

# On Negotiation as Concurrency Primitive

Javier Esparza<sup>1</sup> and Jörg Desel<sup>2</sup>

<sup>1</sup> Fakultät für Informatik, Technische Universität München, Germany

<sup>2</sup> Fakultät für Mathematik und Informatik, FernUniversität in Hagen, Germany

**Abstract.** We introduce negotiations, a model of concurrency close to Petri nets, with multiparty negotiation as primitive. We study the problems of soundness of negotiations and of, given a negotiation with possibly many steps, computing a *summary*, i.e., an equivalent one-step negotiation. We provide a complete set of reduction rules for sound, acyclic, weakly deterministic negotiations and show that, for deterministic negotiations, the rules compute the summary in polynomial time.

## 1 Introduction

Many modern distributed systems consist of components whose behavior is only partially known. Typical examples include open systems where programs (e.g. Java applets) can enter or leave, multi-agent systems, business processes, or protocols for conducting elections and auctions. An interaction between a fixed set of components with not fully known behavior can be abstractly described as a *negotiation* in which several *parties* (the components involved in the negotiation) nondeterministically agree on an *outcome*, which results in a transformation of internal states of the parties. A more technical but less suggestive term would be a *synchronized nondeterministic choice* and, as the name suggests, these interactions can be modelled in any standard process algebra as a combination of parallel composition and nondeterministic choice, or as small Petri nets. We argue that much can be gained by studying formal models with *negotiation atoms* as concurrency primitive. In particular, we show that the negotiation point of view reveals new classes of systems with polynomial analysis algorithms.

Negotiation atoms can be combined into *distributed negotiations*. For instance, a distributed negotiation between a buyer, a seller, and a broker, consists of one or more rounds of atoms involving the buyer and the broker or the seller and the broker, followed by a final atom between the buyer and the seller. We introduce a formal model for distributed negotiations, close to a colored version of van der Aalst's *workflow nets* [1], and investigate two important analysis problems. First, just like workflow nets, distributed negotiations can be *unsound* because of deadlocks or livelocks (states from which no deadlock is reached, but the negotiation cannot be completed). The *soundness* problem consists of deciding if a given negotiation is sound. Second, a sound negotiation is equivalent to a negotiation with only one atom whose state transformation function determines the possible final internal states of all parties as a function of their initial internal states. We call this negotiation a *summary*. The *summarization problem* consists

of computing a summary of a distributed negotiation. Both problems will be shown to be PSPACE-hard for arbitrary negotiations, and NP-hard for acyclic ones. They can be solved by means of well-known algorithms based on the exhaustive exploration of the state space. However, this approach badly suffers from the state-explosion problem: even the analysis of distributed negotiations with a very simple structure requires exponential time.

In this paper we suggest *reduction* algorithms that avoid the construction of the state space but exhaustively apply syntactic reduction rules that simplify the system while preserving some aspects of the behavior, like absence of deadlocks. This approach has been extensively applied to Petri nets or workflow nets, but most of this work has been devoted to the liveness or soundness problems [5,15,16,13,22]. For these problems many reduction rules are known, and some sets of rules have been proved *complete* for certain classes of systems [14,10,11], meaning that they reduce all live or sound systems in the class, and only those, to a trivial system (in our case to a single atomic negotiation). However, many of these rules, like the linear dependency rule of [11], cannot be applied to the summarization problem, because they preserve only the soundness property.

We present a complete set of reduction rules for the summarization problem of *acyclic* negotiations that are either *deterministic* or *weakly deterministic*. The rules are inspired by reduction rules used to transform finite automata into regular expressions by eliminating states [18]. In deterministic negotiations all involved agents are deterministic, meaning that they are never ready to engage in more than one atomic negotiation. Intuitively, nondeterministic agents may be ready to engage in multiple atomic negotiations, and which one takes place is decided by the deterministic parties, which play thus the role of negotiation leaders. In weakly deterministic negotiations not every agent is deterministic, but some deterministic party is involved in every atomic negotiation an agent can engage in next. For *deterministic* negotiations we prove that a sound and acyclic negotiation can be summarized by means of a polynomial number of application of the rules, leading to a polynomial algorithm.

The paper is organized as follows. Section 2 introduces the syntax and semantics of the model. Section 3 introduces the soundness and summarization problems. Section 4 presents our reduction rules. Section 5 defines (weakly) deterministic negotiations. Section 6 proves the completeness and polynomial complexity results announced above. Finally, Section 7 presents some conclusions, open questions and related work. The paper only contains proof sketches. Full proofs can be found in [12].

## 2 Negotiations: Syntax and Semantics

We fix a finite set  $A$  of *agents* representing potential parties of negotiations. Each agent  $a \in A$  has a (possibly infinite) nonempty set  $Q_a$  of *internal states*. We denote by  $Q_A$  the cartesian product  $\prod_{a \in A} Q_a$ . A *transformer* is a left-total relation  $\tau \subseteq Q_A \times Q_A$ , representing a nondeterministic state transforming function. Given  $S \subseteq A$ , we say that a transformer  $\tau$  is an *S-transformer* if, for each

$a_i \notin S$ ,  $\left( (q_{a_1}, \dots, q_{a_i}, \dots, q_{a_{|A|}}), (q'_{a_1}, \dots, q'_{a_i}, \dots, q'_{a_{|A|}}) \right) \in \tau$  implies  $q_{a_i} = q'_{a_i}$ . So an  $S$ -transformer only transforms the internal states of agents in  $S$ .

**Definition 1.** A negotiation atom, or just an atom, is a triple  $n = (P_n, R_n, \delta_n)$ , where  $P_n \subseteq A$  is a nonempty set of parties,  $R_n$  is a finite, nonempty set of outcomes, and  $\delta_n$  is a mapping assigning to each outcome  $r$  in  $R_n$  a  $P_n$ -transformer  $\delta_n(r)$ . We denote the transformer  $\delta_n(r)$  by  $\langle n, r \rangle$ , and, if there is no confusion, by  $\langle r \rangle$ .

Intuitively, if the states of the agents before a negotiation  $n$  are given by a tuple  $q$  and the outcome of the negotiation is  $r$ , then the agents change their states to  $q'$  for some  $(q, q') \in \langle n, r \rangle$ . Only the parties of  $n$  can change their internal states. Each outcome  $r \in R_n$  is possible, independent from the previous internal states of the parties.

For a simple example, consider a negotiation atom  $n_{\text{FD}}$  with parties **F** (Father) and **D** (teenage Daughter). The goal of the negotiation is to determine whether **D** can go to a party, and the time at which she must return home. The possible outcomes are  $\{\text{yes}, \text{no}, \text{ask\_mother}\}$ . Both sets  $Q_{\text{F}}$  and  $Q_{\text{D}}$  contain a state *angry* plus a state  $t$  for every time  $T_1 \leq t \leq T_2$  in a given interval  $[T_1, T_2]$ . Initially, **F** is in state  $t_f$  and **D** in state  $t_d$ . The transformer  $\delta_{n_{\text{FD}}}$  is given by

$$\begin{aligned} \langle \text{yes} \rangle &= \{((t_f, t_d), (t, t)) \mid t_f \leq t \leq t_d \vee t_d \leq t \leq t_f\} \\ \langle \text{no} \rangle &= \{((t_f, t_d), (\text{angry}, \text{angry}))\} \\ \langle \text{ask\_mother} \rangle &= \{((t_f, t_d), (t_f, t_d))\} \end{aligned}$$

That is, if the outcome is **yes**, then **F** and **D** agree on a time  $t$  which is not earlier and not later than both suggested times. If it is **no**, then there is a quarrel and both parties get angry. If the outcome is **ask\_mother**, then the parties keep their previous times.

## 2.1 Combining Atomic Negotiations

If the outcome of the atom above is **ask\_mother**, then  $n_{\text{FD}}$  should be followed by a second atom  $n_{\text{DM}}$  between **D** and **M** (Mother). The complete negotiation is the composition of  $n_{\text{FD}}$  and  $n_{\text{DM}}$ , where the possible internal states of **M** are the same as those of **F** and **D**, and  $n_{\text{DM}}$  is a copy of  $n_{\text{FD}}$ , but without the **ask\_mother** outcome. In order to compose atoms, we add a *transition function*  $\mathcal{X}$  that assigns to every triple  $(n, a, r)$  consisting of an atom  $n$ , a participant  $a$  of  $n$ , and an outcome  $r$  of  $n$  a set  $\mathcal{X}(n, a, r)$  of atoms. Intuitively, this is the set of atomic negotiations agent  $a$  is ready to engage in after the atom  $n$ , if the outcome of  $n$  is  $r$ .

**Definition 2.** Given a finite set of atoms  $N$ , let  $T(N)$  denote the set of triples  $(n, a, r)$  such that  $n \in N$ ,  $a \in P_n$ , and  $r \in R_n$ . A negotiation is a tuple  $\mathcal{N} = (N, n_0, n_f, \mathcal{X})$ , where  $n_0, n_f \in N$  are the initial and final atoms, and  $\mathcal{X}: T(N) \rightarrow 2^N$  is the transition function. Further,  $\mathcal{N}$  satisfies the following properties:

- (1) every agent of  $A$  participates in both  $n_0$  and  $n_f$ ;
- (2) for every  $(n, a, r) \in T(N)$ :  $\mathcal{X}(n, a, r) = \emptyset$  iff  $n = n_f$ .




Fig. 1. An acyclic negotiation and the ping-pong negotiation

We may have  $n_0 = n_f$ . Notice that  $n_f$  has, as all other atoms, at least one outcome  $r \in R_{n_f}$ .

Negotiations are graphically represented as shown in Figure 1. For each atom  $n \in N$  we draw a black bar; for each party  $a$  of  $P_n$  we draw a white circle on the bar, called a *port*. For each  $(n, a, r) \in T(N)$ , we draw a hyperarc leading from the port of  $a$  in  $n$  to all the ports of  $a$  in the atoms of  $\mathcal{X}(n, a, r)$ , and label it by  $r$ . Figure 1 shows on the left the graphical representation of the Father-Daughter-Mother negotiation sketched above. Instead of multiple (hyper)arcs connecting the same input port to the same output ports we draw a single (hyper)arc with multiple labels. In the figure, we write **y** for **yes**, **n** for **no**, and **am** for **ask\_mother**. **st** stands for **start**, the only outcome of  $n_0$ . Since  $n_f$  has no outgoing arc, the outcomes of  $n_f$  do not appear in the graphical representation.

**Definition 3.** *The graph associated to a negotiation  $\mathcal{N} = (N, n_0, n_f, \mathcal{X})$  is the directed graph with vertices  $N$  and edges  $\{(n, n') \in N \times N \mid \exists (n, a, r) \in T(N) : n' \in \mathcal{X}(n, a, r)\}$ . The negotiation  $\mathcal{N}$  is acyclic if its graph has no cycles.*

The negotiation on the left of Figure 1 is acyclic. The negotiation on the right (ignore the black dots on the arcs for the moment) is the ping-pong negotiation, well-known in every family. The  $n_{DM}$  atom has now an extra outcome **ask\_father** (**af**), and Daughter D can be sent back and forth between Mother and Father. After each round, D “negotiates with herself” (atom  $n_D$ ) with possible outcomes **c** (**continue**) and **gu** (**give up**). This negotiation is cyclic because, for instance, we have  $\mathcal{X}(n_{FD}, D, \mathbf{am}) = \{n_{DM}\}$ ,  $\mathcal{X}(n_{DM}, D, \mathbf{af}) = \{n_D\}$ , and  $\mathcal{X}(n_D, D, \mathbf{c}) = \{n_{FD}\}$ .

### 2.2 Semantics

A *marking* of a negotiation  $\mathcal{N} = (N, n_0, n_f, \mathcal{X})$  is a mapping  $\mathbf{x} : A \rightarrow 2^N$ . Intuitively,  $\mathbf{x}(a)$  is the set of atoms that agent  $a$  is currently ready to engage in next. The *initial* and *final* markings, denoted by  $\mathbf{x}_0$  and  $\mathbf{x}_f$  respectively, are given by  $\mathbf{x}_0(a) = \{n_0\}$  and  $\mathbf{x}_f(a) = \emptyset$  for every  $a \in A$ .

A marking  $\mathbf{x}$  enables an atom  $n$  if  $n \in \mathbf{x}(a)$  for every  $a \in P_n$ , i.e., if every agent that parties in  $n$  is currently ready to engage in it. If  $\mathbf{x}$  enables  $n$ , then  $n$  can take place and its parties agree on an outcome  $r$ ; we say that  $(n, r)$  occurs. The occurrence of  $(n, r)$  produces a next marking  $\mathbf{x}'$  given by  $\mathbf{x}'(a) = \mathcal{X}(n, a, r)$  for every  $a \in P_n$ , and  $\mathbf{x}'(a) = \mathbf{x}(a)$  for every  $a \in A \setminus P_n$ . We write  $\mathbf{x} \xrightarrow{(n,r)} \mathbf{x}'$  to denote this, and call it a *small step*.

By this definition,  $\mathbf{x}(a)$  is always either  $\{n_0\}$  or equals  $\mathcal{X}(n, a, r)$  for some atom  $n$  and outcome  $r$ . The marking  $\mathbf{x}_f$  can only be reached by the occurrence of  $(n_f, r)$  ( $r$  being a possible outcome of  $n_f$ ), and it does not enable any atom. Any other marking that does not enable any atom is considered a *deadlock*.

Reachable markings can be graphically represented by placing tokens (black dots) on the forking points of the hyperarcs (or in the middle of an arc). Thus, both the initial marking and the final marking are represented by no tokens, and all other reachable markings are represented by exactly one token per agent.

Figure 1 shows on the right the marking in which Father (F) is ready to engage in the atoms  $n_{FD}$  and  $n_f$ , Daughter (D) is only ready to engage in  $n_{FD}$ , and Mother (M) is ready to engage in both  $n_{DM}$  and  $n_f$ .

We write  $\mathbf{x}_1 \xrightarrow{\sigma}$  to denote that there is a sequence

$$\mathbf{x}_1 \xrightarrow{(n_1, r_1)} \mathbf{x}_2 \xrightarrow{(n_2, r_2)} \dots \xrightarrow{(n_{k-1}, r_{k-1})} \mathbf{x}_k \xrightarrow{(n_k, r_k)} \mathbf{x}_{k+1} \dots$$

of small steps such that  $\sigma = (n_1, r_1) \dots (n_k, r_k) \dots$ . If  $\mathbf{x}_1 \xrightarrow{\sigma}$ , then  $\sigma$  is an *occurrence sequence* from the marking  $\mathbf{x}_1$ , and  $\mathbf{x}_1$  enables  $\sigma$ . If  $\sigma$  is finite, then we write  $\mathbf{x}_1 \xrightarrow{\sigma} \mathbf{x}_{k+1}$  and say that  $\mathbf{x}_{k+1}$  is *reachable* from  $\mathbf{x}_1$ . If  $\mathbf{x}_1$  is the initial marking then we call  $\sigma$  *initial occurrence sequence*. If moreover  $\mathbf{x}_{k+1}$  is the final marking, then  $\sigma$  is a *large step*.

### 3 Analysis Problems

Correct negotiations should be deadlock-free and, in principle, they should not have infinite occurrence sequences either. However, requiring the latter in our negotiation model is too strong, because infinite occurrence sequences may be excluded by fairness constraints. Following [1,2], we introduce a notion of partial correctness independent of termination:

**Definition 4.** A negotiation is sound if (a) every atom is enabled at some reachable marking, and (b) every occurrence sequence from the initial marking is either a large step or can be extended to a large step.

The negotiations of Figure 1 are sound. However, if we set  $\mathcal{X}(n_0, \mathbf{M}, \mathbf{st}) = \{n_{DM}\}$  instead of  $\mathcal{X}(n_0, \mathbf{M}, \mathbf{st}) = \{n_{DM}, n_f\}$ , then the occurrence sequence  $(n_0, \mathbf{st})(n_{FD}, \mathbf{yes})$  leads to a deadlock.

The *final outcomes* of a negotiation are the outcomes of its final atom. Intuitively, two sound negotiations over the same agents are equivalent if they have the same final outcomes, and for each final outcome they transform the same initial states into the same final states.

**Definition 5.** Given a negotiation  $\mathcal{N} = (N, n_0, n_f, \mathcal{X})$ , we attach to each outcome  $r$  of  $n_f$  a summary transformer  $\langle \mathcal{N}, r \rangle$  as follows. Let  $E_r$  be the set of large steps of  $\mathcal{N}$  that end with  $(n_f, r)$ . We define  $\langle \mathcal{N}, r \rangle = \bigcup_{\sigma \in E_r} \langle \sigma \rangle$ , where for  $\sigma = (n_1, r_1) \dots (n_k, r_k)$  we define  $\langle \sigma \rangle = \langle n_1, r_1 \rangle \cdots \langle n_k, r_k \rangle$  (each  $\langle n_i, r_i \rangle$  is a relation on  $Q_A$ ; concatenation is the usual concatenation of relations).

$\langle \mathcal{N}, r \rangle(q_0)$  is the set of possible final states of the agents after the negotiation concludes with outcome  $r$ , if their initial states are given by  $q_0$ .

**Definition 6.** Two negotiations  $\mathcal{N}_1$  and  $\mathcal{N}_2$  over the same set of agents are equivalent, denoted by  $\mathcal{N}_1 \equiv \mathcal{N}_2$ , if they are either both unsound, or if they are both sound, have the same final outcomes, and  $\langle \mathcal{N}_1, r \rangle = \langle \mathcal{N}_2, r \rangle$  for every final outcome  $r$ . If  $\mathcal{N}_1 \equiv \mathcal{N}_2$  and  $\mathcal{N}_2$  consists of a single atom, then  $\mathcal{N}_2$  is a summary of  $\mathcal{N}_1$ .

Notice that, according to this definition, all unsound negotiations are equivalent. This amounts to considering soundness essential for a negotiation: if it fails, we do not care about the rest.

### 3.1 Deciding Soundness

The reachability graph of a negotiation  $\mathcal{N}$  has all markings reachable from  $\mathbf{x}_0$  as vertices, and an arc leading from  $\mathbf{x}$  to  $\mathbf{x}'$  whenever  $\mathbf{x} \xrightarrow{(n,r)} \mathbf{x}'$ .

The soundness problem consists of deciding if a given negotiation is sound. It can be solved by (1) computing the reachability graph of  $\mathcal{N}$  and (2a) checking that every atom appears at some arc, and (2b) that, for every reachable marking  $\mathbf{x}$ , there is an occurrence sequence  $\sigma$  such that  $\mathbf{x} \xrightarrow{\sigma} \mathbf{x}_f$ .

Step (1) needs exponential time, and steps (2a) and (2b) are polynomial in the size of the reachability graph. So the algorithm is single exponential in the number of atoms. This cannot be easily avoided, because the problem is PSPACE-complete, and NP-complete for acyclic negotiations.

Recall that a language  $L$  is in the class DP if there exist languages  $L_1, L_2$  in NP and co-NP, respectively, such that  $L = L_1 \cap L_2$  [20].

**Theorem 1.** *The soundness problem is PSPACE-complete. For acyclic negotiations, the problem is co-NP-hard and in DP (and so at level  $\Delta_2^P$  of the polynomial hierarchy).*

*Proof.* (Sketch) Membership in PSPACE follows easily from Savitch’s theorem and closure of NPSpace under complement. For PSPACE-hardness, given a deterministic linearly bounded automaton  $A$  and a word  $w$ , we construct a negotiation  $\mathcal{N}_{A,w}$  such that  $A$  accepts  $w$  iff  $\mathcal{N}_{A,w}$  is sound.  $\mathcal{N}_{A,w}$  has an agent  $C$  (for control), an agent  $H$  (for head), and an agent  $T_i$  for every tape cell of  $A$ . The negotiation has an atom  $n[q, h, a]$  for every state  $q$ , head position  $h$ , and input letter  $a$ , plus initial and final atoms. The transition function is defined so that  $\mathcal{N}_{A,w}$  simulates the computation of  $A$  on  $w$ .

Membership in DP follows easily from the fact that checking the first (second) condition in the definition of soundness lies in NP (co-NP, respectively). Co-NP-hardness follows by a standard reduction from 3-CNF-UNSAT.  $\square$

### 3.2 A Summarization Algorithm

The *summarization problem* consists of computing a summary of a given negotiation, if it is sound. A straightforward solution follows these steps:

- (1) Compute the reachability graph of  $\mathcal{N}$ . Interpret it as a weighted finite automaton over the alphabet of transformers  $\langle n, r \rangle$ , with  $\mathbf{x}_0$  as initial state, and  $\mathbf{x}_f$  as final state.
- (2) Compute the sum over all paths  $\sigma$  leading from  $\mathbf{x}_0$  to  $\mathbf{x}_f$  of the transformers  $\langle \sigma \rangle$ . We recall a well-known algorithm for this based on state elimination (see e.g. [18]). The algorithm proceeds in phases consisting of the following three steps:
  - (2.1) Exhaustively replace steps  $\mathbf{x} \xrightarrow{f_1} \mathbf{x}'$ ,  $\mathbf{x} \xrightarrow{f_2} \mathbf{x}'$  by one step  $\mathbf{x} \xrightarrow{f_1+f_2} \mathbf{x}'$ .
  - (2.2) Pick a state  $\mathbf{x}$  different from  $\mathbf{x}_0$  and  $\mathbf{x}_f$ . If there is a self-loop  $\mathbf{x} \xrightarrow{f} \mathbf{x}$ , replace all steps  $\mathbf{x} \xrightarrow{g} \mathbf{x}'$ , where  $\mathbf{x}' \neq \mathbf{x}$ , by the step  $\mathbf{x} \xrightarrow{g^*f} \mathbf{x}'$ , and then remove the self-loop.
  - (2.3) For every two steps  $\mathbf{x}_1 \xrightarrow{f_1} \mathbf{x}$  and  $\mathbf{x} \xrightarrow{f_2} \mathbf{x}_2$ , add a step  $\mathbf{x} \xrightarrow{f_1f_2} \mathbf{x}_2$  and remove state  $\mathbf{x}$  together with its incident steps.

Step (1) takes exponential time in the number of atoms. Steps (2.1)-(2.3) can be seen as *reduction rules* that replace an automaton by a smaller one with the same sum over all paths. In the next section we provide similar rules, but at the syntactic level, i.e., rules which act directly on the negotiation diagram, and *not* on the reachability graph. Two of the three rules are straightforward generalizations of (2.1) and (2.3) above, while the third allows us to remove certain useless arcs from a negotiation.

## 4 Reduction Rules

A *reduction rule*, or just a rule, is a binary relation on the set of negotiations. Given a rule  $R$ , we write  $\mathcal{N}_1 \xrightarrow{R} \mathcal{N}_2$  for  $(\mathcal{N}_1, \mathcal{N}_2) \in R$ . A rule  $R$  is *correct* if it preserves equivalence, i.e., if  $\mathcal{N}_1 \xrightarrow{R} \mathcal{N}_2$  implies  $\mathcal{N}_1 \equiv \mathcal{N}_2$ . Notice that, in particular, this implies that  $\mathcal{N}_1$  is sound if and only if  $\mathcal{N}_2$  is sound.

Given a set of rules  $\mathcal{R} = \{R_1, \dots, R_k\}$ , we denote by  $\mathcal{R}^*$  the reflexive and transitive closure of  $R_1 \cup \dots \cup R_k$ . We say that  $\mathcal{R}$  is *complete with respect to a class of negotiations* if, for every negotiation  $\mathcal{N}$  in the class, there is a negotiation  $\mathcal{N}'$  consisting of a single atom such that  $\mathcal{N} \xrightarrow{\mathcal{R}^*} \mathcal{N}'$ .

We describe rules as pairs of a *guard* and an *action*;  $\mathcal{N}_1 \xrightarrow{R} \mathcal{N}_2$  holds if  $\mathcal{N}_1$  satisfies the guard and  $\mathcal{N}_2$  is a possible result of applying the action to  $\mathcal{N}_1$ .

**Merge Rule.** Intuitively, the *merge rule* merges two outcomes with identical next enabled atoms into one single outcome. It corresponds to the rule of step (2.1) in the previous section.

**Definition 7.** *Merge rule*

**Guard:**  $N$  contains an atom  $n$  with two distinct outcomes  $r_1, r_2 \in R_n$ , such that  $\mathcal{X}(n, a, r_1) = \mathcal{X}(n, a, r_2)$  for every  $a \in A_n$ .

**Action:** (1)  $R_n \leftarrow (R_n \setminus \{r_1, r_2\}) \cup \{r_f\}$ , where  $r_f$  is a fresh name.

(2) For all  $a \in P_n$ :  $\mathcal{X}(n, a, r_f) \leftarrow \mathcal{X}(n, a, r_1)$ .

(3)  $\delta(n, r_f) \leftarrow \delta(n, r_1) \cup \delta(n, r_2)$ .

It is easy to see that the merge rule is correct for arbitrary negotiations.

**Shortcut Rule.** The shortcut rule corresponds to the rule of step (2.3) in the previous section. We need a preliminary definition.

**Definition 8.** *Given atoms  $n, n'$ , we say that  $(n, r)$  unconditionally enables  $n'$  if  $P_n \supseteq P_{n'}$  and  $\mathcal{X}(n, a, r) = \{n'\}$  for every  $a \in P_{n'}$ .*

Observe that if  $(n, r)$  unconditionally enables  $n'$  then, for every marking  $\mathbf{x}$  that enables  $n$ , the marking  $\mathbf{x}'$  given by  $\mathbf{x} \xrightarrow{(n,r)} \mathbf{x}'$  enables  $n'$ . Moreover,  $n'$  can only be disabled by its own occurrence.

The shortcut rule merges the outcomes of two atoms that can occur one after the other into one single outcome with the same effect. Consider the negotiation fragment shown on the left of Figure 2. The guard of the rule will state that  $n$  must unconditionally enable  $n'$ , which is the case. For every outcome of  $n'$ , say  $r_1$ , the action of the rule adds a fresh outcome  $r_{1f}$  to  $n$ , and modifies the negotiation so that the occurrence of  $(n, r_{1f})$  has the same effect as the occurrence of the sequence  $(n, r)(n', r_1)$ . In the figure, shortcutting the outcome  $(n, r)$  leaves  $n'$  without any input arc, and in this case the rule also removes  $n'$ . Otherwise we require that at least one input arc of a party  $\tilde{a}$  of  $n'$  is an arc (i.e., not a proper hyperarc) from some atom  $\tilde{n} \neq n$ , annotated by  $\tilde{r}$ . This implies that after the occurrence of  $(\tilde{n}, \tilde{r})$ ,  $n'$  is the only atom agent  $\tilde{a}$  is ready to engage in.

**Definition 9.** *Shortcut rule*

**Guard:**  $N$  contains an atom  $n$  with an outcome  $r$ , and an atom  $n'$ ,  $n' \neq n$ , such that  $(n, r)$  unconditionally enables  $n'$ . Moreover, if  $n' \in \mathcal{X}(\tilde{n}, \tilde{a}, \tilde{r})$  for at least one  $\tilde{n} \neq n$  with  $\tilde{a} \in P_{\tilde{n}}$  and  $\tilde{r} \in R_{\tilde{n}}$ , then  $\{n'\} = \mathcal{X}(\tilde{n}, \tilde{a}, \tilde{r})$  for some  $\tilde{n} \neq n$ ,  $\tilde{a} \in P_{\tilde{n}}$ ,  $\tilde{r} \in R_{\tilde{n}}$ .

**Action:** (1)  $R_n \leftarrow (R_n \setminus \{r\}) \cup \{r'_f \mid r' \in R_{n'}\}$ , where  $r'_f$  are fresh names.

(2) For all  $a \in P_{n'}$ ,  $r' \in R_{n'}$ :  $\mathcal{X}(n, a, r'_f) \leftarrow \mathcal{X}(n', a, r')$ .

For all  $a \in P \setminus P_{n'}$ ,  $r' \in R_{n'}$ :  $\mathcal{X}(n, a, r'_f) \leftarrow \mathcal{X}(n, a, r)$ .

(3) For all  $r' \in R_{n'}$ :  $\langle n, r'_f \rangle \leftarrow \langle n, r \rangle \langle n', r' \rangle$ .

(4) If  $\mathcal{X}^{-1}(n') = \emptyset$  after (1)-(3), then remove  $n'$  from  $N$ , where  $\mathcal{X}^{-1}(n') = \{(\tilde{n}, \tilde{a}, \tilde{r}) \in T(N) \mid n' \in \mathcal{X}(\tilde{n}, \tilde{a}, \tilde{r})\}$ .

**Theorem 2.** *The shortcut rule is correct.*

*Proof.* (Sketch) Let  $\mathcal{N}_2$  be the result of applying the rule to  $\mathcal{N}_1$ . The proof is based on the following observation: Each initial occurrence sequence of  $\mathcal{N}_2$  can be translated to a corresponding initial occurrence sequence of  $\mathcal{N}_1$  by replacing






Fig. 2. An application of the shortcut rule

each occurrence of  $(n, r'_f)$  by  $(n, r)(n', r')$ . Conversely, since  $(n, r)$  unconditionally enables  $n'$ , each  $(n, r)$ -step of an initial occurrence sequence of  $\mathcal{N}_1$  has a corresponding subsequent  $(n', r')$ -step, or the occurrence sequence can be extended by a  $(n', r')$ -step. This  $(n', r')$ -step is independent from all steps in between, i.e., no participant of  $n'$  is involved in any of the steps in between. Therefore, the occurrence sequence of  $\mathcal{N}_1$  (or its extension by  $(n', r')$ ) can be reordered to obtain pairs  $(n, r)(n', r')$  which can be translated to  $(n, r'_f)$ -steps of  $\mathcal{N}_2$ .

This mutual relation of initial occurrence sequences is used to show that an atom occurs in one negotiation if and only if it occurs in the other one, and that an occurrence sequence of  $\mathcal{N}_1$  can be extended by (i.e., is prefix of) another sequence leading to the final marking if the same holds for the corresponding occurrence sequence of  $\mathcal{N}_2$ , and vice versa.

The full proof requires some delicate case distinctions because  $n'$  might still exist in  $\mathcal{N}_2$  or not, and in the first case there are occurrences of  $n'$  in  $\mathcal{N}_2$  that do not belong to pairs  $(n, r)(n', r')$ .  $\square$

**Useless Arc Rule.** Consider the negotiation on the left of Figure 3, in which all atoms have one outcome  $r$ . We have  $\mathcal{X}(n_0, a, r) = \{n_1, n_f\}$ , i.e., after the occurrence of  $(n_0, r)$  agent  $a$  is ready to engage in both  $n_1$  and  $n_f$ . However,  $a$  always engages in  $n_1$ , because the only large step is  $(n_0, r)(n_1, r)(n_2, r)(n_f, r)$ . In other words, we can set  $\mathcal{X}(n_0, a, n_f) = \{n_1\}$  without changing the behavior. Intuitively, we say that the arc (more precisely, the leg of the hyperarc) leading to the  $a$ -port of  $n_f$  is useless. The useless arc rule identifies and removes some useless arcs.




Fig. 3. The useless arc rule can only be applied to the left negotiation

**Definition 10.** *Useless arc rule.*

**Guard:** *There are  $(n, a, r), (n, b, r) \in T(N)$  and two distinct atoms  $n', n'' \in N$  such that  $a, b \in P_{n'} \cap P_{n''}$ ,  $n', n'' \in \mathcal{X}(n, a, r)$  and  $\mathcal{X}(n, b, r) = \{n'\}$ .*

**Action:**  $\mathcal{X}(n, a, r) \leftarrow \mathcal{X}(n, a, r) \setminus \{n''\}$ .

The rule can be applied to the negotiation on the left of Figure 3 by instantiating  $n := n_0$ ,  $n' := n_1$ , and  $n'' := n_f$ . It cannot be applied to the negotiation on the right. If we set  $n := n_0$ ,  $n' := n_1$ , and  $n'' := n_2$ , then  $a \notin P_{n_2}$ .

**Theorem 3.** *The useless arc rule is correct.*

*Proof.* (Sketch) Let  $\mathcal{N}_2$  be the result of applying the rule to  $\mathcal{N}_1$ . We prove that  $\mathcal{N}_1$  and  $\mathcal{N}_2$  have the same initial occurrence sequences, from which the result easily follows. Let  $\mathcal{X}_1, \mathcal{X}_2$  be the transition functions of  $\mathcal{N}_1$  and  $\mathcal{N}_2$ , respectively. Since  $\mathcal{X}_2(n, a, r) \subseteq \mathcal{X}_1(n, a, r)$  for every  $(n, a, r) \in T(N)$ , every initial occurrence sequence of  $\mathcal{N}_2$  is also an initial occurrence sequence of  $\mathcal{N}_1$ . To prove that every initial occurrence sequence  $\sigma$  of  $\mathcal{N}_1$  is also an initial occurrence sequence of  $\mathcal{N}_2$ , assume there is  $\sigma = \sigma_1\sigma_2$  such that  $\sigma_1$  is an initial occurrence sequence of both  $\mathcal{N}_1$  and  $\mathcal{N}_2$  whereas the first step of  $\sigma_2$  is only possible in  $\mathcal{N}_1$ . Since both negotiations only differ w.r.t.  $\mathcal{X}(n, a, r)$ , this step must be an occurrence of agent  $n''$  occurring after an occurrence of  $(n, r)$  such that no other atom of  $\mathcal{X}(n, a, r)$  occurred in between. This holds in particular for  $n' \in \mathcal{X}(n, a, r)$ . But, since  $\mathcal{X}(n, b, r) = \{n'\}$  and  $b$  participates both in  $n'$  and in  $n''$ ,  $n'$  must occur before  $n''$  after  $(n, r)$  in the occurrence sequence – a contradiction.  $\square$

## 5 (Weakly) Deterministic Negotiations

**Definition 11.** *An agent  $a \in A$  is deterministic if for every  $(n, a, r) \in T(N)$  such that  $n \neq n_f$  there exists one atom  $n'$  such that  $\mathcal{X}(n, a, r) = \{n'\}$ .*

*The negotiation  $\mathcal{N}$  is weakly deterministic if for every  $(n, a, r) \in T(N)$  there is a deterministic agent  $b$  that is a party of every atom in  $\mathcal{X}(n, a, r)$ , i.e.,  $b \in P_{n'}$  for every  $n' \in \mathcal{X}(n, a, r)$ . It is deterministic if all its agents are deterministic.*

Graphically, an agent  $a$  is deterministic if no proper hyperarc leaves any port of  $a$ . Consider the negotiations of Figure 1. In the acyclic negotiation both Father and Daughter are deterministic, while Mother is not. In the ping-pong negotiation only Daughter is deterministic. Both negotiations are weakly deterministic, because Daughter participates in all atoms, and so can be always chosen as the party  $b$  required by the definition. Observe that the notion of deterministic agent does not refer to the behavior of atoms, which is intrinsically nondeterministic with respect to its possible outcomes and even to its state transformations. Rather, it refers to the *composition* of negotiations: For each atom  $n$ , the next atom of a deterministic agent is completely determined by the outcome of  $n$ .

Weakly deterministic negotiations have a natural semantic justification. Consider a negotiation with two agents  $a, b$  and three atoms  $\{n_0, n_1, n_f\}$ . All atoms have the same parties  $a, b$  and one outcome  $r$ , such that  $\mathcal{X}(n_0, a, r) = \{n_1, n_f\} = \mathcal{X}(n_0, b, r)$  and  $\mathcal{X}(n_1, a, r) = \{n_f\} = \mathcal{X}(n_1, b, r)$ . After the occurrence of  $(n_0, r)$

the parties  $a$  and  $b$  are ready to engage in both  $n_1$  and  $n_f$ , and so which of them occurs requires a “meta-negotiation” between  $a$  and  $b$ . This meta-negotiation, however, is not part of the model and, more importantly, it can be difficult to implement, since it requires to break a symmetry. In a weakly deterministic negotiation this situation cannot happen. If  $\mathcal{X}(n, a, r)$  contains more than one atom, then some deterministic agent  $b$  is a part of all atoms in  $\mathcal{X}(n, a, r)$ . If some atom of  $\mathcal{X}(n, a, r)$  becomes enabled, say  $n'$ , then because agent  $b$  is ready to engage in it, and, since  $b$  is deterministic,  $b$  is not ready to engage in any other atom. So  $n'$  is the only enabled atom of  $\mathcal{X}(n, a, r)$ , and  $n'$  is the negotiation that  $a$  will engage in next. This is very easy to implement:  $b$  just sends a message to  $a$  telling her that she should commit to  $n'$ .

For deterministic negotiations, the second part of the guard of the shortcut rule is always satisfied. Using the notation of the shortcut rule, this condition requires that the atom  $n'$  is the only atom in  $\mathcal{X}(\tilde{n}, \tilde{a}, \tilde{r})$  for some  $(\tilde{n}, \tilde{r})$ , provided  $n'$  is not only in  $\mathcal{X}(n, a, r)$  for some  $a \in P_n$  (i.e., provided  $n'$  is not removed by the application of the rule). This clearly holds if  $\tilde{n}$  is deterministic, as all agents are deterministic in deterministic negotiations.

In the next section we show that, on top of their semantic justification, weakly deterministic and deterministic negotiations are also interesting from an analysis point of view. We prove that the shortcut and useless arc rules are complete for acyclic, weakly deterministic negotiations, which of course implies that the same rules plus the merge are complete, too. A second result proves that a polynomial number of applications of the merge and shortcut rules suffices to summarize any sound deterministic acyclic negotiation (the useless arc rule is irrelevant for deterministic negotiations).

## 6 Completeness and Complexity

We start with the completeness result for the weakly deterministic case.

**Theorem 4.** *The shortcut and useless arc rules are complete for acyclic, weakly deterministic negotiations.*

*Proof.* (Sketch) Let  $\mathcal{N}$  be a sound and acyclic weakly deterministic negotiation. The proof has two parts.

(1) If  $\mathcal{N}$  has more than one atom, then the shortcut rule or the useless arc rule can be applied to it.

Since  $\mathcal{N}$  is acyclic, its graph generates a partial order on atoms in the obvious way ( $n < n'$  if there is a path from  $n$  to  $n'$ ). Clearly  $n_0$  is the unique minimal element. We choose an arbitrary linearisation of this partial order. Since  $\mathcal{N}$  has more than one atom, this linearisation begins with  $n_0$  and has some second element, say  $n_1$ .

The full proof shows that, if  $\{n_1\} = \mathcal{X}(n_0, a, r_0)$  for every party  $a$  of  $n_1$ , then the shortcut rule can be applied to  $n_0, n_1$ , and if  $\{n_1\} \neq \mathcal{X}(n_0, a, r_0)$  for some party  $a$  of  $n_1$ , then the useless arc rule is applicable.

(2) The shortcut and useless arc rules cannot be applied infinitely often.

This follows from two facts: the application of the shortcut rule does not increase the length of any large step of a negotiation, and strictly decreases the length of at least one large step (by replacing a sequence of two outcomes by one single outcome); the application of the useless arc rule does not change the large steps but decreases the number of arcs.

By (2) every maximal sequence of applications of the rule terminates. By (1) it terminates with a negotiation containing one single atom.  $\square$

Next we prove that a *polynomial number* of applications of the merge and shortcut rules suffice to summarize any sound *deterministic* acyclic negotiation (SDAN). For this we have to follow a strategy in the application of the shortcut rule.

**Definition 12.** *The deterministic shortcut rule, or d-shortcut rule, is the result of adding to the guard of the shortcut rule a new condition: (3)  $n'$  has at most one outcome (the actions of the shortcut and d-shortcut rules coincide).*

We say that a SDAN is *irreducible* if neither the merge nor the d-shortcut rule can be applied to it. In the rest of the section we prove that irreducible SDANs are necessarily atomic. The proof, which is rather involved, proceeds in three steps. First, we prove a technical lemma showing that SDANs can be reduced so that all agents participate in every atom with more than one outcome. Then we use this result to prove that, loosely speaking, every SDAN can be reduced to a “replication” of a negotiation with only one agent: if  $\mathcal{X}(n, a, r) = \{n'\}$  for some agent  $a$ , then  $\mathcal{X}(n, a, r) = \{n'\}$  for every agent  $a$ . In the third step, we show that replications can be reduced to atomic negotiations. Finally, we analyze the number of rule applications needed in each of these three steps, and conclude.

**Lemma 1.** *Let  $\mathcal{N}$  be an irreducible SDAN and let  $n \neq n_f$  be an atom of  $\mathcal{N}$  with more than one outcome. Then every agent participates in  $n$ .*

*Proof.* (Sketch) We first prove the following key claim: The atom  $n$  has an outcome  $r$  such that either  $(n, r)$  unconditionally enables  $n_f$ , or  $(n, r)$  unconditionally enables some atom with more than one outcome.

To prove the claim, we first show that it suffices to prove: if some outcome  $(n, r)$  unconditionally enables some atom, then the claim holds. For this we assume the contrary, and prove that  $\mathcal{N}$  then contains a cycle, contradicting the hypothesis.

By repeated application of the claim we find a chain  $(n_1, r_1) \dots (n_k, r_k)$  such that  $n_1 = n$ ,  $n_k = n_f$ , and  $(n_i, r_i)$  unconditionally enables  $n_{i+1}$  for every  $1 \leq i \leq k - 1$ . By the definition of “unconditionally enabled” we have  $P_1 \supseteq P_2 \supseteq \dots \supseteq P_k = P_f$ . Since  $P_f = A$ , we obtain  $P_1 = A$ .  $\square$

**Lemma 2.** *Let  $\mathcal{N}$  be an irreducible SDAN. Every agent participates in every atom, and for every atom  $n \neq n_f$  and every outcome  $r$  there is an atom  $n'$  satisfying  $\mathcal{X}(n, a, r) = \{n'\}$  for every agent  $a$ .*

*Proof.* We first show that every agent participates in every atom. By Lemma 1, it suffices to prove that every atom  $n \neq n_f$  has more than one outcome. Assume the contrary, i.e., some atom different from  $n_f$  has only one outcome. Since, by soundness, every atom can occur, there is an occurrence sequence  $(n_0, r_0)(n_1, r_1) \cdots (n_k, r_k)$  such that  $n_k$  has only one outcome and all of  $n_0, \dots, n_{k-1}$  have more than one outcome. By Lemma 1, all agents participate in all of  $n_0, n_1, n_{k-1}$ . It follows that  $(n_i, r_i)$  unconditionally enables  $(n_{i+1}, r_{i+1})$  for every  $0 \leq i \leq k-1$ . In particular,  $(n_{k-1}, r_{k-1})$  unconditionally enables  $(n_k, r_k)$ . But then, since  $n_k$  only has one outcome, the d-shortcut rule can be applied to  $n_{k-1}, n$ , contradicting the hypothesis that  $\mathcal{N}$  is irreducible.

For the second part, assume there is an atom  $n \neq n_f$ , an outcome  $r$  of  $n$ , and two distinct agents  $a_1, a_2$  such that  $\mathcal{X}(n, a_1, r) = \{n_1\} \neq \{n_2\} = \mathcal{X}(n, a_2, r)$ . By the first part, every agent participates in  $n, n_1$  and  $n_2$ . Since  $\mathcal{N}$  is sound, some reachable marking  $\mathbf{x}$  enables  $n$ . Moreover, since all agents participate in  $n$ , and  $\mathcal{N}$  is deterministic, the marking  $\mathbf{x}$  only enables  $n$ . Let  $\mathbf{x}'$  be the marking given by  $\mathbf{x} \xrightarrow{(n,r)} \mathbf{x}'$ . Since  $a_1$  participates in all atoms, no atom different from  $n_1$  can be enabled at  $\mathbf{x}'$ . Symmetrically, no atom different from  $n_2$  can be enabled at  $\mathbf{x}'$ . So  $\mathbf{x}'$  does not enable any atom, contradicting that  $\mathcal{N}$  is sound.  $\square$

**Theorem 5.** *Let  $\mathcal{N}$  be an irreducible SDAN. Then  $\mathcal{N}$  contains only one atom.*

*Proof.* (Sketch) Assume  $\mathcal{N}$  contains more than one atom. For every atom  $n \neq n_f$ , let  $l(n)$  be the length of the longest path from  $n$  to  $n_f$  in the graph of  $\mathcal{N}$ . Let  $n_{\min}$  be any atom such that  $l(n_{\min})$  is minimal, and let  $r$  be an arbitrary outcome of  $n_{\min}$ . By Lemma 2 there is an atom  $n'$  such that  $\mathcal{X}(n_{\min}, a, r) = \{n'\}$  for every agent  $a$ . If  $n' \neq n_f$  then by acyclicity we have  $l(n') < l(n_{\min})$ , contradicting the minimality of  $n_{\min}$ . So we have  $\mathcal{X}(n_{\min}, a, r) = \{n_f\}$  for every outcome  $r$  of  $n_{\min}$  and every agent  $a$ . If  $n_{\min}$  has more than one outcome, then the merge rule is applicable. If  $n_{\min}$  has one outcome, then, since it unconditionally enables  $n_f$ , the d-shortcut rule is applicable. In both cases we get a contradiction to irreducibility.  $\square$

**Definition 13.** *For every atom  $n$  and outcome  $r$ , let  $shoc(n, r)$  be the length of a shortest maximal occurrence sequence containing  $(n, r)$  minus 1, and let  $Shoc(\mathcal{N}) = \sum_{n \in \mathcal{N}, r \in R} shoc(n, r)$ . Finally, let  $Out(\mathcal{N}) = \sum_{(P,R,\delta) \in \mathcal{N} \setminus \{n_f\}} |R|$  be the total number of outcomes of  $\mathcal{N}$ , excluding those of the final atom.*

Notice that if  $\mathcal{N}$  has  $K$  atoms then  $shoc(n, r) \leq K - 1$  holds for every atom  $n$  and outcome  $r$ . Further, if  $K = 1$  then  $Shoc(\mathcal{N}) = 0 = Out(\mathcal{N})$ .

**Theorem 6.** *Every SDAN  $\mathcal{N}$  can be completely reduced by means of  $Out(\mathcal{N})$  applications of the merge rule and  $Shoc(\mathcal{N})$  applications of the d-shortcut rule.*

*Proof.* Let  $\mathcal{N}$  and  $\mathcal{N}'$  be negotiations such that  $\mathcal{N}'$  is obtained from  $\mathcal{N}$  by means of the merge or the d-shortcut rule. For the merge rule we have  $Out(\mathcal{N}') < Out(\mathcal{N})$  and  $Shoc(\mathcal{N}') \leq Out(\mathcal{N})$  because the rule reduces the number of outcomes by one. For the d-shortcut rule we have  $Out(\mathcal{N}') = Out(\mathcal{N})$  because if it is applied to pairs  $n, n'$  such that  $n'$  has one single outcome, and  $Shoc(\mathcal{N}') < Shoc(\mathcal{N})$ , because  $Shoc(n, r'_f) < Shoc(n, r)$ .  $\square$

## 7 Conclusions

We have introduced negotiations, a formal model of concurrency with negotiation atoms as primitive. We have defined and studied two important analysis problems: soundness, which coincides with the soundness notion for workflow nets, and the new *summarization problem*. We have provided a complete set of rules for sound acyclic, weakly deterministic negotiations, and we have shown that the rules allow one to compute a summary of a sound deterministic negotiations in polynomial time.

Several open questions deserve further study. Our results show that summarization can be solved in polynomial time for deterministic, acyclic negotiations, and is co-NP-hard for arbitrary acyclic negotiations. The precise complexity of the weak deterministic case is still open. We are currently working on a generalization of Rule 2.2. of Section 3.2, such that we can completely reduce in polynomial time *arbitrary* deterministic negotiations, even if they contain cycles.

**Related Work.** Previous work on Petri net analysis by means of reductions has already been discussed in the Introduction.

A number of papers have modelled specific distributed negotiation protocols with the help of Petri nets or process calculi (see [21,19,3]). However, these papers do not address the issue of negotiation as concurrency primitive.

The feature of summarizing parts of a negotiation to single negotiation atoms has several analogies in Petri net theory, among these the concept of zero-safe Petri nets. By abstracting from reachable markings which mark distinguished “zero-places”, transactions can be modelled by zero-safe Petri nets [8]. Reference [7] extends this concept to reconfigurable, dynamic high-level Petri nets.

A related line of research studies global types and session types to model multi-party sessions [17]. See [9] for an overview that also covers choreography-based approaches for web services. This research emphasises communication aspects in the formal setting of mobile processes. Thus, the aim differs from our aim. However, it might be worth trying to combine the two approaches.

Finally, the graphical representation of negotiations was partly inspired by the BIP component framework [4,6], where a set of sequential components (i.e., the agents) interact by synchronizing on certain actions (i.e., the atoms).

**Acknowledgement.** We thank the reviewers, in particular for their hints to related papers. We also thank Stephan Barth, Eike Best, and Jan Kretinsky for very helpful discussions.

## References

1. van der Aalst, W.M.P.: The application of Petri nets to workflow management. J. Circuits, Syst. and Comput. 8(1), 21–66 (1998)
2. van der Aalst, W.M.P., van Hee, K.M., ter Hofstede, A.H.M., Sidorova, N., Verbeek, H.M.W., Voorhoeve, M., Wynn, M.T.: Soundness of workflow nets: classification, decidability, and analysis. Formal Asp. Comp. 23(3), 333–363 (2011)

3. Bacarin, E., Madeira, E.R.M., Medeiros, C.B., van der Aalst, W.M.P.: *SpiCa's* multi-party negotiation protocol: Implementation using YAWL. *Int. J. Cooperative Inf. Syst.* 20(3), 221–259 (2011)
4. Basu, A., Bensalem, S., Bozga, M., Combaz, J., Jaber, M., Nguyen, T.-H., Sifakis, J.: Rigorous component-based system design using the BIP framework. *IEEE Software* 28(3), 41–48 (2011)
5. Berthelot, G.: Transformations and decompositions of nets. In: Brauer, W., Reisig, W., Rozenberg, G. (eds.) APN 1986. LNCS, vol. 254, pp. 359–376. Springer, Heidelberg (1987)
6. Bliudze, S., Sifakis, J.: The algebra of connectors - structuring interaction in BIP. *IEEE Trans. Computers* 57(10), 1315–1330 (2008)
7. Bruni, R., Melgratti, H.C., Montanari, U.: Extending the zero-safe approach to coloured, reconfigurable and dynamic nets. In: Desel, J., Reisig, W., Rozenberg, G. (eds.) ACPN 2003. LNCS, vol. 3098, pp. 291–327. Springer, Heidelberg (2004)
8. Bruni, R., Montanari, U.: Executing transactions in zero-safe nets. In: Nielsen, M., Simpson, D. (eds.) ICATPN 2000, pp. 83–102 (2000)
9. Castagna, G., Dezani-Ciancaglini, M., Padovani, L.: On global types and multi-party session. *Logical Methods in Computer Science* 8(1) (2012)
10. Desel, J.: Reduction and design of well-behaved concurrent systems. In: Baeten, J.C.M., Klop, J.W. (eds.) CONCUR 1990. LNCS, vol. 458, pp. 166–181. Springer, Heidelberg (1990)
11. Desel, J., Esparza, J.: *Free choice Petri nets*. Cambridge University Press, New York (1995)
12. Desel, J., Esparza, J.: On negotiation as concurrency primitive. Technical report, Technische Universität München, Germany (2013) Available via arxiv.org
13. van Dongen, B.F., van der Aalst, W.M.P., Verbeek, H.M.W.: Verification of EPCs: Using reduction rules and Petri nets. In: Pastor, Ó., Falcão e Cunha, J. (eds.) CAiSE 2005. LNCS, vol. 3520, pp. 372–386. Springer, Heidelberg (2005)
14. Genrich, H.J., Thiagarajan, P.S.: A theory of bipolar synchronization schemes. *Theor. Comput. Sci.* 30, 241–318 (1984)
15. Haddad, S.: A reduction theory for coloured nets. In: Rozenberg, G. (ed.) APN 1989. LNCS, vol. 424, pp. 209–235. Springer, Heidelberg (1990)
16. Haddad, S., Pradat-Peyre, J.-F.: New efficient Petri nets reductions for parallel programs verification. *Parallel Processing Letters* 16(1), 101–116 (2006)
17. Honda, K., Yoshida, N., Carbone, M.: Multiparty asynchronous session types. In: Necula, G.C., Wadler, P. (eds.) POPL, pp. 273–284. ACM (2008)
18. Hopcroft, J.E., Motwani, R., Ullman, J.D.: *Introduction to Automata Theory, Languages, and Computation*, 3rd edn. Addison-Wesley Longman Publishing Co., Inc., Boston (2006)
19. Ji, S., Tian, Q., Liang, Y.: A Petri-net-based modeling framework for automated negotiation protocols in electronic commerce. In: Lukose, D., Shi, Z. (eds.) PRIMA 2005. LNCS, vol. 4078, pp. 324–336. Springer, Heidelberg (2009)
20. Papadimitriou, C.H., Yannakakis, M.: The complexity of facets (and some facets of complexity). In: Lewis, H.R., Simons, B.B., Burkhard, W.A., Landweber, L.H. (eds.) STOC, pp. 255–260. ACM (1982)
21. Salaün, G., Ferrara, A., Chirichello, A.: Negotiation among web services using LOTOS/CADP. In: Zhang, L.-J., Jeckle, M. (eds.) ECOWS 2004. LNCS, vol. 3250, pp. 198–212. Springer, Heidelberg (2004)
22. Verbeek, H.M.W., Wynn, M.T., van der Aalst, W.M.P., ter Hofstede, A.H.M.: Reduction rules for reset/inhibitor nets. *J. Comp. Syst. Sci.* 76(2), 125–143 (2010)