
Now or Never: Negotiating Efficiently with Unknown Counterparts

Toni Mancini

Computer Science Department, Sapienza University of Rome, Italy
<http://tmancini.di.uniroma1.it>

Abstract. We define a new protocol rule, Now or Never (NoN), for bilateral negotiation processes which allows self-motivated competitive agents to *efficiently* carry out multi-variable negotiations with remote untrusted parties, where privacy is a major concern and agents know *nothing* about their opponent. By building on the geometric concepts of convexity and convex hull, NoN ensures a continuous progress of the negotiation, thus neutralising malicious or inefficient opponents. In particular, NoN allows an agent to derive in a finite number of steps, and *independently* of the behaviour of the opponent, that there is *no hope* to find an agreement. To be able to make such an inference, the interested agent may *rely on herself only*, still keeping the *highest freedom* in the choice of her strategy.

We also propose an actual NoN-compliant strategy for an automated agent and evaluate the computational feasibility of the overall approach on instances of practical size.

1 Introduction

Automated negotiation among rational agents is crucial in Distributed Artificial Intelligence domains as, e.g., resource allocation [3], scheduling [16], e-business [7], and applications where: (i) no agent can achieve her own goals without interaction with the others (or she is expected to achieve more utility with interaction), and (ii) constraints of various kinds (e.g., security or privacy) forbid the parties to communicate their desiderata to others (the opponent or a trusted authority), hence centralised approaches cannot be used.

We present a framework which allows two *self-motivated, competitive* agents to negotiate *efficiently* and find a mutually satisfactory agreement in a particularly *hostile* environment, where each party has *no information* on constraints, preferences, and *willingness to collaborate* of the opponent. This means that also the *bounds* of the domains of the negotiation variables are *not* common knowledge. Our framework deals with negotiations over multiple *constrained* variables over the type of *real numbers*, regarding integer or categorical variables as special cases.

The present setting is very different from what is often assumed in the literature: the set of possible agreements is *infinite* and agents do not even know (or probabilistically estimate) possible opponent's types, variable domain bounds or most preferred values. It is not a split-the-pie game as, e.g., in [6] although with incomplete information, as in [11], and computing *equilibrium* or evaluating *Pareto-optimality* is not possible.

A major problem in our setting is that even *termination* of the negotiation process is not granted: it is in general *impossible* for the single agent to recognise whether the negotiation is making some progress, or if the opponent is just wasting time or arbitrarily delaying the negotiation outcome.

We solve this problem by proposing a new protocol rule, Now or Never (NoN) (Section 3), explicitly designed as to ensure a continuous progress of the negotiation. The rule (whose fulfilment can be assessed *independently* by each party using only the exchanged information) forces the agents to never reconsider already taken decisions, thus injecting a minimum, but sufficient amount of *efficiency* in the process. This leads to the *monotonic shrinking* of the set of possible agreements, which in turn allows each agent to derive in a finite number of steps, *independently* of the behaviour of the opponent, that there is *no hope* to find an agreement.

Furthermore, we discuss the notion of *non-obstructionist* agents, i.e., agents who genuinely aim at efficiently finding an agreement, even sacrificing their preferences (among the agreements they would accept). If both agents are non-obstructionist, the NoN rule guarantees that, whenever the termination condition arises, then *no agreement actually exists*. Hence, in presence of non-obstructionist agents, our approach is both *complete* and *terminates*.

We also propose (Section 4) a full NoN-compliant strategy for an agent which ensures termination independently of the behaviour of the opponent. The strategy, which takes into full account the presence of a utility function on the set of acceptable deals, is inspired to the well-known mechanism of Monotonic Concessions (MC) [13] and allows the agent to perform a sophisticated reasoning, based on the evidence collected so far on the behaviour of the opponent, to select the best deals to offer at each step and keep the process as efficient as possible.

Section 5 specialises NoN to discrete and categorical variables and Section 6 presents experimental results showing that enforcing the NoN rule in practical negotiation instances is computationally feasible.

2 Preliminaries and Negotiation Framework

In the following, we denote with \mathbb{R} the set of real numbers and with \mathbb{N}^+ the set of strictly positive integers.

Our framework deals with (possibly multi-deal) negotiations between *two* agents (agent 0 and agent 1) over *multiple constrained* variables. Agents do *not* have any information about constraints, goals, preferences, reasoning capabilities, willingness to collaborate, and strategy of the opponent. *The only knowledge common to both agents is the set of negotiation variables and the protocol rules.*

Definition 1 introduces the main concepts of our framework. Some of them are standard in the literature and are adapted to our framework to ease presentation of the following definitions and results.

Definition 1 (Negotiation process). A negotiation process *is* a tuple $\pi = \langle \mathcal{V}, s, k, \mathcal{R} \rangle$ where \mathcal{V} is a finite set of negotiation variables, $s \in \{0, 1\}$ is the agent starting the negotiation, and \mathcal{R} is the set of protocol rules.

The negotiation space *is* the multi-dimensional real vector space $\mathbb{R}^{|\mathcal{V}|}$. Each point $D \in \mathbb{R}^{|\mathcal{V}|}$ is a deal. A proposal for π *is* a set of at most k deals or the

distinguished element \perp . Value $k \in \mathbb{N}^+$ is the maximum number of deals that can be included in a single proposal.

Negotiation proceeds in steps (starting from step 1) with agents (starting from agent s) alternately exchanging proposals. The proposal exchanged at any step $t \geq 1$ is sent by agent $\text{ag}(t)$, defined as s if t is odd and $1 - s$ if t is even.

The status of negotiation process π at step $t \geq 1$ is the sequence $\mathcal{P} = \mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_t$ of proposals exchanged up to step t .

At each step, the status of π must satisfy the set \mathcal{R} of protocol rules, a set of boolean conditions on sequences of proposals.

A strategy for agent $A \in \{0, 1\}$ for π is a function σ_A that, for each step t such that $\text{ag}(t) = A$ and each status $\mathcal{P} = \mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_{t-1}$ of π at step $t - 1$, returns the proposal \mathcal{P}_t to be sent by agent A at step t , given the sequence of proposals already exchanged (σ_A is constant for $t = 1$ and $A = s$).

Our alternating offers [14] based framework primarily focuses on *real variables*. In Section 5 we discuss how more specialised domains (e.g., integers, categories) can be handled as special cases, and which is the added value of primarily dealing with real variables. Also, as each proposal can contain up to k deals, our framework supports *multi-deal* negotiations when $k > 1$. Section 4 discusses the added value given by the possibility of exchanging multi-deal proposals.

Protocol rules are important to prevent malicious or inefficient behaviour. Well-designed rules are of paramount importance when the process involves self-motivated and/or unknown/untrusted opponents. For protocol rules to be effective, agents must be able at any time to verify them using the current negotiation status only.

We will use Example 1 as a running example throughout the paper.

Example 1 (Alice vs. Bob). Alice wants to negotiate with her supervisor Bob to schedule a meeting. At the beginning, agents agree on the relevant variables \mathcal{V} : (i) the start day/time t ; (ii) the meeting duration d .

Deals are assignments of values to variables, as, e.g., $D = \langle t = \text{“Mon 11 am”}, d = \text{“30 min”} \rangle$. Deals can be easily encoded as points in \mathbb{R}^2 .

Definition 2 (Negotiation outcomes). Let $\pi = \langle \mathcal{V}, s, k, \mathcal{R} \rangle$ be a negotiation process whose status at step $T > 1$ is $\mathcal{P} = \mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_T$. We say that π terminates at step T if and only if T is the smallest value such that one of the following two cases holds:

- success: $\mathcal{P}_T = \{D\} \subseteq \mathcal{P}_{T-1}$ ($\text{ag}(T)$ accepts deal D proposed by $\text{ag}(T - 1)$ at step $T - 1$)
- opt-out: $\mathcal{P}_T = \perp$ ($\text{ag}(T)$ opts-out).

If no such a T exists, then π is non-terminating (non-term).

Success, opt-out and non-term are the possible negotiation outcomes.

A negotiation process can be infinite (case *non-term*) or terminate in a finite number of steps, either with an *agreement* found (case *success*, where point D is the *agreement*) or with a failure (case *opt-out*, where one of the agents proposes \perp , which aborts the process).

For a deal to be *acceptable* to an agent, some *constraints* must be satisfied. Such constraints, which are *private information* of the single agent, are formalised by Definition 3.

Definition 3 (Feasibility region). Let $\pi = \langle \mathcal{V}, s, k, \mathcal{R} \rangle$ be a negotiation process. The feasibility region of agent $A \in \{0, 1\}$, denoted by R_A , is the subset of the negotiation space $\mathbb{R}^{|\mathcal{V}|}$ of deals acceptable to A .

For agent A , any deal in R_A is better than failure. An *agreement* is thus any deal $D \in R_0 \cap R_1$.

Example 2 (Alice vs. Bob (cont.)). Alice wants the meeting no later than Wednesday. Normally she needs at least 30 minutes and does not want the meeting to last more than one hour; however, if she has to wait until Wednesday, she would have time during her Tuesday’s trip to prepare new material to show; in this case she wants the meeting to last at least one hour, but no more than 75 minutes. Conversely, Bob has his own, private, constraints.

Fig. 1a shows Alice’s feasibility region in a 2D space, as the areas delimited by the three polygons. The region takes into account duties in her agenda (e.g., Alice is busy on Monday from 1pm to 4pm).

Agents may have *preferences* on the deals in their feasibility region. Such preferences are often represented by a private *utility function*. Fig. 1a shows that, e.g., Alice prefers a long meeting on Monday. We will handle the agent utility function in Section 4 when we present a full strategy for an agent.

We assume (as typically done, see, e.g., [5,12]) that agents offer only deals in their feasibility region (i.e., agents do not offer deals they are not willing to accept). This does not limit our approach, as suitable ex-post measures (e.g., penalties) can be set up to cope with the case where a deal offered by an agent (but *not* acceptable to her) is accepted by the other.

3 Now or Never

We are interested in negotiations which are *guaranteed* to terminate in a finite number of steps (note that, being negotiation variables real-valued, the set of potential agreements is infinite), so we want to avoid case *non-term* of Definition 2. In this section we define a protocol rule, the Now or Never (NoN) rule, which is our key to drive a negotiation process towards termination, avoiding malicious or inefficient agents behaviour. The rule relies on the notions of Definition 4.

Definition 4 (Convex region, convex hull, operator \cup). Let \mathbb{R}^n be the n -dimensional real vector space (for any $n > 0$). Region $R \subseteq \mathbb{R}^n$ is convex if, for any two points D_1 and D_2 in R , the straight segment $\overline{D_1 D_2}$ is entirely in R .

Given a finite set of points $\mathcal{D} \subset \mathbb{R}^n$, the convex hull of \mathcal{D} , $\text{conv}(\mathcal{D})$, is the smallest convex region of \mathbb{R}^n containing \mathcal{D} .

Given a collection of finite sets of points \mathcal{D} , $\cup \mathcal{D}$ is the union of the convex hulls of all sets in \mathcal{D} : $\cup \mathcal{D} = \bigcup_{\mathcal{D} \in \mathcal{D}} \{\text{conv}(\mathcal{D})\}$.

Convexity arises often in feasibility regions of agents involved in negotiations. An agent feasibility region is convex if, for any two acceptable deals D_1 and D_2 , all *intermediate* deals (i.e., those lying on $\overline{D_1 D_2}$) are acceptable as well. In some cases [2,12,4] the feasibility region of an agent is *entirely* convex (consider, e.g., a negotiation instance over a single variable, the price of a good). In other cases

this does not hold. However, a feasibility region may always be considered as the *union* of a number of convex sub-regions. Furthermore, in most real cases, this number is *finite* and *small*. Also, in most practical situations, the closer two acceptable deals D_1 and D_2 , the higher the likelihood that intermediate deals are acceptable as well.

Example 3 (Alice vs. Bob (cont.)). Knowing that deals $\langle t = \text{“Mon at 11am”}, d = \text{“30 min”} \rangle$ and $\langle t = \text{“Wed at 3pm”}, d = \text{“1 hour”} \rangle$ are both acceptable to Bob would not be a strong support for Alice to assume that also $\langle t = \text{“Tue at 1pm”}, d = \text{“45 min”} \rangle$ would be acceptable to him. On the other hand, if $\langle t = \text{“Mon at 9am”}, d = \text{“40 min”} \rangle$ and $\langle t = \text{“Mon at 9.30am”}, d = \text{“20 min”} \rangle$ are both acceptable to Bob, it would not be surprising if also $\langle t = \text{“Mon at 9.15am”}, d = \text{“30 min”} \rangle$ is acceptable.

Before formalising the NoN rule (Definition 6), we introduce it using our example.

Example 4 (Alice vs. Bob (cont.)). Steps below are shown in Fig. 1.

Steps 1 and 2. Alice starts the negotiation by sending proposal $\mathcal{P}_1 = \{A_1^a, A_1^b\}$. As a reply, she receives $\mathcal{P}_2 = \{B_2^a, B_2^b\}$ (see Fig. 1b). As none of Bob’s counteroffers, B_2^a and B_2^b , belong to $\text{conv}(\{A_1^a, A_1^b\}) = A_1^a A_1^b$, all such deals are *removed* from further consideration (by exploiting NoN). The rationale is as follows:

(a) Bob had *no evidence* that $\text{conv}(\{A_1^a, A_1^b\})$ includes deals outside R_{Alice} (i.e., at the end of step 1 Bob had no evidence that this portion of R_{Alice} is not convex).

(b) Given that Bob has not proposed any such deal therein, then either $R_{Bob} \cap \text{conv}(\{A_1^a, A_1^b\}) = \emptyset$ (in which case, Bob has no interest at all in proposing there), or Bob has chosen not to go for any such a deal *now* (as, e.g., he currently aims at higher utility).

(c) In the latter case, NoN forbids Bob to reconsider that decision anymore (*never*).

Step 3. Alice, having *no evidence* that $\text{conv}(\{B_2^a, B_2^b\})$ includes deals outside R_{Bob} , proposes \mathcal{P}_3 containing deal $A_3^a \in \text{conv}(\{B_2^a, B_2^b\}) \cap R_{Alice}$ (see Fig. 1c): by proposing A_3^a she aims at closing the negotiation successfully *now*, believing that such a deal (intermediate to B_2^a and B_2^b) is likely to be acceptable also to Bob. Alice also includes in \mathcal{P}_3 deal A_3^b .

Step 4. It’s Bob’s turn again. By receiving $\mathcal{P}_3 = \{A_3^a, A_3^b\}$, Bob knows that such deals belong to R_{Alice} . Assume that Bob rejects \mathcal{P}_3 by sending a counteroffer. As there is no evidence that $\text{conv}(\{A_1^a, A_3^a, A_3^b\})$, $\text{conv}(\{A_1^b, A_3^a, A_3^b\})$, or $\text{conv}(\{A_1^a, A_3^a, A_3^b\})$ (the 3 light-grey areas in Fig. 1c) include deals outside R_{Alice} , NoN forces him to take a decision: either his counteroffer \mathcal{P}_4 contains some deals in one of such regions, or he must forget those regions forever. Note that NoN does not apply to, e.g., $\text{conv}(\{A_1^b, A_3^a, A_3^b\})$, as this region contains B_2^b , which was part of a Bob’s proposal already rejected by Alice. Hence, there is already evidence that some of the deals in $\text{conv}(\{A_1^b, A_3^a, A_3^b\})$ are acceptable to Bob and NoN does not forbid agents to further explore that region.

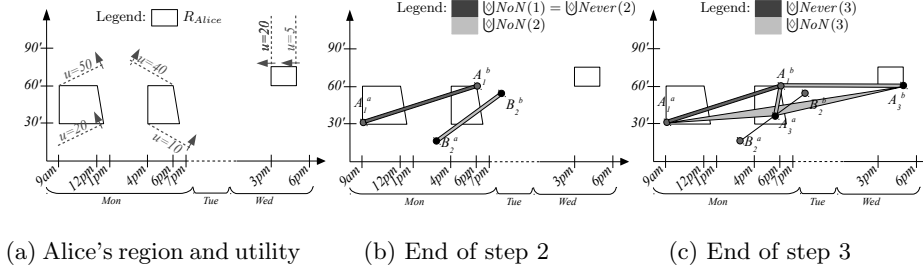


Fig. 1: Alice vs. Bob (Example 4)

Definition 5 (Sets Never and NoN). Let $\pi = \langle \mathcal{V}, s, k, \mathcal{R} \rangle$ be a negotiation process and $\mathcal{P} = \mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_T$ its status at step $T \geq 1$. For each agent $A \in \{0, 1\}$ and each step $1 \leq t \leq T$, let $\text{deals}_A(t)$ be the set of all the deals in \mathcal{P} proposed by A up to step t (included). Sets $\text{Never}(t)$ and $\text{NoN}(t)$ are defined inductively for each $t \geq 1$ as follows:

$$t=1: \text{Never}(1) = \emptyset, \text{NoN}(1) = \{\mathcal{P}_1\}$$

$$t>1:$$

$$\text{Never}(t) = \begin{cases} \text{Never}(t-2) \cup \text{NoN}(t-1) & \text{if } \mathcal{P}_t \cap \text{NoN}(t-1) = \emptyset \\ \text{Never}(t-2) \cup \{\{D\} \mid D \in \text{NoN}(t-1)\} & \text{otherwise} \end{cases}$$

$$\text{NoN}(t) = \{D \subseteq \text{deals}_{\text{ag}(t)}(t) \mid \text{conv}(D) \cap \text{NoN}(t-1) = \emptyset\}$$

where \mathcal{P}_t is the proposal sent by $\text{ag}(t)$ at step t and $\text{Never}(0) = \emptyset$.

At each step t , $\text{NoN}(t)$ represents the region, defined by $\text{ag}(t)$'s deals, for which the other agent $1 - \text{ag}(t)$ needs, in the next step ($t + 1$) to take a NoN decision: to offer a deal therein (showing to $\text{ag}(t)$ that she is potentially interested to that region) or to neglect that region forever. Similarly, $\text{Never}(t)$ represents the region, defined by $(1 - \text{ag}(t))$'s deals, for which $\text{ag}(t)$ has taken a never decision. Deals therein cannot be offered any more. Note that $\text{Never}(t) \supseteq \text{Never}(t-2)$ for all $t \geq 2$ (i.e., sequences $\text{Never}(t)$ for odd and even values of t are monotonically non-decreasing). Fig. 1 shows NoN and Never regions at all steps of the previous example.

Definition 6 formalises our NoN protocol rule, which forbids agents to reconsider never decisions already taken.

Definition 6 (Now or Never rule). Status $\mathcal{P} = \mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_T$ of negotiation process $\pi = \langle \mathcal{V}, s, k, \mathcal{R} \rangle$ satisfies the NoN protocol rule if, for all steps $2 \leq t \leq T$, $\mathcal{P}_t \cap \text{Never}(t-2) = \emptyset$.

Proposition 1 shows that the NoN rule of Definition 6 allows agents to infer when no further agreements are possible. All proofs are omitted for lack of space.

Proposition 1 (Termination condition). Let $\pi = \langle \mathcal{V}, s, k, \mathcal{R} \rangle$ be a negotiation process where the NoN rule is enforced and let $\mathcal{P} = \mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_T$ be the status of π at step $T \geq 2$.

If $R_{\text{ag}(T)} \subseteq \text{Never}(T-1) \cup \text{NoN}(T-2)$ and \mathcal{P}_T is not a singleton $\{D\} \subseteq \mathcal{P}_{T-1}$, then:

- (a) there exists no extension $\mathcal{P}' = \mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_{T-1}, \mathcal{P}_T, \dots, \mathcal{P}_{T'}$ of \mathcal{P} to step $T' > T$ such that $\mathcal{P}_{T'} = \{D\} \subseteq \mathcal{P}_{T'-1}$
 (b) for all $D \in R_0 \cap R_1$, there exists $1 < t_D < T$ such that $D \in \bigcup \text{NoN}(t_D - 1) \cap \bigcup \text{Never}(t_D)$.

A consequence of (a) is that, if at step $T \geq 2$, $R_{\text{ag}(T)} \subseteq \bigcup \text{Never}(T - 1) \cup \bigcup \text{Never}(T - 2)$ and agent $\text{ag}(T)$ cannot or does not want to accept a deal offered in the last incoming proposal \mathcal{P}_{T-1} , she can safely opt-out by proposing $\mathcal{P}_T = \perp$, as she has no hope to reach an agreement in the future. Also, from (b), for every mutually acceptable agreement D , there was a step $t_D < T$ in which agent $\text{ag}(t_D)$ took a *never* decision on a NoN region containing D . This means that $\text{ag}(t_D)$, although knowing that $D \in R_{\text{ag}(t_D)}$ was likely to be acceptable also to the opponent (because she had no evidence, at that time, that the portions of the opponent region defined by deals in $\text{NoN}(t_D - 1)$ were not convex), explicitly decided not to take that chance and proposed elsewhere.

As a matter of fact, NoN can be thought as a *deterrent*, for each agent, to delay the negotiation by ignoring plausible agreements which, although acceptable to her, do not grant herself the utility she currently aims at. As NoN forbids the agents to propose such deals in the future, any such “obstructionist” behaviour has a price in terms of opportunities that must be sacrificed forever.

Definition 7 defines *non-obstructionist agents*.

Definition 7 (Non-obstructionist agent). Let $\pi = \langle \mathcal{V}, s, k, \mathcal{R} \rangle$ be a negotiation process where the NoN rule is enforced. Agent $A \in \{0, 1\}$ is non-obstructionist if her strategy satisfies the following conditions for all $t \geq 2$ such that $\text{ag}(t) = A$:

1. if $\mathcal{P}_{t-1} \cap R_A \neq \emptyset$, then $\mathcal{P}_t = \{D\} \subseteq \mathcal{P}_{t-1}$
2. else if $\bigcup \text{NoN}(t-1) \cap R_A \neq \emptyset$, then $\mathcal{P}_t \cap \bigcup \text{NoN}(t-1) \neq \emptyset$.

A non-obstructionist agent A accepts *any* acceptable deal $D \in R_A$ and takes a *now* decision at all steps t when $\bigcup \text{NoN}(t-1)$ intersects R_A . Non-obstructionist agents genuinely aim at finding an agreement *efficiently*, even sacrificing their preferences among deals they would accept. However, they are *not* necessarily collaborative, as they do not disclose to the opponent their constraints and preferences.

Proposition 2 shows that, in a negotiation process between two non-obstructionist agents, if one of the parties reaches the *termination condition* of Proposition 1, then no agreement exists (i.e., $R_0 \cap R_1 = \emptyset$).

Proposition 2 (Completeness). Let $\pi = \langle \mathcal{V}, s, k, \mathcal{R} \rangle$ be a negotiation process between two non-obstructionist agents where the NoN rule is enforced.

If π reaches, at step $T - 1 \geq 2$, status $\mathcal{P} = \mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_{T-1}$ s.t. $R_{\text{ag}(T)} \subseteq \bigcup \text{Never}(T - 1) \cup \bigcup \text{Never}(T - 2)$, then $R_0 \cap R_1 = \emptyset$.

4 A Terminating Strategy Based on Monotonic Concessions

Propositions 1 and 2 show that the Now or Never (NoN) rule allows each agent to detect when the negotiation process can be safely terminated, as no agreement can be found in the sequel. However, still the termination condition may

not arise in a finite number of steps. In this section we show that, with NoN, *termination can be enforced by any agent alone, without relying on the willingness to terminate of the counterpart*. To this end, from now on we focus on one agent only, which we call agent A (A can be either 0 or 1). To ease presentation, the other agent, agent $1 - A$, will be called agent B .

We make some assumptions on the feasibility region of agent A : (a) R_A is *bounded* and defined as the union $P_1 \cup \dots \cup P_q$ of a *finite* number q of convex sub-regions; (b) each convex sub-region P_i ($1 \leq i \leq q$) of R_A is defined by *linear constraints*, hence is a (bounded) *polyhedron* in $\mathbb{R}^{|\mathcal{V}|}$. Any bounded feasibility region can be approximated arbitrarily well with a (sufficiently large) union of bounded polyhedra. However, in many practical cases, a finite and small number of polyhedra suffices.

Deals in R_A may not be equally worth for agent A , who may have a (again, *private*) utility function u_A to *maximize*. We assume that u_A is *piecewise-linear* and defined (without loss of generality) by a linear function u_A^i for each polyhedron P_i of R_A ($1 \leq i \leq q$). For this definition to be well founded, if a deal D belongs to two different polyhedra P_i and P_j of R_A , it must be $u_A^i(D) = u_A^j(D)$. Note that, again, any differentiable utility function can be approximated arbitrarily well with a piecewise-linear utility, provided R_A is decomposed in an enough number of polyhedra.

In this setting, we define a full strategy for agent A for negotiation processes $\pi = \langle \mathcal{V}, s, k, \mathcal{R} \rangle$ for which $k \geq 2$, i.e., in which exchanged proposals can contain multiple deals. Although our strategy is correct independently of the opponent region shape, it is designed for the common cases where agent A believes that the opponent feasibility region is the union of a small number of convex sub-regions (not necessarily polyhedra). Hence, a task of agent A while following the strategy is to *discover* non-convexities of the opponent region during negotiation and take them into account.

Our strategy is inspired by (but different from) the well-known mechanism of Monotonic Concessions (MC) [13]. It has three phases, *utility-driven*, *non-obstructionist*, and *terminating* phases, which are executed in the given order.

4.1 Utility-Driven Phase

Agent A keeps and dynamically revises two utility thresholds, α and u , which are, respectively, the *responding* and the *proposing threshold*. At each step t such that $\text{ag}(t) = A$, agent A uses: (a) threshold α to decide whether to take a *now* decision (if $t > 1$), by including, in the proposal \mathcal{P}_t she will propose next, a deal in $\bigcup \text{NoN}(t-1)$ (possibly accepting one deal in \mathcal{P}_{t-1}), and (b) threshold u to select the other deals to include in \mathcal{P}_t ($t \geq 1$).

By generalising [6,12], α is a function of the agent A utility of the best deal D_{next} that would be chosen in step (b). In particular, α is $u_A(D_{\text{next}}) - \text{span} \cdot \xi$, where span is the absolute difference of the extreme values of u_A in R_A and $0 \leq \xi \leq 1$ is a parameter (possibly varying during the negotiation) called *respond policy*. Hence, if $\xi = 1$, agent A accepts all acceptable deals and takes a *now* decision whenever possible, behaving in a non-obstructionist way (Definition 7). On the other extreme, if $\xi = 0$ the agent accepts only incoming deals $D \in R_A$ that are not worse than the best proposal D_{next} that would be chosen next in step (b), upon rejection of D .

Our strategy for this phase is decomposed into *responding*, *proposing* and *conceding* sub-strategies as in [8], after an *initialisation* phase where agent A sets u to the highest utility of deals in R_A (as in the spirit of MC).

Responding. At step $t \geq 2$, after that agent A has received proposal $\mathcal{P}_{t-1} \neq \perp$, proposal \mathcal{P}_t is chosen as follows. Let $R_A^\alpha = \{D \in R_A \mid u_A(D) \geq \alpha\}$.

(1) If \mathcal{P}_{t-1} contains deals in $R_A^\alpha - \bigcup \text{Never}(t-2)$, then $\mathcal{P}_t = \{D\}$, where D is one such a deal giving agent A the highest utility (i.e., agent A accepts the best deal D among those acceptable in \mathcal{P}_{t-1} granting herself at least utility α). Otherwise:

(2) \mathcal{P}_t contains a deal in $(\bigcup \text{NoN}(t-1) \cap R_A^\alpha) - \bigcup \text{Never}(t-2)$ if and only if this region is not empty (*now* decision taken).

Given that the closer deals in a set \mathcal{D} defining $\text{NoN}(t-1)$ (see Definition 5) the more likely they belong to a single convex sub-region of R_B , as for (2) agent A selects a deal with the highest utility in a set \mathcal{D} having the minimum *diameter*.

Proposing. At any step $t \geq 1$ such that $\text{ag}(t) = A$, if agent A has not accepted an incoming deal (case (1) of the *responding* sub-strategy), proposal \mathcal{P}_t contains *additional* deals (as to make $|\mathcal{P}_t| = k \geq 2$). Let $R_A^u = \{D \in R_A \mid u_A(D) \geq u\}$ (which is again a union of bounded polyhedra, as u_A is piecewise-linear). Deals to be proposed in \mathcal{P}_t are carefully selected among *vertices* of R_A^u (some of them can be vertices of the overall region R_A) which do not belong to $\bigcup \text{Never}(t-2)$, as agent A needs to comply with the NoN rule. If $t > 1$, vertices of R_A^u to be proposed will be carefully selected by reasoning on the *evidence* provided by the past opponent behaviour. The reasoning is as follows.

Let $\hat{n}(t)$ be the *minimum* number of convex sub-regions that *must* compose $R_B - \bigcup \text{Never}(t-1)$, i.e., the opponent region minus the regions for which the opponent has taken a *never* decision (and in which, by the NoN rule, no agreements can be found in the sequel): $\hat{n}(t)$ is the minimum value such that there exists a $\hat{n}(t)$ -partition $\{\mathcal{D}_1, \dots, \mathcal{D}_{\hat{n}(t)}\}$ of $\text{deals}_B(t-1)$ (i.e., a mapping of each opponent deal to one sub-region) such that for all $1 \leq j \leq \hat{n}(t)$, $\text{conv}(\mathcal{D}_j) \cap \bigcup \text{Never}(t-1) = \emptyset$.

Agent A temporarily focuses on $\hat{n}(t)$, assuming that $R_B - \bigcup \text{Never}(t-1)$ is the union of *exactly* $\hat{n}(t)$ convex sub-regions. We call this assumption *Non-obstructionist Opponent Assumption (NOA)*. Under NOA, agent A tries to regard the past opponent behaviour as non-obstructionist, hence interprets the already taken *never* decisions as an admission that $R_B \cap \bigcup \text{Never}(t-1) = \emptyset$ (Proposition 2). Value $\hat{n}(t)$ is the minimum number of convex sub-regions that must compose R_B which is consistent with this (optimistic) hypothesis.

Agent A computes the subsets \mathcal{D} of the opponent deals that *might* belong to the same convex sub-region of $R_B - \bigcup \text{Never}(t-1)$, provided that NOA is correct. We call these sets of deals *Possible Opponent Clusters (POCs)*:

$$\mathcal{K}(t) = \left\{ \mathcal{D} \subseteq \text{deals}_B(t-1) \mid \begin{array}{l} \exists \hat{n}(t)\text{-partition } \{\mathcal{D}_1, \dots, \mathcal{D}_{\hat{n}(t)}\} \text{ of } \text{deals}_B(t-1) \\ \text{s.t. } \forall j \in [1, \hat{n}(t)] \text{ conv}(\mathcal{D}_j) \cap \bigcup \text{Never}(t-1) = \emptyset \end{array} \right\} \quad (1)$$

Let $\text{proj}(R, R')$ (the *projection* of region R onto R') be the set of points X for which there exists $Y \in R$ such that \overline{XY} intersects R' [2]. Region $\text{proj}(R, R')$ is an unbounded polyhedron if both R and R' are polyhedra (see Fig. 2a, where $\text{proj}(R, R')$ is the unbounded grey area) and $\text{proj}(R, R' \cup R'') = \text{proj}(R, R') \cup$

$\text{proj}(R, R'')$. Provided that NOA is correct, agent A can derive (Proposition 3) that region

$$\Pi(t) = \bigcap_{\mathcal{D} \in \mathcal{K}(t)} \text{proj}(\text{conv}(\mathcal{D}), \mathcal{N}Never(t-1))$$

does not contain agreements that can be still reached.

Proposition 3. *If, at step $t \geq 3$ s.t. $\text{ag}(t) = A$, NOA is correct, then $\Pi(t) \cap (R_B - \mathcal{N}Never(t-1)) = \emptyset$.*

Example 5 (Alice vs. Bob (cont.)). Consider Fig. 2b. At step 4 Bob sent Alice proposal $\mathcal{P}_4 = \{B_4^a\}$. At step 5 (Alice's turn), $\hat{n}(5)$ is 3, as it is clear that B_2^a, B_2^b , and B_4^a belong to all-different convex sub-regions of $R_{Bob} - \mathcal{N}Never(4)$. POCs are $\mathcal{K}(5) = \{\{B_2^a\}, \{B_2^b\}, \{B_4^a\}\}$. Region $\Pi(5)$ is the area in light-grey: if NOA is correct ($R_{Bob} - \mathcal{N}Never(4)$ or, equivalently, R_{Bob} if Bob is non-obstructionist, consists of exactly 3 convex sub-regions), then no $X \in \Pi(5)$ can belong to $R_{Bob} - \mathcal{N}Never(4)$.

Besides always ignoring vertices in $\mathcal{N}Never(t-2)$ (as to comply with the NoN rule), as a result of Proposition 3 agent A (exploiting NOA) can *temporarily ignore* vertices of R_A^u in $\Pi(t)$ while choosing deals to propose at step t . By exploiting the fail-first principle, we define the following criterion (*best vertex under NOA*) to select the next vertices in $R_A^u - \Pi(t)$ (and not in $\mathcal{N}Never(t-2)$) to propose: those that, if rejected, would make the *highest* number of vertices be excluded in the next step, under NOA.

Conceding. When no more vertices in $R_A^u - \Pi(t)$ (and not in $\mathcal{N}Never(t-2)$) can be proposed, agent A reduces threshold u , if possible, by a given amount Δu , whose value, possibly varying during time (see, e.g, [5]), depends on the application. Reducing u is in the spirit of MC (where the agent *increases* during time the *opponent* utility of the proposed deals). Differently from MC, here agent A *reduces own* utility of the deals she proposes (with the goal of approaching opponent's demand), as she has no information about opponent utility.

Let \hat{T} ($\text{ag}(\hat{T}) = A$) be the step in which agent A reduces u and R_A^u becomes equal to R_A (i.e., u cannot be further reduced). From step \hat{T} onwards, the strategy of agent A moves to the *non-obstructionist* phase.

4.2 Non-Obstructionist Phase

Our strategy for this phase is decomposed into *responding* and *proposing* sub-strategies. As utility threshold u has already reached its minimum, in this phase there is no *conceding* sub-strategy.

Responding. The *responding* sub-strategy is identical to that of the *utility-driven* phase with $\alpha = u$. Given that in the non-obstructionist phase u is at its minimum, agent A accepts any incoming acceptable deal and takes a *now* decision whenever possible. Thus, the agent is now certainly non-obstructionist, independently of the value of her respond policy ξ .

Proposing. As a result of acting in a non-obstructionist way, from step \hat{T} onwards the following result holds:

Proposition 4. *For each step $t \geq \hat{T}$ such that $\text{ag}(t) = A$, $R_A \cap \mathcal{N}ever(t-2) = R_A \cap \mathcal{N}ever(\hat{T}-2)$.*

Hence, for each step $t \geq \hat{T}$ such that $\text{ag}(t) = A$, if agent A has not accepted an incoming deal, the region in which the additional deals to propose will be selected (as to make $|\mathcal{P}_t| = k \geq 2$ whenever possible), i.e., $R_A - \mathcal{N}ever(t-2)$, is steadily equal to $R_A - \mathcal{N}ever(\hat{T}-2)$.

In this phase, agent A aims at proposing *vertices* of $R_A - \mathcal{N}ever(\hat{T}-2)$ with the goal of eventually covering it within the *never* set of the opponent, as to reach the termination condition of Proposition 1. Unfortunately, as both R_A and $\mathcal{N}ever(\hat{T}-2)$ are unions of polyhedra, their difference might *not* be represented as a union of polyhedra. Anyway, it can be always represented as a union of Not Necessarily Closed polyhedra (i.e., polyhedra possibly defined by some *strict* inequalities, with some of their faces and vertices *not* belonging to them). In order to comply with the NoN rule, the agent must not propose vertices of $R_A - \mathcal{N}ever(\hat{T}-2)$ not belonging to that region, as they would belong to $\mathcal{N}ever(\hat{T}-2)$. The problem is solved by computing a suitable *under-approximation* $\lfloor R_A - \mathcal{N}ever(\hat{T}-2) \rfloor \subseteq R_A - \mathcal{N}ever(\hat{T}-2)$ which can be defined as a union of bounded (and closed) polyhedra. Note that such an under-approximation can be computed in order to make the error

$$R_A^{\text{err}} = (R_A - \mathcal{N}ever(\hat{T}-2)) - \lfloor R_A - \mathcal{N}ever(\hat{T}-2) \rfloor$$

arbitrarily small. As a special case, if agent A was non-obstructionist from the beginning of the negotiation process, $R_A \cap \mathcal{N}ever(\hat{T}-2) = \emptyset$ and $R_A^{\text{err}} = \emptyset$.

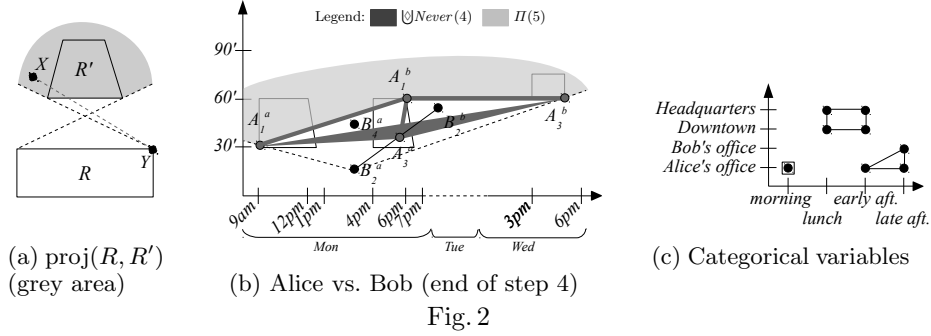
Agent A continues to use both NOA and $\Pi(t)$ as defined in the *utility-driven* phase. In particular, the agent proposes *vertices* of $\lfloor R_A - \mathcal{N}ever(\hat{T}-2) \rfloor$ which are not in $\Pi(t)$. When no more such vertices can be proposed, NOA is gradually relaxed (i.e., $\hat{n}(t)$ is gradually increased) and the remaining vertices of $\lfloor R_A - \mathcal{N}ever(\hat{T}-2) \rfloor$ are enabled. By construction, $\hat{n}(t)$ cannot grow beyond the number of deals proposed by the opponent so far. If also in that case $\Pi(t)$ covers $\lfloor R_A - \mathcal{N}ever(\hat{T}-2) \rfloor$, the agent sets $\Pi(t)$ to $\mathcal{N}ever(t-1)$, hence assumes that R_B consists of at least one convex sub-region *not yet* disclosed by the opponent (i.e., not containing any of the past incoming deals).

As it happens in the *utility-driven* phase, given that *multi-deal* proposals are allowed ($k \geq 2$), all vertices will be proposed within a finite number of steps independently of the number of *now* decisions taken. When all vertices have been proposed and no agreement has been reached, agent A enters the *terminating phase*.

4.3 Terminating Phase

In this phase, agent A continues by sending *empty proposals* until she receives and accepts an acceptable deal or infers $R_A - R_A^{\text{err}} \subseteq \mathcal{N}ever(T-1) \cup \mathcal{N}ever(T-2)$. Proposition 5 states that also this condition will arise in a finite number of steps.

Proposition 5. *Let $\pi = \langle \mathcal{V}, s, k, \mathcal{R} \rangle$ be a negotiation process ($k \geq 2$) where the NoN rule is enforced. If any agent $A \in \{0, 1\}$ uses the strategy above, then, within a finite number of steps $T \geq \hat{T} \geq 2$ such that $\text{ag}(T) = A$, either an agreement is found or condition $R_A - R_A^{\text{err}} \subseteq \mathcal{N}ever(T-1) \cup \mathcal{N}ever(T-2)$ is satisfied.*



Condition of Proposition 5 can be considered the *termination condition* of Proposition 1 in case agent A had admissible region $R_A - R_A^{\text{err}}$. Given that region R_A^{err} can be chosen as to be *arbitrarily small*, agent A can terminate the negotiation when this condition is reached. Any possible remaining acceptable deals would be in the (arbitrarily small) region R_A^{err} .

We stress again that, in case agent A is non-obstructionist from the beginning, for all $t \geq \hat{T}$ such that $\text{ag}(t) = A$, R_A^{err} can be made *empty*. Hence, as it happens for any acceptable deal in $R_A \cap \cup\text{Never}(t-2) = R_A \cap \cup\text{Never}(\hat{T}-2)$, any acceptable deal in R_A^{err} can be considered as an opportunity (with arbitrarily small Euclidean distance to $R_A \cap \cup\text{Never}(\hat{T}-2)$) that agent A had to sacrifice for having behaved in an obstructionist way (at most) up to step $\hat{T}-2$.

5 Handling Discrete and Categorical Variables

The NoN rule works also when (some of) the variables are discrete (e.g., integer), if we consider the union of the *integer hulls* [15] of the polyhedra in the *NoN* and *Never* sets of Definition 5. Integer Linear Programming results tell us that the integer hull of a polyhedron can still be represented with linear (plus integrality) constraints. Vertices of this new polyhedron have integer coordinates. Hence, the NoN rule as well as the strategy above and the underlying projection-based reasoning can be adapted to *prune* the space of the possible agreements: only the branches that deal with *now* decisions need to be refined (deals in $\cup\text{NoN}(t-1)$ proposed at step t need to have integer coordinates). Also, $R_A - \cup\text{Never}(\hat{T}-2)$ can always be represented by an union of closed polyhedra, hence R_A^{err} can be always made empty. Categorical variables can be tackled by fixing an ordering of their domain (common to both parties) and mapping them onto integers. Fig. 2c shows Alice's region in a variation of Example 1 where variables are categorical.

6 Implementation and Experiments

Our framework has at its core well-studied tasks in computational geometry and (integer) linear programming [15]. However, to our knowledge, the exact complexity of the agent's reasoning is unknown, as these core tasks must be repeated

on sets of exchanged deals. Still, existing libraries of computational geometry algorithms can manage the size of instances needed for practical scenarios. We have implemented a system that uses the Parma Polyhedra Library (PPL) [1] to compute polyhedra, convex hulls, and projections, and an all-solutions SAT solver [9] to revise $\hat{n}(t)$ and Possible Opponent Clusters (POCs) (the problem is reduced to hypergraph-colouring). Performance on these sub-problems are very good: PPL completes most of the required tasks within very few seconds (on a reasonably small set of variables, e.g., 3–4) and the generated SAT instances are trivial. Although the number of sets in *NoN* and *Never* can grow exponentially, by keeping only (depending on the case) their \subseteq -maximal or \subseteq -minimal members (which is enough to enforce the NoN rule and to perform the needed reasoning), the overall memory requirements become, in the instances we consider below, compatible with the amount of RAM available on an ordinary PC.

In the following we present an empirical evaluation of the *computational feasibility* of the approach. Negotiations have been performed between two identical agents. We evaluated our implementation on both random and structured instances using a *single* computer (a PC with a dual-core AMD Opteron 3GHz and 8GB RAM) for both agents. At each step, agents can exchange contracts of at most $k = 2$ deals. Note that, as our approach requires agents to comply with the NoN rule, it *cannot* be evaluated against other negotiators.

Random Instances. We generated 100 random negotiation instances over 3 variables. Feasibility regions are unions of 3 random polyhedra, each with at most 10 vertices. In about 44% of the instances $R_0 \cap R_1 \neq \emptyset$. The average volume of the intersection is 2.19% of the volume of each agent’s region (stddev is 4.5%). Agents have random piecewise-linear utilities and concede constant $\Delta u = 0.2span$ each time all vertices of R_a^u ($a \in \{0, 1\}$) belong to Π .

Such negotiations terminate in < 5 minutes and 20–30 steps. Agreements were found in $> 95\%$ of the instances for which $R_0 \cap R_1 \neq \emptyset$. Fig. 3 shows average time, success rate (i.e., number of negotiations closed successfully / number of negotiations such that $R_0 \cap R_1 \neq \emptyset$), and average quality of the agreement found for each agent (the quality of an agreement D for agent $a \in \{0, 1\}$ is $(u_a(D) - L_a)/(H_a - L_a)$, where H_a and L_a are, respectively, the highest and lowest values of agent a utility in $R_0 \cap R_1$) as a function of the respond policies used (ξ_0 and ξ_1).

It can be seen that moderate respond policies (intermediate values of ξ) lead to very high probabilities ($> 97\%$) of finding an agreement if one exists; moreover, the quality of such agreements for the two agents is similar if their respond policies are similar (fairness). Conversely, if agents use very different values for ξ , the more conceding agent unsurprisingly gets lower utility with the agreement, but negotiations are more often aborted by the other, more demanding, agent.

Structured Instances. We evaluated our system on the 6 scenarios in Table 1a: two scenarios (AB1, AB2) of the Alice vs. Bob example, two scenarios (SU1, SU2) of a negotiation problem regarding the rental of a summerhouse, and two scenarios (EZ1, EZ2) of a variation of the England-Zimbabwe problem of [10] adapted to our domain (real variables and no known bounds for their domains). A description of these negotiation scenarios is omitted for space reasons.

Table 1a shows also some relevant properties of these negotiation scenarios. Column “vars” gives the number of negotiation variables. Columns “polys” and “con” give, respectively, the number of polyhedra and the overall number of

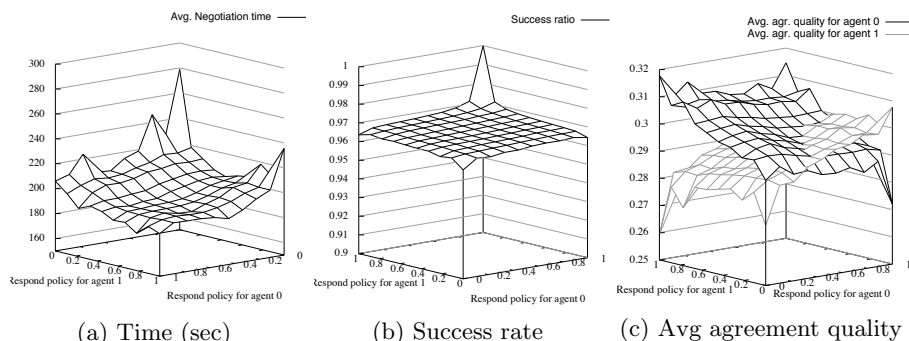


Fig. 3: Experimental results for random instances

Scenario	R_0		R_1		$\frac{vol(R_0 \cap R_1)}{vol(R_0)}$	$\frac{vol(R_0 \cap R_1)}{vol(R_1)}$	Scenario	ξ_0	ξ_1	agr. found	steps	time (sec)	polys	
	polys	con.	polys	con.										
AB1	2	3	12	3	12	$< 10^{-8}\%$	$< 10^{-8}\%$	AB1	1	1	Y	20	1.09	321
AB2	2	3	12	5	20	—	—	AB1	0	0	N	24	1.21	450
SU1	3	3	12	4	16	1.5%	2.0%	AB2	1	1	N	20	1.33	466
SU2	3	3	12	4	16	—	—	SU1	0	0	N	417	38.44	20413
EZ1	4	2	8	2	8	0.10%	0.05%	SU1	0.4	0	Y	347	15.22	5679
EZ2	4	2	8	2	8	0.15%	0.04%	SU2	1	1	N	513	63.12	29148
								EZ1	0	0	Y	80	1237	3115508
								EZ2	0	0	Y	92	2836	8057011

(a) Properties of negotiation scenarios

(b) Experimental results

Table 1: Structured instances

linear constraints defining each agent feasibility region, R_0 and R_1 . The two last columns give the ratio of the volume of $R_0 \cap R_1$ (i.e., the volume of the space of the possible agreements) with respect to the volume of the feasibility region of each agent (“—” means that $R_0 \cap R_1$ is empty, hence no agreement is possible).

Table 1b shows some results on the above negotiation scenarios, under different values of the respond policies of each agent (ξ_0 and ξ_1). All instances have been run with k (the maximum number of deals in a proposal) equal to 2. For each instance, column “agr. found” tells whether an agreement has been found (an agreement exists if and only if $R_0 \cap R_1 \neq \emptyset$, see Table 1a), column “steps” gives the number of negotiation steps needed to conclude the negotiation process, column “time” gives the overall negotiation time, and column “polys” gives the overall number of polyhedra computed by PPL during the process. For each negotiation instance, the number of all-SAT instances solved to compute POCs (see formula (1)) is equal to the number of negotiation steps.

Our results show that enforcing NoN is computationally feasible: negotiation processes with hundreds of interaction steps could be performed in minutes, even when NoN enforcement and agents reasoning require the computation of millions of polyhedra and the resolution of hundreds of all-SAT instances.

7 Conclusions

In this paper we defined a new protocol rule, Now or Never (NoN), for bilateral negotiation processes which allows self-motivated competitive agents to

efficiently carry out multi-variable negotiations with remote untrusted parties, where privacy is a major concern and agents know *nothing* about their opponent. NoN has been explicitly designed as to ensure a continuous progress of the negotiation, thus neutralising malicious or inefficient opponents.

We have also presented a NoN-compliant strategy for an agent that, under mild assumptions on her feasibility region, allows her to derive, in a finite number of steps and *independently* of the behaviour of her opponent, that there is *no hope* to find an agreement. We finally evaluated the computational feasibility of the overall approach on random and structured instances of practical size.

Acknowledgements. This research was funded by the EU 7th Framework Programme under grant agreements n. 317761 (SmartHG) and n. 600773 (PAEON).

References

1. R. Bagnara, P. M. Hill, and E. Zaffanella. The Parma Polyhedra Library: Toward a complete set of numerical abstractions for the analysis and verification of hardware and software systems. *Sci. of Comp. Progr.*, 72(1–2):3–21, 2008.
2. M. Cadoli. Proposal-based negotiation in convex regions. In *Proc. of CIA 2003*, v. 2782 of *LNCS*, pp. 93–108. Springer, 2003.
3. S. E. Conry, K. Kuwabara, and R. A. Lesser, V. R. Meyer. Multistage negotiation for distributed constraint satisfaction. *IEEE Transactions on Systems, Man and Cybernetics*, 21(6):462–477, 1991.
4. S. Costantini, G. De Gasperis, A. Provetti, and P. Tsintza. A heuristic approach to proposal-based negotiation: with applications in fashion supply chain management. *Math. Problems in Engin.*, 2013. Article ID 896312.
5. P. Faratin, C. Sierra, and N. R. Jennings. Negotiation decision functions for autonomous agents. *Intl. J. Robotics & Autonomous Systems*, 24(3–4):159–182, 1998.
6. P. Faratin, C. Sierra, and N. R. Jennings. Using similarity criteria to make issue trade-offs in automated negotiations. *Artif. Intell.*, 142(2):205–237, 2002.
7. N. R. Jennings, T. J. Norman, P. Faratin, P. O’Brien, and B. Odgers. Autonomous agents for business process management. *Appl. Artif. Intell.*, 14(2):145–189, 2000.
8. G. Lai and K. Sycara. A generic framework for automated multi-attribute negotiation. *Group Decision and Negotiation*, 18(2):169–187, 2009.
9. D. Le Berre and A. Parrain. The sat4j library, release 2.2. *Journal on Satisfiability, Boolean Modeling and Computation*, 7:59–64, 2010.
10. R. Lin, S. Kraus, D. Tykhonov, K. Hindriks, and C. M. Jonker. Supporting the design of general automated negotiators. In *Proc. of ACAN 2009*, 2009.
11. R. Lin, S. Kraus, J. Wilkenfeld, and J. Barry. Negotiating with bounded rational agents in environments with incomplete information using an automated agent. *Artif. Intell.*, 172(6–7):823–851, 2008.
12. T. Mancini. Negotiation exploiting reasoning by projections. In *Proc. of PAAMS 2009*, Advances in Intelligent and Soft Computing, 2009. Springer.
13. J. S. Rosenschein and G. Zlotkin. *Rules of Encounter: Designing Conventions for Automated Negotiations Among Computers*. The MIT Press, 1994.
14. A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97–109, 1982.
15. A. Schrijver. *Theory of linear and integer programming*. John Wiley & Sons, 1998.
16. K. P. Sycara, S. Roth, N. Sadeh, and M. Fox. Distributed constrained heuristic search. *IEEE Transactions on Systems, Man and Cybernetics*, 21(6):446–461, 1991.