in fact, it is even a flow-graph), we can compute it as the least solution of the following fixed point equation, where we abbreviate  $\text{Prob}(\ell)$  to  $P(\ell)$  and  $\text{Cost}(\ell)$  to  $C(\ell)$ :

$$(p, c) = \left( p \cdot P(n_3, b) \cdot P(n_5, a) + P(n_3, a), \right. \\ \left. P(n_3, b) \cdot P(n_5, a) \cdot (C(n_3, b) + C(n_5, a) + c) + \right. \\ \left. P(n_3, a) \cdot C(n_3, a) \right)$$

which gives

$$(p, c) = \left( \frac{1}{2}p + \frac{1}{2}, \frac{1}{2}(2 + c) + \frac{1}{2} \right)$$

with least fixed point  $(1, 3)$ , which of course can be computed by just solving the linear equation. So  $\text{MOP}(\mathcal{N}|_{n_3}) = (1, 3)$ . We obtain  $\llbracket n_3, a_{n_3} \rrbracket = \llbracket n_4, a_{n_4} \rrbracket = 3$ ,  $\llbracket n_5, a_{n_5} \rrbracket = \llbracket n_6, a_{n_6} \rrbracket = 4$ , and the reduced negotiation diagram at the top of Figure 4. Observe that nodes  $n_5$  and  $n_6$  are no longer reachable.

The minimal non-reduced nodes are now  $n_2$  and  $n_7$ . The locations of  $n_7$  satisfy  $\llbracket \mathcal{N}|_\ell \rrbracket = \llbracket \ell \rrbracket$ . For the location  $(n_2, a)$  we obtain  $\text{MOP}(\mathcal{N}|_{(n_2, a)}) = (1, 1 + 3 + 3) = (1, 7)$ ; this we can easily do because  $\mathcal{N}|_{(n_2, a)}$  is a one-trace negotiation; we just pick any successful run, for example  $(n_2, a)(n_3, a_{n_3})(n_4, a_{n_4})$ , and add the costs. After reducing  $\mathcal{N}|_{(n_2, a)}$  we obtain the negotiation diagram in the middle of Figure 4; the non-reachable nodes  $n_3, \dots, n_6$  are no longer displayed, and all locations have cost 1 but  $(n_2, a_{(n_2, a)})$ , which has cost 7. Further, all locations  $\ell$  of  $n_2$  and  $n_7$

satisfy  $\llbracket \mathcal{N} \rrbracket_\ell = \llbracket \ell \rrbracket$ . We compute  $\mathcal{N}|_{n_2}$  and  $\mathcal{N}|_{n_7}$  by a fixed point calculation analogous to the one above (observe that both of them are replications), and after reduction obtain the negotiation diagram at the bottom of the figure, with  $\llbracket n_7, a_{n_7} \rrbracket = (1, 9)$  and  $\llbracket n_2, a_{n_2} \rrbracket = (1, 7 + 9) = (1, 16)$ . Now, the minimal non-reduced node is  $n_0$ . We compute  $MOP(\mathcal{N}|_{(n_0, a)})$  ( $\mathcal{N}|_{(n_0, a)}$  is a one-trace negotiation), and obtain  $\llbracket n_0, a \rrbracket = (1, 1 + 1 + 16) = (1, 18)$ , which is the final result.  $\square$

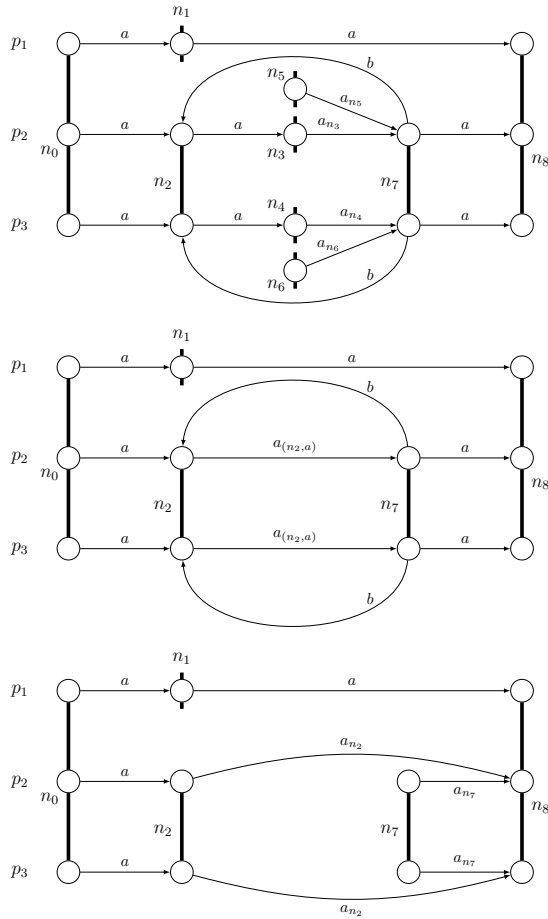


Figure 4. Three reduced negotiation diagrams constructed by Algorithm 1 started on the negotiation diagram of Figure 2.

Finally, to estimate the complexity of the algorithm, we should also examine how to calculate  $Red_\ell(\mathcal{N})$  and  $Red_n(\mathcal{N})$  in lines 7 and 14 of the algorithm. In order to calculate  $Red_\ell(\mathcal{N})$  we need to know  $F(\ell)$ . The invariant says that  $\mathcal{N}|_\ell$  at that point has one trace. So it is enough to execute this trace in  $\mathcal{N}|_\ell$  to reach  $F(\ell)$ . Similarly, for  $Red_n(\mathcal{N})$  we need to know  $F(n)$ . The invariant says that  $\mathcal{N}_n$  is a replication, so it is just a flow graph with one final node  $n_{fin}$ . We can execute in  $\mathcal{N}$  any path leading to the exit to calculate  $F(n)$ . To sum up, the calculations of  $Red_\ell(\mathcal{N})$  and  $Red_n(\mathcal{N})$  in lines 7 and 14 can be done in linear time with respect to the size of  $\mathcal{N}$ .

We summarize the results of this section:

**Theorem 3.** *Let  $\mathcal{N}$  be a sound deterministic negotiation diagram, and  $\llbracket \_ \rrbracket$  a Mazurkiewicz invariant analysis framework. Algorithm 1 stops and outputs  $\llbracket \mathcal{N} \rrbracket$ . The complexity of the algorithm is  $\mathcal{O}(|\mathcal{N}|(C + |\mathcal{N}|))$  where  $|\mathcal{N}|$  is the size of  $\mathcal{N}$ , and  $C$  is the cost of  $\llbracket \_ \rrbracket$  analysis for flow-graphs of size  $|\mathcal{N}|$ .*

## 5. Anti-Patterns and Gen/Kill Analyses

We saw in Section 2.1 that the problem of detecting an anti-pattern can be naturally captured as an analysis framework, which however is not Mazurkiewicz-invariant. We now show that, while the *natural* framework is not Mazurkiewicz invariant, an equivalent framework that returns the same result is. Then we sketch how this result generalizes to arbitrary Gen/Kill analysis frameworks, a much studied class [20].

We need some basic notions of Mazurkiewicz trace theory [4]. Given a sequence  $w = w_1 \dots w_n \in \mathcal{L}^*$ , define  $i \sqsubseteq' j$  for two positions  $i, j$  of  $w$  if  $i \leq j$  and  $w_i$  is not independent from  $w_j$  (see Definition 1). Further, define  $\sqsubseteq$  is the transitive closure of the relation  $\sqsubseteq'$ . It is well-known that if  $w \equiv v$ , i.e., if  $w, v$  are Mazurkiewicz equivalent, then  $\sqsubseteq_w$  and  $\sqsubseteq_v$  are isomorphic as labeled partial orders (the labels being the locations). We write  $btw(i, j)$  for the set of positions between  $i$  and  $j$ , i.e.,  $\{k : i \sqsubseteq k \sqsubseteq j\}$ .

Recall that anti-pattern analysis asked if there is an execution with the property “ $w \in L$ ” for the language  $L = \mathcal{L}^* \ell_1 (\overline{K})^* \ell_2 \mathcal{L}^*$ . Instead of this property consider the following property of  $w$ :

- (\*) there are two positions  $i, j$  such that  $w_i = \ell_1$ ,  $w_j = \ell_2$ ,  $btw(i, j) \cap K = \emptyset$ , and not  $j \sqsubseteq i$ .

The special case of this condition is when  $btw(i, j) = \emptyset$ ; then, since not  $j \sqsubseteq i$ , we actually know that the two positions  $i, j$  are concurrent, so  $w$  is Mazurkiewicz equivalent to  $u \ell_1 \ell_2 v$ , for some  $u, v$ .

**Lemma 9.** *A negotiation diagram  $\mathcal{N}$  has a successful run  $w \in L$  iff it has a successful run  $v$  with property (\*).*

It remains to set a static analysis framework for tracking property (\*). Since this is a property of Mazurkiewicz traces, the framework is Mazurkiewicz-invariant.

We need one more piece of notation. For a word  $w$  and a process  $p$  we write  $btw(\ell_1, p)$  for the set  $btw(i, j)$  where  $i$  is the last occurrence of  $\ell_1$ , and  $j$  is the last occurrence of a location using process  $p$ .

We define now an auxiliary function  $\alpha$  from sequences to  $\mathcal{P}(\text{Proc})^2 \cup \{\top\}$ . We set  $\alpha(w) = \top$  if  $w$  has property (\*), otherwise  $\alpha(w) = (P_A, P_B)$  where

- $p \in P_A$  if  $btw(\ell_1, p) \neq \emptyset$  and  $btw(\ell_1, p) \cap K = \emptyset$ ,
- $p \in P_B$  if  $btw(\ell_1, p) \neq \emptyset$  and  $btw(\ell_1, p) \cap K \neq \emptyset$ .

We describe a PTIME computable function  $F$  such that for every sequence  $w$  and location  $\ell$ :

$$\alpha(w \ell) = F(\alpha(w), \ell)$$

For defining  $F$  we first describe the update of  $btw(\ell_1, p)$  when extending  $w$  by  $\ell$ . Let us detail the more interesting

case where  $\ell \neq \ell_1$ . Observe that the set of positions  $btw(\ell_1, p)$  does not change if  $p \notin \text{dom}(\ell)$ . If  $p \in \text{dom}(\ell)$  then the update of  $btw(\ell_1, p)$  is the union of  $btw(\ell_1, q)$  over all  $q \in \text{dom}(\ell)$ , plus  $\ell$ . According to these observations, we define the update of  $F$  in the case  $\ell \neq \ell_1$  as follows. If  $p \notin \text{dom}(\ell)$  then process  $p$  remains in its set,  $P_A$  or  $P_B$ . If  $p \in \text{dom}(\ell)$ , then:  $p$  goes into the set  $P'_B$  if either there was some  $q \in \text{dom}(\ell)$  in  $P_B$ , or  $\ell \in K$  and there is some  $q \in \text{dom}(\ell)$  in  $P_A$ ;  $p$  goes into the set  $P'_A$  if  $\ell \notin K$ , no  $q \in \text{dom}(\ell)$  is in  $P_B$  and there is at least one  $q \in \text{dom}(\ell)$  in  $P_A$ .

The function  $F$  can be extended to a monotonic and distributive function  $\widehat{F}$  on  $\mathcal{P}(\text{Proc})^3 \cup \{\top\}$  ordered componentwise, by turning  $\alpha(w)$  into a partition of  $\text{Proc}$  (adding a component  $P_C = \text{Proc} \setminus (P_A \cup P_B)$ ) and embedding it into a suitable function over  $\mathcal{P}(\text{Proc})^3 \cup \{\top\}$ .

With the help of the function  $\widehat{F}$  we define now the value of each location:

$$\llbracket \ell \rrbracket(P_A, P_B, P_C) = \widehat{F}((P_A, P_B, P_C), \ell)$$

Observe that this gives us  $\llbracket w \rrbracket = \alpha(w)$  for every sequence  $w$ . The above discussion yields two lemmas showing that  $\llbracket \cdot \rrbracket$  is a Mazurkiewicz invariant analysis framework that can be computed in PTIME, since  $\widehat{F}$  can be computed in PTIME.

**Lemma 10.**  $\llbracket \cdot \rrbracket$  is Mazurkiewicz-invariant.

**Lemma 11.** Consider a negotiation diagram  $\mathcal{N}$  over set of locations  $\mathcal{L}$ . For every sequence  $w \in \mathcal{L}^*$ :  $\llbracket w \rrbracket(\emptyset, \emptyset, \text{Proc}) = \top$  iff  $w \in L$ . Moreover,  $\llbracket \mathcal{N} \rrbracket(\emptyset, \emptyset, \text{Proc}) = \top$  iff  $\mathcal{N}$  has a successful execution in  $L$ .

## 5.1. Generalization to Gen/Kill analyses

We consider general Gen/Kill analyses<sup>8</sup>. We are given a set of locations  $G \subseteq \mathcal{L}$  that *generate* something, and a set of locations  $K \subseteq \mathcal{L}$  (not necessarily disjoint with  $G$ ) that *kill* this something. The lattice  $\mathcal{D}$  has just two elements  $\{0, 1\}$ , with  $\wedge$  and  $\vee$  as lattice operations, and the transformer of a program instruction  $\ell$  is of the form  $\llbracket \ell \rrbracket(v) = (v \wedge (\ell \notin K)) \vee (\ell \in G)$ . Classical examples from the static analysis of programs are the computation of reaching definitions, available expressions, live variables, very busy expressions, where the “something” are values assigned to a variable or an expression [24]. The four main classes of Gen/Kill analyses differ only on whether control-flow is interpreted forward or backwards, and on whether we do “merge over all paths” or “meet over all paths”.

- **may/forward.** For some configuration  $C$  there is an execution  $C_{init} \xrightarrow{w} C$  with  $w \in \mathcal{L}^* G(\overline{K})^* \ell$ .
- **must/forward.** For every configuration  $C$  and every execution  $C_{init} \xrightarrow{w} C$ , if  $w$  ends with  $\ell$  then  $w \in \mathcal{L}^* G(\overline{K})^* \ell$ .
- **may/backward.** For some reachable configuration  $C$  there is an execution  $C \xrightarrow{w} C_{fin}$  with  $w \in \ell(\overline{K})^* G \mathcal{L}^*$ .

8. Although not in bitvector form, which we leave for future work (see the conclusions).

- **must/backward.** For every reachable configuration  $C$  and every execution  $C \xrightarrow{w} C_{fin}$ , if  $w$  starts with  $\ell$  then  $w \in \ell(\overline{K})^* G \mathcal{L}^*$ .

Observe that in backward properties we require that a configuration  $C$  is reachable from the initial configuration. For forward properties we do not need to assume that a configuration is co-reachable from the final configuration, as this will be immediately implied by soundness.

The two existential properties above can be expressed in terms of the existence of successful executions of a particular form: executions  $C_{init} \xrightarrow{w} C_{fin}$  with  $w$  belonging to some language. The same is true for the negation of the universal properties. Consider the languages given by regular expressions:

- 1)  $E_1 = \mathcal{L}^* G(\overline{K})^* \ell \mathcal{L}^*$ ,
- 2)  $E_2 = (\overline{K} \cap \overline{G})^* \ell \mathcal{L}^* \cup \mathcal{L}^* (K \cap \overline{G})(\overline{K} \cap \overline{G})^* \ell \mathcal{L}^*$ ,
- 3)  $E_3 = \mathcal{L}^* \ell(\overline{K})^* G \mathcal{L}^*$ ,
- 4)  $E_4 = \mathcal{L}^* \ell(\overline{K} \cap \overline{G})^* \cup \mathcal{L}^* \ell(\overline{K} \cap \overline{G})^* (K \cap \overline{G}) \mathcal{L}^*$ .

**Lemma 12.** For a sound negotiation diagram we have the following:

- *may/forward* is equivalent to  $\exists C_{init} \xrightarrow{w} C_{fin}$  with  $w \in E_1$ .
- *negation of must/forward* is equivalent to  $\exists C_{init} \xrightarrow{w} C_{fin}$  with  $w \in E_2$ .
- *may/backward* is equivalent to  $\exists C_{init} \xrightarrow{w} C_{fin}$  with  $w \in E_3$ .
- *negation of must/backward* is equivalent to  $\exists C_{init} \xrightarrow{w} C_{fin}$  with  $w \in E_4$ .

The resource analysis at the beginning of this section corresponds to  $E_3$ . For each one of  $E_1, E_2, E_4$  it is easy to produce an analogue of Lemma 9 reformulating the property in trace terms. This allows us to check all properties in polynomial time using our algorithm for Mazurkiewicz invariant analysis frameworks.

## 6. Conclusions

Previous work had identified deterministic negotiations – a model of concurrency essentially isomorphic to free-choice workflow Petri nets – as a class that has both practical relevance for business process modeling, and admits PTIME analysis for several important properties once negotiations are assumed to be sound. Moreover soundness is a natural prerequisite that can be checked in PTIME.

We have proposed a general notion of Mazurkiewicz-invariant analysis frameworks. We have shown that computing the MOP in such frameworks for sound deterministic negotiations is as easy as computing it for sequential flow graphs (while computing the MOP of general frameworks takes exponentially longer, unless PTIME=NP). This result not only subsumes all previous PTIME results on analysis of sound deterministic negotiations, but also yields PTIME algorithms for new problems, like the computation of the best-case/worst-case execution time, the detection of anti-patterns, and general gen/kill analysis problems. The result is particularly interesting for gen/kill problems: While

their natural formulation is not in terms of Mazurkiewicz-invariant frameworks, we have shown that they can be reformulated as such.

In future work we plan to improve the degree of the polynomial bounding the runtime of our algorithm. Since our decomposition does not partition a negotiation into disjoint parts, when computing MOPs of subnegotiations we are redoing computations. Bounding the size of overlaps looks like a promising way of bringing the complexity on a par with the sequential case. Section 5 on gen/kill analyses raises a further question. In the sequential case, a gen/kill analysis can be simultaneously computed for all program points and all program variables (for example, one can compute for each program point the set of live variables at that point). This is not yet the case in our algorithm. In fact, it is not clear what is a program point in a negotiation: If one takes a configuration as a program point, then, since the number of reachable configurations can grow exponentially in the size of the negotiation diagram, any algorithm that explicitly computes the MOP for each reachable configuration has exponential worst-case complexity.

**Acknowledgments.** We thank the anonymous reviewers for useful remarks, and Jörg Desel, Philipp Hoffmann, and Denis Kuperberg for helpful discussions.

## References

- [1] A. Bouajjani and M. Emmi. Analysis of recursively parallel programs. *ACM Trans. Program. Lang. Syst.*, 35(3):10:1–10:49, 2013.
- [2] R. Chugh, J. W. Voung, R. Jhala, and S. Lerner. Dataflow analysis for concurrent programs using datarace detection. In *SIGPLAN 2008*, pages 316–326. ACM, 2008.
- [3] J. Desel and J. Esparza. Negotiations and Petri nets. In *PNSE’15*, volume 1372 of *CEUR Workshop Proceedings*, pages 41–57. CEUR-WS.org, 2015.
- [4] V. Diekert and G. Rozenberg, editors. *The Book of Traces*. World Scientific, 1995.
- [5] C. Eisentraut, H. Hermanns, J. Katoen, and L. Zhang. A semantics for every GSPN. In *PETRI NETS 2013*, volume 7927 of *LNCS*, pages 90–109, 2013.
- [6] J. Esparza and J. Desel. On negotiation as concurrency primitive. In *CONCUR*, volume 8052 of *LNCS*, pages 440–454, 2013. Extended version in arXiv:1307.2145.
- [7] J. Esparza and J. Desel. On negotiation as concurrency primitive II: Deterministic cyclic negotiations. In *FoSSaCS*, volume 8412 of *LNCS*, 2014.
- [8] J. Esparza and P. Hoffmann. Reduction rules for colored workflow nets. In *FASE 2016*, volume 9633 of *LNCS*, pages 342–358, 2016.
- [9] J. Esparza, P. Hoffmann, and R. Saha. Polynomial analysis algorithms for free choice probabilistic workflow nets. In *QEST 2016*, volume 9826 of *LNCS*, pages 89–104, 2016.
- [10] J. Esparza, D. Kuperberg, A. Muscholl, and I. Walukiewicz. Soundness in negotiations. In *CONCUR 2016*, volume 59 of *LIPICs*, pages 12:1–12:13, 2016.
- [11] J. Esparza, A. Muscholl, and I. Walukiewicz. Static analysis of deterministic negotiations. arXiv:1704.04190, 2017.
- [12] J. Esparza and A. Podelski. Efficient algorithms for pre\* and post\* on interprocedural parallel flow graphs. In *POPL 2000*, pages 1–11. ACM, 2000.
- [13] D. Fahland, C. Favre, B. Jobstmann, J. Koehler, N. Lohmann, H. Völzer, and K. Wolf. Instantaneous soundness checking of industrial business process models. In *Business Process Management*, pages 278–293. Springer, 2009.
- [14] D. Fahland and H. Völzer. Dynamic skipping and blocking and dead path elimination for cyclic workflows. In *Business Process Management*, volume 9850 of *LNCS*, pages 234–251. Springer, 2016.
- [15] A. Farzan and Z. Kincaid. Compositional bitvector analysis for concurrent programs with nested locks. In *SAS 2010*, volume 6337 of *LNCS*, pages 253–270. Springer, 2010.
- [16] A. Farzan and P. Madhusudan. Causal dataflow analysis for concurrent programs. In *TACAS 2007*, volume 4424 of *LNCS*, pages 102–116, 2007.
- [17] C. Favre, D. Fahland, and H. Völzer. The relationship between workflow graphs and free-choice workflow nets. *Inf. Syst.*, 47:197–219, 2015.
- [18] C. Favre, H. Völzer, and P. 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.
- [19] P. Ganty and R. Majumdar. Algorithmic verification of asynchronous programs. *ACM Trans. Program. Lang. Syst.*, 34(1):6, 2012.
- [20] M. S. Hecht. *Flow Analysis of Computer Programs*. Elsevier Science Inc., 1977.
- [21] R. M. Karp and R. E. Miller. Parallel program schemata. *J. Comput. Syst. Sci.*, 3(2):147–195, 1969.
- [22] J. Katoen. GSPNs revisited: Simple semantics and new analysis algorithms. In *ACSD 2012*, pages 6–11. IEEE Computer Society, 2012.
- [23] P. Lammich and M. Müller-Olm. Conflict analysis of programs with procedures, dynamic thread creation, and monitors. In *SAS 2008*, *Valencia*, volume 5079 of *LNCS*, pages 205–220, 2008.
- [24] F. Nielson, H. R. Nielson, and C. Hankin. *Principles of program analysis*. Springer, 1999.
- [25] G. Ramalingam. Context-sensitive synchronization-sensitive analysis is undecidable. *ACM Trans. Program. Lang. Syst.*, 22(2):416–430, 2000.
- [26] M. D. Schwarz, H. Seidl, V. Vojdani, P. Lammich, and M. Müller-Olm. Static analysis of interrupt-driven programs synchronized via the priority ceiling protocol. In *POPL 2011*, pages 93–104. ACM, 2011.
- [27] H. Seidl and B. Steffen. Constraint-based inter-procedural analysis of parallel programs. In *ESOP 2000*, volume 1782 of *LNCS*, pages 351–365, 2000.
- [28] N. Trecka, W. M. P. van der Aalst, and N. Sidorova. Data-flow anti-patterns: Discovering data-flow errors in workflows. In *CAISE 2009*, volume 5565 of *LNCS*, pages 425–439, 2009.
- [29] W. M. P. van der Aalst. The application of Petri nets to workflow management. *J. Circuits, Syst. and Comput.*, 08(01):21–66, 1998.
- [30] W. M. P. van der Aalst. Workflow verification: Finding control-flow errors using Petri-net-based techniques. In *Business Process Management, Models, Techniques, and Empirical Studies*, volume 1806 of *LNCS*, pages 161–183. Springer, 2000.
- [31] W. M. P. van der Aalst and K. M. van Hee. *Workflow management: models, methods, and systems*. MIT press, 2004.
- [32] B. F. van Dongen, M. H. Jansen-Vullers, H. Verbeek, and W. M. 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.

## Proofs from Section 2.3

**Lemma 1.** *If  $C_1 \xrightarrow{w} C_2$  and  $v \equiv w$ , then  $C_1 \xrightarrow{v} C_2$ . In particular, if  $w$  is a (successful) run, then  $v$  is.*

**Proof:** It is easy to see that if  $\ell_1$  and  $\ell_2$  are independent and  $C \xrightarrow{\ell_1 \ell_2} C'$  for some configurations  $C, C'$ , then  $C \xrightarrow{\ell_2 \ell_1} C'$ . It follows that the same holds for any two Mazurkiewicz equivalent runs  $v, w$ . Indeed, if  $w$  is successful, then by definition  $C_{init} \xrightarrow{w} C_{fin}$ . Hence  $C_{init} \xrightarrow{v} C_{fin}$ , and so  $v$  is also successful.  $\square$

**Lemma 2.** *Let  $\mathcal{N}$  be a deterministic negotiation diagram and let  $S$  be a scheduler of  $\mathcal{N}$ . For every successful run  $w$  there is exactly one successful run  $v \equiv w$  that is compatible with  $S$ .*

**Proof:** Let  $w$  be a successful run, i.e.,  $C_{init} \xrightarrow{w} C_{fin}$ . Observe that, since  $\mathcal{N}$  is deterministic and the domain of  $n_{fin}$  contains all processes,  $n_{fin}$  is the only node enabled at  $C_{fin}$ .

We first prove that *at most* one run  $v \equiv w$  is compatible with  $S$ . Assume there are two different such runs  $v_1, v_2$ . Then  $v_1 = ul_1u_1$  and  $v_2 = ul_2u_2$  for some  $u, u_1, u_2$  and  $\ell_1 \neq \ell_2$ . Since  $v_1 \equiv v_2 \equiv w$ , the locations  $\ell_1$  and  $\ell_2$  are independent. But then  $S(u) \neq \ell_1$  or  $S(u) \neq \ell_2$ , and so at least one of the two runs is not compatible with  $S$ .

We now construct a run  $v \equiv w$  that is compatible with  $S$ . Suppose  $w \equiv v_1w_2$  and  $v_1$  is compatible with  $S$ . If  $w_2$  is not empty then we will show how to prolong  $v_1$  to an  $S$  compatible run  $v'_1$ , and shorten  $w_2$  to  $w'_2$  so that still  $w \equiv v'_1w'_2$ . Let  $S(v_1) = n$ . If we look at the run  $C_{init} \xrightarrow{v_1} C_1 \xrightarrow{w_2} C_{fin}$  then  $n$  is enabled in  $C_1$ . If  $n = n_{fin}$  then  $w_2$  is the empty sequence and we are done. Otherwise, since  $n$  is enabled in  $C_1$ ,  $w_2$  must be of the form  $ulu'$  where no location in  $u$  uses a process from  $dom(n)$ , and  $\ell = (n, b)$  for some outcome  $b$ . But then  $w_2 \equiv lluu'$ , and  $w \equiv v_1lluu'$ . We have found our desired  $v'_1 = v_1\ell$  and  $w'_2 = uu'$ .  $\square$

**Example:** Observe that Lemma 2 may not hold for runs that are not successful nor for non-deterministic negotiation diagrams. We explain this second point in more detail. Consider the diagram of Figure 5 with  $n_0$  and  $n_4$  as initial and final nodes (the hyperarcs indicate  $\delta(n_0, p_1, a) = \{n_1, n_2\}$  and  $\delta(n_0, p_2, a) = \{n_2, n_3\}$ ). Consider the scheduler that after the occurrence of  $(n_0, a)$  selects the node  $n_1$ . The successful trace  $[(n_0, a)(n_2, a)]$  does not have any run compatible with this scheduler.  $\square$

**Theorem 1.** *Let  $\mathcal{N}$  be a negotiation diagram, and let  $\llbracket \_ \rrbracket$  be an analysis framework for  $\mathcal{N}$ . If  $\mathcal{N}$  is deterministic and  $\llbracket \_ \rrbracket$  is Mazurkiewicz-invariant, then  $\llbracket \mathcal{N}, S \rrbracket = \llbracket \mathcal{N}, S' \rrbracket$  for every two schedulers  $S, S'$ , and so  $\llbracket \mathcal{N} \rrbracket = \llbracket \mathcal{N}, S \rrbracket$  for every scheduler  $S$ .*

**Proof:** By Lemma 2, there exists a bijection  $\phi$  between the successful runs compatible with  $S$  and  $S'$ : Given a successful run  $w$  compatible with  $S$ , we define  $\phi(w)$  as the unique run such that  $w \equiv \phi(w)$ . Since  $\llbracket \_ \rrbracket$  is Mazurkiewicz

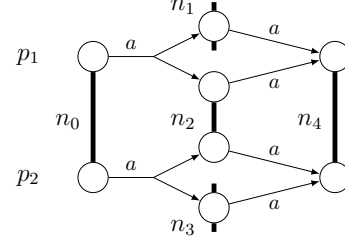


Figure 5. A non-deterministic negotiation diagram for which Lemma 2 does not hold.

invariant, we further have  $\llbracket w \rrbracket = \llbracket \phi(w) \rrbracket$ . Letting  $w$  and  $w'$  range over the successful runs of  $\mathcal{N}$  compatible with  $S$  and  $S'$ , respectively we obtain  $\llbracket \mathcal{N}, S \rrbracket = \llbracket \_ \rrbracket_w \llbracket w \rrbracket = \llbracket \_ \rrbracket_w \sum \phi(w) = \llbracket \_ \rrbracket_{w'} \llbracket w' \rrbracket = \llbracket \mathcal{N}, S' \rrbracket$ .  $\square$

## Proofs from Section 3

We recall some definitions and two results, one from [7] and one from [10]. A local path  $n_0 \xrightarrow{p_0, a_0} n_1 \cdots n_{k-1} \xrightarrow{p_{k-1}, a_{k-1}} n_k$  is *realizable* from a configuration  $C$  if there is a run  $C \xrightarrow{w}$  with  $w$  of the form  $(n_0, a_0)w_1(n_1, a_1) \cdots w_k(n_k, a_k)$ , such that  $p_i$  does not belong to the domain of any location of  $w_{i+1}$ .

**Lemma 13.** ([10]) *Let  $C$  be a reachable configuration of a sound deterministic negotiation diagram  $\mathcal{N}$ . Every local path whose initial node is enabled at  $C$  is realizable from  $C$ .*

**Lemma 14.** ([7]) *Let  $C$  be a reachable configuration of a sound deterministic negotiation diagram  $\mathcal{N}$  and let  $C \xrightarrow{w} C'$  where  $w = \ell_1 \cdots \ell_k$ ,  $k > 0$ . There is an index  $1 \leq i \leq k$  such that  $\ell_j \preceq \ell_i$  for every  $1 \leq j \leq k$ .*

No we are ready to prove domination lemma restated below.

**Lemma 3. (Domination Lemma)** *Let  $\mathcal{N}$  be a deterministic sound negotiation diagram. Every reachable local circuit of  $\mathcal{N}$  contains a dominant node, i.e. a node  $n$  such that  $m \preceq n$ , for every node  $m$  of the circuit.*

**Proof:** Let  $\pi$  be a reachable local circuit of  $\mathcal{N}$ . By Lemma 13, every number of iterations of  $\pi$  is realizable. If we take a sufficiently big number of iterations, say  $l$ , and a reachable configuration  $C_1$ , then we get an execution  $C_1 \xrightarrow{u} C \xrightarrow{v} C \xrightarrow{u'} C_2$  with  $uvu'$  realizing  $\pi^l$ . We have that the looping part  $v$  is of the form  $v = (n_0, a_0)w_1(n_1, a_1) \cdots w_k(n_k, a_k)$  where  $n_0 \xrightarrow{p_0, a_0} n_1 \xrightarrow{p_1, a_1} \dots \xrightarrow{p_{k-1}, a_{k-1}} n_k$  is a, continues, subsequence of  $\pi^l$ . Moreover realizability guarantees us that  $p_i$  does not belong to the domain of any location of  $w_{i+1}$ . By Lemma 14 some location  $\ell = (n, a)$  of  $v$  satisfies  $dom(\ell') \subseteq dom(\ell)$  for every location  $\ell'$  of  $v$ .

We claim that  $\ell = (n_j, a_j)$  for some  $1 \leq j \leq k$ , which implies that  $n_j$  is a dominating node of  $\pi$ . Let  $C'$  be a configuration reached during the execution of the loop  $C \xrightarrow{v}$

$C$ , that enables  $\ell$ . Let  $(n_i, a_i)$  be the last location of the path  $\pi$  occurring in the loop before  $C'$ . By the definition of  $\ell$  and of  $C'$  we have  $C'(p) = n$  for every  $p \in \text{dom}(n_i)$ , and so in particular  $C'(p_i) = n$ . Since  $p_i$  does not belong to the domain of any location of  $w_{i+1}$ , we also have  $C'(p_i) = n_{i+1}$ . So  $n = n_{i+1}$ , hence the claim is proved by taking  $j = i + 1$ .  $\square$

Next we prove the unique configuration lemma, that is the main technical tool for the theorem that follows.

**Lemma 4. (Unique Configuration Lemma)** *Let  $X \subseteq \text{Proc}$  be a set of processes. Let  $C_1, C_2$  be reachable configurations such that (1)  $C_1(p) = C_2(p)$  for every  $p \in X$ , and (2) every node  $n$  enabled at  $C_1$  or  $C_2$  satisfies  $\text{dom}(n) \cap X \neq \emptyset$ . Then  $C_1 = C_2$ .*

**Proof:** First we note that the assumptions about  $C_1, C_2$  imply the following, more general, property that we refer to as property II:

For every node  $m$  that is enabled either at  $C_1$  or at  $C_2$ , there is some process  $p \in \text{dom}(m)$  such that  $C_1(p) = C_2(p)$ .

The second observation is that for every pair of reachable configurations  $C_1, C_2$  the following holds:

If  $C_1, C_2$  satisfy property II, and  $C_1 \xrightarrow{w} C'_1$ ,  $C_2 \xrightarrow{w} C'_2$  are runs with the same labeling, then  $C'_1, C'_2$  satisfy II, too.

The above observation is easily seen: let us assume wlog. that  $n$  is enabled in  $C'_1$ . Either  $\text{dom}(n) \cap \text{dom}(w) = \emptyset$ , in which case  $n$  is enabled in  $C_1$ , and we apply the assumption on  $C_1, C_2$ . Or  $\text{dom}(n) \cap \text{dom}(w) \neq \emptyset$ , in which case we obtain  $C'_1(p) = C'_2(p)$  for any  $p \in \text{dom}(n) \cap \text{dom}(w)$ .

We show now the following claim:

If  $\mathcal{N}$  is sound and  $C_1, C_2$  are reachable configurations satisfying property II, then  $C_1 = C_2$ .

Let us assume by contradiction that  $C_1 \neq C_2$  are reachable configurations satisfying property II. By soundness, there is some run  $C_1 \xrightarrow{u} C_{fin}$ . Since  $C_2 \neq C_1$ , we cannot execute  $u$  entirely from  $C_2$ . Let  $v$  be the maximal trace prefix of  $u$  that can be executed from  $C_2$ . Let  $C'_1, C'_2$  and  $m_1$  be such that  $C_1 \xrightarrow{v} C'_1$ ,  $C_2 \xrightarrow{v} C'_2$  and  $m_1$  is a node enabled in  $C'_1$  but not in  $C'_2$ . Since  $C'_1, C'_2$  also satisfy property II, we know that there is some process  $p_1 \in \text{dom}(m_1)$  such that  $C'_2(p_1) = C'_1(p_1) (= m_1)$ . This implies, again by soundness, that from  $C'_2$ , node  $m_1$  must be executed at some point, although it is not enabled in  $C'_2$ . So there exists some node  $m' \neq m_1$  with  $\text{dom}(m') \cap \text{dom}(m_1) \neq \emptyset$  that must be executed before  $m_1$  from  $C'_2$ . Equivalently, there exists an execution of the form

$$C'_2 \xrightarrow{u'} C_{fin}, \quad \text{where } u' = w_1 m' w_2 m_1 w_3$$

for some sequences  $w_1, w_2, w_3$  such that  $\text{dom}(m_1) \cap \text{dom}(w_1) = \emptyset$ . Let us assume that  $u'$  is such that  $w_1$  is the shortest possible. This implies that all nodes in  $w_1$  are dependent, possibly indirectly on  $m'$ . In other words, there is a local path from every node in  $w_1$  to  $m'$ .

Let  $v'$  be the maximal prefix of  $u'$  that can be executed from  $C'_1$ . Note that  $v'$  must be a prefix of  $w_1$  since for some process  $p'$  of  $\text{dom}(m_1)$  we have  $C'_1(p') = m_1$ , but  $C'_2(p') = m'$ .

Let  $C''_1, C''_2$  be such that  $C'_1 \xrightarrow{v'} C''_1$  and  $C'_2 \xrightarrow{v'} C''_2$  - in particular,  $C''_1, C''_2$  satisfy property II. Let  $m_2$  be the first node after  $v'$  on  $u'$ ; it may be  $m'$  or a node from  $w_1$ . By the choice of  $v'$ , node  $m_2$  that is enabled in  $C''_2$ , but not in  $C''_1$ . Observe also that there is a local path from  $m_2$  to  $m_1$ , since there is a local path from  $m_2$  to  $m'$  and  $\text{dom}(m') \cap \text{dom}(m_1) \neq \emptyset$ . Note moreover that  $C''_1(p_1) = C''_2(p_1) = m_1$  and that process  $p_1$  cannot occur on the path from  $m_2$  to  $m_1$  (except at  $m_1$ ).

Let us summarize: from the pair  $C_1, C_2$  we obtained the pair  $C'_1, C'_2$ , node  $m_1$  and process  $p_1 \in \text{dom}(m_1)$ . From  $C'_1, C'_2$  we obtained  $C''_1, C''_2$ , and node  $m_2$  such that there is a path from  $m_2$  to  $m_1$  without process  $p_1$  (except at  $m_1$ ). All these pairs satisfy property II. Iterating this argument we build a sequence of nodes  $m_1, \dots, m_k$ , and processes  $p_i \in \text{dom}(m_i)$  such that for every  $i$  there is a path from  $m_{i+1}$  to  $m_i$  without process  $p_i$  (except at  $m_i$ ). Since the number of nodes is finite we must find ultimately a cycle, so wlog.  $m_k = m_1$ . Note that such a cycle cannot contain a dominating node, so it violates Lemma 3, contradiction. Therefore,  $C_1 = C_2$  and we showed thus the claim.  $\square$

**Theorem 2.** *Let  $m$  be a reachable node of a sound deterministic negotiation diagram  $\mathcal{N}$ .*

- (i) *There is a unique reachable configuration  $I(m)$  of  $\mathcal{N}$  that enables  $m$ , and no other node.*
- (ii) *There is a unique configuration  $F(m)$  such that*
  - *$F(m)$  is reachable from  $I(m)$  by means of an  $m$ -sequence, and*
  - *for every node  $n$  enabled at  $F(m)$ ,  $\text{dom}(n)$  is not included in  $\text{dom}(m)$ .*
- (iii) *For every location  $\ell$  of  $m$  there is a unique configuration  $F(\ell)$  such that*
  - *$F(\ell)$  is reachable from  $I(m)$  by means of a strict  $m$ -sequence starting with  $\ell$ , and*
  - *for every node  $n$  enabled at  $F(\ell)$ ,  $\text{dom}(n)$  is not strictly included in  $\text{dom}(m)$ .*

**Proof:** For (i), take  $X = \text{dom}(m)$ . Suppose that there are two configurations  $I_1$  and  $I_2$  as in (i). The hypotheses of Lemma 4 are satisfied, and so  $I_1 = I_2$ .

For (ii) take  $X = \text{Proc} \setminus \text{dom}(m)$ . Suppose that there are two configurations  $F_1$  and  $F_2$  as in (ii). We get that the two configurations agree on  $X$  since they are obtained from  $I(m)$  by  $m$ -sequences. By definition of the two  $F$  configurations, every node enabled in one of  $F_1, F_2$  needs a process from  $X$ , so  $F_1 = F_2$  by Lemma 4.

For (iii) we take  $m$  and  $\ell$  as in the statement. We prove a stronger property by induction. Let  $C$  be a reachable configuration of  $\mathcal{N}$ , and  $v, w$  two strict  $m$ -sequences leading to  $F_v(\ell)$  and  $F_w(\ell)$ , respectively, satisfying the two conditions of (iii). We want to show that  $F_v(\ell) = F_w(\ell)$ . The proof

is by induction on the sum of the lengths of  $v$  and  $w$ . The conclusion then follows by taking as  $C$  the configuration obtained from  $I(m)$  after doing  $\ell$ .

For the induction step, let us take the first location  $(n, a)$  of  $v$ , i.e.  $v = (n, a)v'$ . Now  $w$  must be Mazurkiewicz equivalent to  $(n, a)w'$ , since  $n$  is enabled in  $C$ . Thanks to Lemma 1, we obtain two computations  $C \xrightarrow{(n, a)} C' \xrightarrow{v'} F_v(\ell)$ , and  $C \xrightarrow{(n, a)} C' \xrightarrow{w'} F_w(\ell)$ . By induction hypothesis we have  $F_v(\ell) = F_w(\ell)$ .  $\square$

Before proving that  $\mathcal{N}|_n$  is sound we need an important technical lemma.

**Lemma 15.** *If  $m = F(n)(p)$  for some process  $p$  then there is no reachable configuration of  $\mathcal{N}|_n$  where  $m$  is enabled.*

**Proof:** Suppose by contradiction that  $m$  is executed in  $\mathcal{N}|_n$  and  $F(n)(p) = m$  for some process  $p$ . So we can find an  $n$ -sequence  $u$  such that

$$I(m) \xrightarrow{u} F(n)$$

is an execution of  $\mathcal{N}$ . We choose  $u$  minimal, so that  $m$  is executed only once, at the beginning of  $u$ . Since in  $F(n)$  no node  $m' \preceq n$  is enabled, we have  $F(n)(p_1) = m_1 \neq m$  for some  $p_1 \in \text{dom}(m)$ .

If  $m_1 \preceq n$ , then since  $m_1$  is not enabled in  $F(n)$  we have  $F(n)(p_2) = m_2$  for some  $p_2 \in \text{dom}(m_1)$ . By induction we get a sequence of processes  $p_1, \dots, p_k$  and of nodes  $m_1, \dots, m_k$  such that:

- $F(n)(p_i) = m_i$  and  $p_i \in \text{dom}(m_{i-1})$  for all  $i$ ,
- $m_i \preceq n$  for all  $i < k$ , and  $m_k \not\preceq n$ .

Intuitively, each  $p_i$  has to wait for  $p_{i+1}$  in order to execute  $m_i$ . Since  $\mathcal{N}$  is sound there is some execution from  $F(n)$  that enables  $m$ . Let us consider such an execution

$$I(m) \xrightarrow{u} F(n) \xrightarrow{u_k} C_k \xrightarrow{u_{k-1}} \dots \xrightarrow{u_1} C_1 \xrightarrow{u_0} C$$

such that

- $C \models m$ ,
- for all  $i \leq k$ ,  $u_{i-1}$  starts with  $\ell_i = (m_i, a_i)$  for some  $a_i$ ,
- process  $p$  does not occur in  $u_k \dots u_1 u_0$ .

Recall that  $u$  starts by  $(m, a)$  for some  $a$ . The above execution yields some local path  $\pi$  from  $m$  to  $m$  containing  $m_k$ . This path should have a dominant node by Lemma 3. Since  $m_k \not\preceq n$  the set of processes occurring in  $\pi$  is not included in  $\text{dom}(n)$ . So the dominant node cannot be a part of  $u$  since the domains of nodes in  $u$  are included in  $\text{dom}(n)$ . The dominant node cannot be a part of  $u_k \dots u_0$  either since process  $p$  does not occur in  $u_k \dots u_0$ . We obtain thus a contradiction.  $\square$

**Lemma 5.** *If  $\mathcal{N}$  is a sound deterministic negotiation diagram then so is  $\mathcal{N}|_n$ .*

**Proof:** Consider a run  $C_{init}|_n \xrightarrow{v} C_1$  in  $\mathcal{N}|_n$ , where  $C_{init}|_n$  is the initial configuration of  $\mathcal{N}|_n$ , i.e.,  $C_{init}|_n(p) = \{n\}$  for

every  $p \in \text{dom}(n)$ . We prove that this run can be extended to as successful run of  $\mathcal{N}|_n$ .

Since  $n$  is reachable in  $\mathcal{N}$ , by Theorem 2(i) we can take a run  $C_{init} \xrightarrow{u} I(n)$  and prolong it to  $C_{init} \xrightarrow{u} I(n) \xrightarrow{v} C'_1$  such that (i)  $C'_1(p) = I(n)(p)$  for every process  $p \notin \text{dom}(n)$ , and (ii)  $C'_1(p) = C_1(p)$  or  $(C_1(p) = n|_n^{fin}$  and  $C'_1(p) = F(n))$  for every process  $p \in \text{dom}(n)$ . Since  $\mathcal{N}$  is sound, we can prolong the run further to an accepting one, say  $C_{init} \xrightarrow{u} I(n) \xrightarrow{v} C'_1 \xrightarrow{w} C_{fin}$ . We now permute exhaustively consecutive independent outcomes  $\ell \ell'$  in  $w$  such that  $\text{dom}(\ell) \subseteq \text{dom}(n)$  and  $\text{dom}(\ell')$  is not included in  $\text{dom}(n)$ ; say the result is  $w_1 w_2$  with  $\text{dom}(w_1) \subseteq \text{dom}(n)$ . Then we have  $C_{init} \xrightarrow{u} I(n) \xrightarrow{v} C'_1 \xrightarrow{w_1} C'_2$ . Lemma 15 gives us then an execution  $C_{init}|_n \xrightarrow{v} C_1 \xrightarrow{w_1} C_2$  for configurations  $C_2$  and  $C'_2$  satisfying the conditions (i) and (ii) above (that is,  $C_2$  and  $C'_2$  satisfy the same two conditions as  $C_1$  and  $C'_1$ ). Moreover, since the outcomes are permuted exhaustively, either  $C'_2 = C_{fin}$ , or every node enabled in  $C'_2$  needs a process outside of  $\text{dom}(n)$ . By Theorem 2(ii), in both cases we have  $C'_2 = F(n)$ . By condition (ii) above, we have  $C_2(p) = n|_n^{fin}$  for every process  $p \in \text{dom}(n)$ . So  $C_2$  is the final configuration of  $\mathcal{N}|_n$ , and the run can be prolonged to a successful run.  $\square$

**Lemma 6.** *If  $\mathcal{N}$  is a sound deterministic negotiation diagram, then so is  $\mathcal{N}|_\ell$ .*

**Proof:** Analogous to the proof of Lemma 5, replacing Theorem 2(ii) by Theorem 2(iii).  $\square$

## Proofs from Section 4

**Lemma 7.** *Let  $\mathcal{N}$  and  $\ell = (n, a)$  be as in Definition 6. Assign to the new location  $\ell' = (n, a_\ell)$  the mapping  $\llbracket \ell' \rrbracket := \llbracket \mathcal{N}|_\ell \rrbracket$ . Then  $\llbracket \mathcal{N} \rrbracket = \llbracket \text{Red}_\ell(\mathcal{N}) \rrbracket$ . Analogously,  $\llbracket \mathcal{N} \rrbracket = \llbracket \text{Red}_n(\mathcal{N}) \rrbracket$  when we assign  $\llbracket (n, a_n) \rrbracket = \llbracket \mathcal{N}|_n \rrbracket$ .*

**Proof:** We will consider only the first statement. The proof of the second is analogous.

Since  $\mathcal{N}$  is deterministic and the framework is Mazurkiewicz-invariant, we have  $\llbracket \mathcal{N} \rrbracket = \llbracket \mathcal{N}, S \rrbracket$  for every scheduler  $S$ . Let  $S$  be the scheduler that gives priority to nodes outside  $\mathcal{N}|_\ell$  over nodes of  $\mathcal{N}|_\ell$ . By Theorem 1, every successful run  $C_{init} \xrightarrow{w} C_{fin}$  of  $\mathcal{N}$  compatible with  $S$  can be split into  $w = w_0 u_1 w_1 \dots u_k w_k$  such that  $C_{init} \xrightarrow{w_0} I(\ell)$ ,  $I(\ell) \xrightarrow{u_i} F(\ell) \xrightarrow{w_i} I(\ell)$  for every  $1 \leq i \leq k-1$ , and  $I(\ell) \xrightarrow{u_k} F(\ell) \xrightarrow{w_k} C_{fin}$ , where  $u_1, \dots, u_k$  are successful runs of  $\mathcal{N}|_\ell$ , and  $w_k$  does not contain  $\ell$ . Let  $\mathcal{R}(w_1, \dots, w_k)$  stand for the set of all successful runs of this form for some fixed  $w_1, \dots, w_k$ . By distributivity of  $\llbracket \cdot \rrbracket$ , the total contribution of  $\mathcal{R}(w_1, \dots, w_k)$  to  $\llbracket \mathcal{N} \rrbracket$  is  $\llbracket \mathcal{R}(w_1, \dots, w_k) \rrbracket = \llbracket w_0 \rrbracket \circ \llbracket \mathcal{N}|_\ell \rrbracket \circ \llbracket w_1 \rrbracket \circ \dots \circ \llbracket \mathcal{N}|_\ell \rrbracket \circ \llbracket w_k \rrbracket$ .

Consider now a successful run  $w'$  of  $\text{Red}_\ell(\mathcal{N})$ . It is of the form  $w_0 \ell' w_1 \ell' w_2 \dots \ell' w_k$  with  $w_k$  not containing  $\ell' = (n, a_\ell)$ . Since  $\llbracket \ell' \rrbracket = \llbracket \mathcal{N}|_\ell \rrbracket$ , the contribution of  $w'$  to  $\llbracket \text{Red}_\ell(\mathcal{N}) \rrbracket$  is  $\llbracket w' \rrbracket = \llbracket \mathcal{R}(w_1, \dots, w_n) \rrbracket$ . Abusing language, let us write  $\mathcal{R}(w')$  instead of  $\mathcal{R}(w_1, \dots, w_n)$ .

Letting  $w'$  range over the successful runs of  $Red_\ell(\mathcal{N})$ , we get  $\llbracket Red_\ell(\mathcal{N}) \rrbracket = \sqcap_{w'} \llbracket w' \rrbracket = \sqcap_{w'} \llbracket \mathcal{R}(w') \rrbracket = \llbracket \mathcal{N} \rrbracket$ .  $\square$

**Lemma 8.** *Let  $\mathcal{N}$  be a sound deterministic negotiation diagram, and  $n$  a node such that all nodes  $m \prec n$  are reduced in  $\mathcal{N}$ .*

- (1) *if  $a$  is an outcome of  $n$ , then for  $\ell = (n, a)$  the negotiation diagram  $\mathcal{N}|_\ell$  is a one-trace negotiation.*
- (2) *if all locations  $\ell' \preceq n$  are reduced, then  $\mathcal{N}|_n$  is a replication.*

**Proof:** (1) Let  $n$  and  $\ell$  be as in the assumption. By the definition of  $\mathcal{N}|_\ell$ , all nodes  $m \neq n$  of  $\mathcal{N}|_\ell$  (except for the final node) have a smaller domain than  $n$  and so, by the minimality of  $n$ , are reduced. In particular, they all have a single outcome. Since  $n$  also has one single outcome in  $\mathcal{N}|_\ell$ , namely  $\ell$ , every node of  $\mathcal{N}|_\ell$  has one single outcome. Finally, we observe that  $\mathcal{N}|_\ell$  is acyclic: any circuit  $\mathcal{N}|_\ell$  would contain a dominant node  $n' \neq n$  by Lemmas 6, 3. Since  $n'$  is reduced we have  $\delta(n', a, p) = F(n')(p)$  for every  $p \in dom(n')$ . So it cannot be the case that  $n'$  is dominant, by the definition of  $F(n')$ .

(2) Let  $C_{init}|_n \xrightarrow{\ell_1 \dots \ell_k} C$  be an arbitrary run of  $\mathcal{N}|_n$ , where  $\ell_i = (n_i, a_i)$  for every  $1 \leq i \leq k$ . We prove that for every  $1 \leq i \leq k-1$ , the outcome  $\ell_i$  satisfies  $\delta|_n(n_i, a, p) = n_{i+1}$  for every process  $p \in dom(n)$ , which implies that  $\mathcal{N}|_n$  is a replication.

Assume that the above property does not hold, and let  $\ell_i$  be the first location that does not satisfy the property. By the definition of  $\mathcal{N}|_n$  we have  $dom(n_j) \subseteq dom(n)$  for every  $1 \leq j \leq k$ . Since  $n_{i+1}$  is enabled after the occurrence of  $\ell_i$ , we have  $\delta(n_i, a_i, p) = n_{i+1}$  for every process  $p \in dom(n_{i+1})$ . Since  $\ell_i$  does not satisfy the property, we have  $dom(n_{i+1}) \subset dom(n_i)$ . But since  $\ell_i$  is reduced we have  $\delta(n_i, a_i, p) = F(\ell_i)(p) = n_{i+1}$ . By definition of  $F(\ell_i)$  it cannot be that  $n_{i+1}$  is enabled after doing  $\ell_i$ . A contradiction.  $\square$