Negotiation exploiting reasoning by projections

Toni Mancini

Dipartimento di Informatica, Università di Roma "La Sapienza", Italy
`tmancini@di.uniroma1.it`

Abstract We present a framework that allows two self-motivated distributed agents to perform, in an efficient way, negotiations aiming at achieving mutually satisfactory agreements, when privacy of information is an issue, and no central authority could be used. In particular, each agent has her own constraints to satisfy, as well as her own utility function. Such issues are kept private, and cannot be disclosed to the counterpart. Negotiation is hence carried out by exchanging proposals and by performing sophisticated forms of reasoning on the remote agent's offers, by trying to infer some characteristics of the counterpart, in order to achieve efficient process convergence.

1 Introduction

Automated negotiation among rational agents is an important topic in Distributed Artificial Intelligence, being necessary in several domains, e.g., distributed resource allocation [3], distributed scheduling [11], e-business [9] and, in general, in applications in which: *(i)* no single 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 counterpart or a trusted authority), hence traditional centralized approaches, like Mathematical or Constraint Programming cannot be used. In the literature, various protocols and algorithms for negotiation have been proposed (cf., e.g., [13, 9, 12, 8, 10]), all of which can be classified according to several factors, like the negotiation objects, the agents' decision making models, the degree of cooperation and the level of privacy about each agents' constraints and preferences, or the communication and computation costs.

In this paper, we focus on negotiation between two self-motivated, competitive agents, which aim to find a mutually satisfactory agreement without disclosing their private information (constraints and preferences). Differently from much work existing in the literature, *we explicitly deal with the fact that, in many scenarios, agents face with unknown counterparts, for which no strong assumptions (even probabilistic ones) can be made*: understanding their characteristics and behavior is part of the job of the agent. One of the major prices of this setting is that even the termination of the negotiation process is not guaranteed. To this end, we focus on the local agent only (the only one upon which we have control), which aims at *maximizing* her own private utility function (*competitiveness*) while *guaranteeing termination and efficiency* of the negotiation process. In order to do so, she behaves in a sort of *non-obstructionist way*, and acts very carefully in order to avoid the process to go thrashing and last indefinitely, relying on initial expectations about the reasoning capabilities of the counterpart, but being ready to revise them, by means of a continuous monitoring of her actions, in case a contradicting behavior is observed.

The starting point of our research is the (cooperative) framework described in [1], in which the task of finding a mutually satisfactory agreement is particularly efficient. In such a framework, in which privacy of information is a major concern, negotiation proceeds by exchanging *proposals*. The other agent can either accept the deal or propose a counter-offer. An agent is also able to *reason* on the other agent's proposals. This form of reasoning is called *reasoning by projections*.

In this paper we extend the framework above in the following directions (Sec. 2): *(i)* we provide the agent with the ability to *dynamically estimate and reason on the quality of reasoning made by the counterpart*, and to consequently adapt her behavior *exploiting different heuristics* suitable for the different cases; *(ii)* we provide the agent with the possibility of *defining her own* (private) *utility function* to maximize. Given that the presence of utility functions could make the agent interested in refusing acceptable deals in order to pursue better ones, guarantee of termination can easily be lost. We equip the agent with capabilities that allow her to *still guarantee convergence*, by reasoning on the impact that her refusals have on the counterpart potential reasoning. A *full functional implementation* is briefly described and experimentally evaluated (Sec. 3, details in Appendix[1] due to space reasons).

2 The framework

When dealing with negotiation, there are several aspects that must be taken into account. We address the reader to [1] for a discussion on the most important ones. Here, we only describe the assumptions behind our theoretical framework.

Negotiation framework. When two agents start a negotiation, they already agree on the relevant *variables* (or issues). Following [1], we assume that for each negotiation there is a finite list of *real variables* that are involved. Hence, negotiation spaces can be regarded as multi-dimensional real vector spaces. Agents have their own private *feasibility regions* (R_{loc} and R_{rem}). This setting is different from what is often assumed in the literature, since agents don't even know (or probabilistically estimate) possible counterpart's *types*, *bounds* or *most preferred values* for the variables (in other words, it is not a *split-the-pie* game [5] although with incomplete information [10]). Negotiation proceeds with agents alternatively exchanging *proposals* (aka *deals*), as single points in the space of variable assignments. The counterpart can either accept the deal or decline it, by sending a *counter-offer*.

Characteristics of the local agent. We assume that R_{loc} is *convex* (i.e., all points between two acceptable points are acceptable as well), and more particularly *limited* and defined by means of *linear constraints*, hence it is a *bounded polyhedron*. We also assume that R_{loc} is *stable* during time. Convexity and stability, which play a key

[1] Available online at `http://www.dis.uniroma1.it/~tmancini`.

role in our approach, are very common in many scenarios of practical utility (cf., e.g., [1] and Appendix A.2 for an example). Admissible proposals are in general not equally worth for the agent, that may have her own *utility function*, preferring some solutions to other ones. We assume that the utility function, to be *maximized* by the local agent, is *linear*. Reservation utility can be, wlog, embedded in R_{loc} by additional constraints: hence, any agreement in R_{loc} would be better than failure.

The agent has two *goals*: *(i)* to pursue deals that give her high utility, and *(ii)* to guarantee termination of the negotiation, i.e., either to find a point (agreement) in $R_{\mathrm{loc}} \cap R_{\mathrm{rem}}$, or to logically prove that none exists. Of course, such goals cannot be formally guaranteed, since they depend on the unknown characteristics and behavior of the remote agent. In other words, local agent follows a *heuristic approach*.

The agent is *logically omniscient*, i.e., she is able to compute all logical consequences of the information she has. Moreover, she is *non-obstructionist*, in the sense that she exploits all her knowledge and reasoning capabilities on the behavior of the counterpart in order to keep the negotiation efficient: *(i)* she never proposes a deal that she believes it will be rejected, and *(ii)* never refuses an acceptable offer, if she believes that no better agreement (for her) can be reached. However, given that, in our scenario, the local agent has no information about constraints, goals, preferences, and strategy of the other, this form of collaboration turns out to be much weaker than that which can be exploited in other frameworks, and desirable properties like agreements' Pareto optimality or maximum social welfare cannot be formally guaranteed (but can be reached by a suitable post-negotiation cooperative step, starting from an already mutually satisfiable agreement).

Beliefs about the remote agent. Even if agent has no knowedge about the counterpart, she makes some *initial* (and favorable) assumptions on her, ready to be revised in case an unexpected behavior is observed. In particular she initially assumes that the counterpart: *(i)* Knows that R_{loc} is convex and stable; *(ii)* Has *at least* her own reasoning capabilities (hence is able to reason on the local agent's own behavior, avoiding obstructionism, cf. later), and *(iii)* In turn, believes the same about her (hence, e.g., knows that local agent assumes R_{rem} convex). Local agent maintains and continuously revises a *knowledge base* with the following information: *(i)* All proposals made ($\mathbb{P}$) and offers received ($\mathbb{O}$); *(ii)* A Boolean flag (X) indicating whether R_{rem} is believed conve$X$; *(iii)* A Boolean flag (C) indicating whether the counterpart is believed non-obstructionist (partially Collaborative). Such flags are initially both set to *true*, but will be revised in case of unexpected remote behaviors.

Protocol rules. These are formal rules supposed to always be respected by the agents. In particular, cf. [1], "No cheat" seems to be fundamental in order to guarantee convergence. There (fully collaborative framework and no utility functions), agents cannot propose deals they are not willing to accept, and agree on the first offer that belongs to their feasibility region. Here, given that the presence of utility function makes the local agent interested in rejecting acceptable offers in order to pursue better deals, and given that local agent has no strong guarantees about the counterpart, we relax the "No cheating" rule as follows: *(i)* Local agent never proposes deals she is not willing to accept, and *(ii)* never behaves in ways that would prove herself to be obstructionist or not convex to a counterpart with the same reasoning

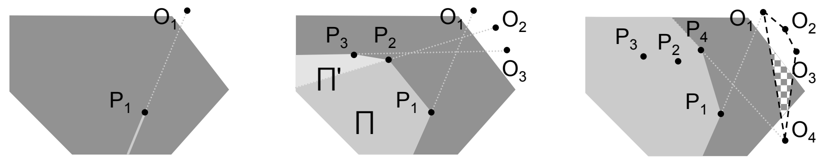

(a) (b) (c)

Fig. 1 (a) Agent's feasibility region along with the first proposal. (b) Regions that can be excluded from proposals after the second (Π) and third (Π') iteration (light areas). (c) 4th iteration: local agent proposes P_4, and receives counter-offer O_4. $R_{\text{loc}} \cap conv(\mathbb{O})$ (textured) now is non-empty: any point there should be good for both agents. Agent may successfully resolve by proposing P_5 there.

capabilities. However, given that the actual characteristics of the remote agent are unknown, local agent cannot enforce symmetric protocol rules to her, and is ready to tune her behavior when the counterpart proves to be less cooperative.

Negotiation with no utility function. In order to describe the form of reasoning used by the local agent in the simplest case, when she has no utility function and believes the same about the counterpart, we show an example concerning a negotiation that involves just two real variables (only to graphically show the ongoing process). R_{loc} is shown in Fig. 1(a); all points are equally worth. Fig. 1(a) also shows the first deal P_1 proposed by the agent, which is not accepted by the counterpart, that in turn offers O_1. Information stored in KB allows the agent to infer new constraints, which add on those of R_{loc}. An obvious example is $\mathbb{P} \wedge \mathbb{O} \vdash (P_1 \notin R_{\text{rem}})$, since the remote agent, offering O_1, implicitly rejects P_1. Interestingly, the local agent, relying on the current assumptions she has about the remote party (initially believed convex), is able to infer that all the points lying on the same line joining P_1 and O_1, below P_1 (highlighted in solid light gray) do not belong to R_{rem}. The reasoning is as follows: *(i)* in case there would be a point $S \in R_{\text{rem}}$ there, by convexity of R_{rem} all points in the segment $\overline{SO_1}$ would belong to R_{rem}; *(ii)* however, since also P_1 lies on $\overline{SO_1}$, a contradiction arises.

As for the next iteration, the local agent proposes P_2, which is again refused by the counterpart, which in turn offers O_2 (cf. Fig. 1(b)). The reasoning that agent can perform at this stage is much more interesting (cf. [1]), and exploits also the initial assumption of non-obstructionism ($C = true$) of the counterpart:

1. By convexity of R_{rem}, all points in $\overline{O_1 O_2} \in R_{\text{rem}}$. Hence, in case $\overline{O_1 O_2} \cap R_{\text{loc}} \neq \emptyset$, she would expect that any point in such set will be acceptable for the counterpart, and would terminate the negotiation process successfully by choosing one of them. So, let us assume that the intersection is empty. The agent may perform additional reasoning in order to exclude, from the set of all possible new proposals, the whole area Π highlighted in medium gray:
   - Since the counterpart offered O_2 after having received P_2, the agent knows that $R_{\text{rem}} \cap \overline{P_1 P_2} = \emptyset$, since also the remote one is assumed to be collaborative ($C = true$), hence should assume convexity of R_{loc} (in absence of a remote utility function, a non-obstructionist remote agent would accept the first acceptable deal, simply because no better agreements are possible);

- Let us assume that there is a point $O \in R_{\mathrm{rem}} \cap \Pi$; from the last assertion and the convexity hypothesis of R_{rem}, R_{rem} must contain all points of the triangle $\widehat{O_1 O O_2}$; now, since such a triangle contains at least one point of the line segment $\overline{P_1 P_2}$, a contradiction arises.

Such form of reasoning can be performed at *any* iteration, in order to reduce the so called *active region*, i.e., the set of all candidate points in R_{loc} for the new proposal. In particular, at any step, when the agent has proposed points $\mathbb{P}$ and received offers $\mathbb{O}$, the choice of next proposal P is made by considering, in order, two alternatives:

(Resolve) If $conv(\mathbb{O}) \cap R_{\mathrm{loc}} \neq \emptyset$, where $conv(\mathbb{O})$ is the convex hull of the received offers, then (under current assumptions) any point in such intersection will be acceptable by both agents. The agent will choose P among them, expecting the successful termination of the process.

(Propose) Otherwise, P is chosen in $R_{\mathrm{loc}} - \Pi(conv(\mathbb{O}), conv(\mathbb{P}))$, where, $\Pi(R, R')$ is the *projection* of region R over R' (defined as the set of all points Q' for which there exists $Q \in R$ such that $\overline{QQ'}$ intersects R'). In case also such set is empty, current assumptions imply that $R_{\mathrm{loc}} \cap R_{\mathrm{rem}} = \emptyset$, hence no agreements exist.

Formal definitions of convex hull and projection are given in Appendix A.1. Fig. 1(b) shows the scenario resulting after points P_3 and O_3 have been exchanged (the whole area Π' colored in light and medium gray can be safely excluded), while Fig. 1(c) shows the situation after the next step, where point O_4 has been received: since now the intersection of R_{loc} with the convex hull of the received offers is non-empty, any point in the textured area is (by current assumptions) a guaranteed good deal.

Such form of reasoning allows to exclude large areas from R_{loc}. In [1], the agent is supposed to propose *vertices*, to guarantee convergence. In fact, assuming a collaborative remote party, if all vertices have been rejected, then no possible agreement can be found (the whole R_{loc} –convex– is excluded from the active region). This can be seen in Fig. 1: if $P_1, \ldots, P_4$ were vertices, the active (dark) region would be much smaller, and convergence faster. However, convergence does not ensure efficiency of termination: in the worst case, a number of proposals proportional to the number of vertices, hence easily exponential in that of the variables is needed to terminate [1].

The question that arises is then the following: *which is the best vertex to propose, in case the alternative* (resolve) *is not feasible?* According to the *fail-first* principle, we equip the local agent with the following heuristic for next vertex proposal:

Best vertex. *The new proposal is the vertex (of R_{loc} and active) which, in case of rejection, will make the highest number of vertices excluded by the active region of the next step.*

This heuristic (that of course cannot take into account the effects, on the exclusion of other vertices, of the counter-offer that will be received upon rejection of the proposal being computed), clearly tends at keeping the negotiation efficient (cf. also Sec. 3), maximally reducing the number of remaining vertices to propose.

Revision of belief about remote agent. Assume we are in the situation of Fig. 1(c), i.e., alternative *(resolve)* is feasible ($conv(\mathbb{O}) \cap R_{\mathrm{loc}} \neq \emptyset$). The agent will propose an arbitrary deal in such intersection, e.g., point P_5 of Fig. 2(a), expecting acceptance. *What if the remote agent refuses this deal?* Unsurprisingly, local agent infers that

R_{rem} is not convex, and revises KB assuming $\neg X$. Consequently, local active region becomes larger, cf. Fig. 2(a): only points in $conv(\mathbb{P})$ can be excluded (the collaborative counterpart would already offered them if possible, because they would be assumed in R_{loc}).

As a second example, assume that, at a given point, e.g., from Fig. 1(b), *remote agent offers a deal belonging to the projection she is supposed to compute*, i.e., $\Pi(conv(\mathbb{P}), conv(\mathbb{O}))$, e.g., offer O_4 in the textured area of Fig. 2(b). Local agent would immediately notice that her counterpart is obstructionist. The reasoning mimics the one that a collaborative counterpart would perform: *(i)* A collaborative remote agent would believe that the local one has a convex region, because she never refused a point O_i in $conv(\{P_1, \ldots, P_{i-1}\})$, and is collaborative, because she never behaved in contradiction with such hypothesis. *(ii)* Hence, she would infer that, if O_4 was acceptable, so would be all points in $conv(\{O_4, P_1, P_2, P_3\})$. *(iii)* Since by construction, such a region has non-empty intersection with $conv(\{O_1, O_2, O_3\})$, the remote agent would expect that the collaborative local one should have already proposed a deal in such intersection, which is not the case. *(iv)* Hence, the remote agent cannot be collaborative (since she doesn't properly assume convexity and collaborativeness of the local agent that have not been violated so far). The local agent thus revises her KB, inferring $\neg C$. Also in this case the active region becomes different, cf. Fig. 2(b). Two more scenarios may arise that let the agent discover violations, by the counterpart, of the convexity or collaborativeness assumptions, one of which is shown in Fig. 2(c). Due to space reasons, they are discussed in Appendix A.3.

In case the counterpart is believed to be either not convex or not collaborative, Best vertex heuristic loses its main advantages: if the remote agent is not convex, the local one cannot exclude any projection, but only the convex hull of her own proposals ($conv(\mathbb{P})$, cf. Fig 2(a&c), light areas). If not collaborative, proposing vertices does not even guarantee termination (although, if still believed convex, the agent could exclude points in $\bigcup_{P \in \mathbb{P}} \Pi(conv(\mathbb{O}), P)$, cf. Fig. 2(b)). To this end, focusing on the case of a non-convex but collaborative remote party, in principle proposing either vertex is equally worth. However, by following the intuitive approach of trying to meet as much as possible the counterpart's last offer we devise a second heuristic:

Closest vertex. *The new proposal is the vertex (of R_{loc} and active) closest to the last offer.*

Since the remote agent is believed collaborative, when all vertices of R_{loc} (bounded polyhedron) have been proposed, the active local region becomes empty. In case of a non-collaborative remote agent, no guarantee of termination is possible. Hence, the local agent could either *(i)* abort the negotiation; *(ii)* propose random points in her active region (**Any point** heuristic) or, as suggested in [4], *(iii)* approach the opponent last offers proposing internal points of R_{loc}. This last heuristic (which mimics the approach typical of local search) is claimed to work well in several cases.

It is worth noting that, every time the local active region becomes empty, the agent has a proof that, if her current assumptions are valid, then no agreement actually exists. But of course, it might be the case they are wrong, but have never been violated by the counterpart so far. As above, the local agent may in principle decide either *(i)* to abort the negotiation (trusting in her assumptions), or *(ii)* to deliberately relax them in order to have additional deals to propose. However, in the latter case

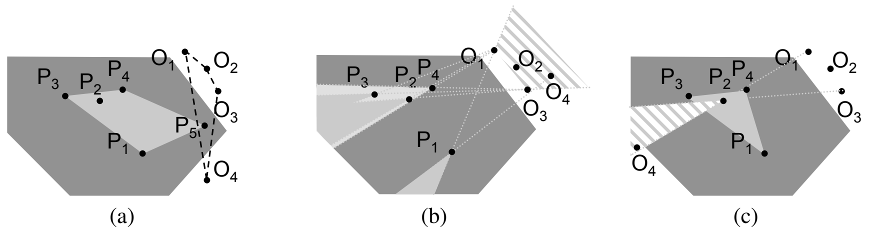

Fig. 2 (a) 5th iteration: Remote agent refuses $P_5 \in R_{\text{loc}} \cap conv(\mathbb{O})$, making a counter-offer (not shown). The local agent infers $\neg X$. (b) 4th iteration: Remote agent offers O_4 in $\Pi(conv(\mathbb{P}), conv(\{O_1, O_2, O_3\}))$ (textured area, unlimited). Local agent infers $\neg C$. (c) 4th iteration: Remote agent offers O_4 in $\Pi(conv(\mathbb{O}), P))$ for some $P \in \mathbb{P}$ (textured area, with $P = P_2$, cf. Appendix A.3). Local agent infers $\neg X$.

she may herself appear obstructionist (if her assumptions were indeed valid). Also, switching to $\neg C$ would make any guarantee of convergence lost.

Negotiation in presence of utility function. Heuristics above might behave very bad in case an utility function does exist (e.g., it would be not hard to figure out situations in which the Best vertex is the point with lowest utility). However, facing with utility functions is problematic in general, since the guarantee of convergence can be easily lost on the way, because agents may have interest in rejecting acceptable offers, hoping to increase their utility (in the following, we call these actions *lies*, because agents are deliberatively making the negotiation longer than that it could be). In order to still guarantee termination, lies should have no negative impact in the effectiveness of the counterpart's potential reasoning. In particular, we assume that local agent always acts in order to *appear* (rather than to be, given the presence of an utility function) convex and non-obstructionist: she will never refuse a deal if such action would give the counterpart enough information to infer she is not convex. On the other hand, as for other offers and her own proposals, she will follow more careful strategies, deeply exploiting, by additional reasoning, what privacy of information guarantees, i.e., the potential discrepancy between her actual behavior and that *perceived* by the counterpart.

In particular, when rejecting an acceptable offer O may have an impact on process convergence? The answer is pretty simple, if we consider what a collaborative counterpart could infer. Rejecting $O \in conv(\mathbb{P})$ would give to remote agent evidence that the local one is *not* convex (cf. Fig. 2(a), reversing the roles). Hence, such offers *cannot be rejected*, in order to maintain guarantee of convergence. [2] In the other cases ($O \in R_{\text{loc}} - conv(\mathbb{P})$), acceptable offers may safely be rejected (we call these: *safe* lies). However, in such situations, other constraints must be taken into account by local agent when making subsequent proposals in order to "lying with impunity" and still appear convex and collaborative: *(i)* The agent must never propose a point P in $\bigcup_{O \in \mathbb{O}} \Pi(conv(\mathbb{P}), O)$, because this would show to counterpart that (rejected) $O \in conv(\mathbb{P} \cup \{P\})$, hence R_{loc} is not convex; *(ii)* In case the counterpart can be believed as convex, the agent must never propose a point P in $\Pi(conv(\mathbb{P}), conv(\mathbb{O})) - (conv(\mathbb{O}) - conv(\mathbb{O} - \{O_{last}\}))$, with O_{last} being the last received offer (proof in Appendix A.4).

---

[2] This is not a big deal, since, by construction, $u_{loc}(O) \geq u_{loc}(P)$ for some past proposal $P \in \mathbb{P}$.

In order to allow the agent to autonomously decide whether to lie or not, an *acceptance policy* can be defined for her. Following (and slightly generalizing) the main approach [5] of accepting an offer if *better* than the counter-proposal the agent would make next, we define an acceptance policy as a real $\xi \in [0..1]$: offer $O \in R_{\text{loc}}$ would be accepted iff $u_{loc}(O) \geq u_{loc}(P_{next}) - span \cdot \xi$, where *span* is the absolute difference of the extreme values of $u_{loc}()$ in R_{loc}, and P_{next} is the point that would be chosen by following alternative *(propose)*. Thus, acceptance policies vary between two extremes: *(i)* Any acceptable offer is accepted ($\xi = 1$ –this neutralizes the need for lies and minimizes the negotiation length); and *(ii)* An acceptable offer is accepted iff its utility is higher than that of the counter-proposal P_{next} that would be chosen next, upon rejection ($\xi = 0$).

As for the overall strategy in presence of a linear utility function $u_{loc}()$ to maximize, the agent maintains two distinct feasibility regions, R_{loc} (the real one), and $R^k_{\text{loc}} = R_{\text{loc}} \cap \{P : u_{loc}(P) \geq k\}$ (with k varying during the process). A linear $u_{loc}()$ guarantees that also R^k_{loc} is a bounded polyhedron. The agent will use $R_{\text{loc}} \cap \{P : u_{loc}(P) \geq u_{loc}(P_{next}) - span \cdot \xi\}$ to check whether an offer is acceptable or for resolving a negotiation (according to the acceptance policy), and, as for alternative *(propose)*, she will search for deals in R^k_{loc} only.

1. At the beginning, k is fixed to the highest utility value: $k = \max(u_{loc}(P) : P \in R_{\text{loc}})$.
2. The local agent proceeds as in the case with no utility function (using R_{loc} and the acceptance policy as for alternative *(resolve)* –as long as $X = \mathit{true}$– and R^k_{loc} as for *(propose)*, taking also care of her lies), until one of two conditions holds: *(i)* the active region is empty; or *(ii)* the counterpart proves to be not collaborative.
3. In case *(i)*, the agent has a proof that no agreements (modulo safe lies on remote side) can be found in R^k_{loc}. Hence, she will perform a *utility concession* to the counterpart, by lowering k by a given $\delta > 0$, hence considering $R^{k-\delta}_{\text{loc}} \supseteq R^k_{\text{loc}}$. This will produce (in case of $\supset$) new vertices to propose. Negotiation aborts when $R^k_{\text{loc}} = R^{k-\delta}_{\text{loc}} = R_{\text{loc}}$, meaning that no agreements exist at all (modulo lies at both sides, and except for wrong assumptions, dealt with by unilateral relaxation).
4. As with no utility function, in case *(ii)*, any guarantee of termination is lost (the agent cannot count on the capabilities of the remote one). Also here, she can in principle proceed in different ways, but needs to have a mechanism for conceding in terms of utility even if the current active region is non-empty (to avoid to remain stuck forever with unfruitful proposals). As an example, she can decrease k every a fixed number of steps, or immediately reduce herself to the case with no utility function. The efficacy of either option is expected to strongly depend on the application domain, hence it will not be considered here.

From point 2 it can be observed that, in order to lie, it is not necessary to reject an acceptable offer. Consider the situation in Fig. 1(c). The local agent believes, by the convexity assumptions about the counterpart, that all points in $conv(\mathbb{O}) \cap R_{\text{loc}} \neq \emptyset$ are acceptable deals. To behave collaboratively, she should propose there. However, she can safely lie, proposing in a different region, as long as the counterpart has no evidence of this non-collaborative behavior.

The overall strategy causes a quantization of the utility function: for a given k, all points with utility between k and $k + \delta$ (δ being the *quantization value*) are consid-

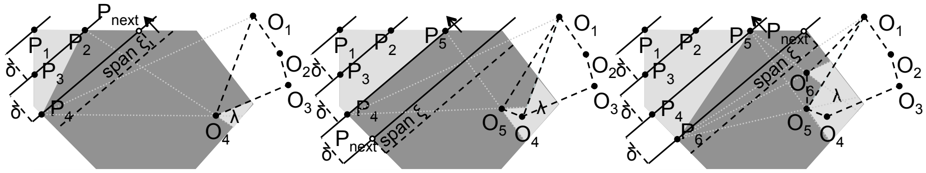

(a) (b) (c)

Fig. 3 Negotiation in which the agent has a linear utility function to maximize (black solid line, utility growing North-West), exploits Best vertex heuristic, and has an acceptance policy $\xi < 1$ (never accepting deals below the dashed line). (a) 4th step, she lowers her utility threshold to $U_{max} - 2\delta$, having a proof that no agreement can be found in $R_{\text{loc}}^{U_{max}-\delta}$ –previous steps not shown. The best vertex is P_4, which is rejected by the counterpart that in turn offers $O_4 \in R_{\text{loc}}$. Local agent may safely refuse it, as long as she still appears convex and collaborative, not proposing in $\lambda = \bigcup_{O \in \mathbb{O}} \Pi(conv(\mathbb{P}), O) \cup (\Pi(conv(\mathbb{P}), conv(\mathbb{O})) - (conv(\mathbb{O}) - conv(\mathbb{O} - \{O_{last}\})))$, O_{last} being the last offer). (b)&(c) 5th and 6th steps, with additional lies.

ered equally worth. This is important to preserve convergence (the process reduces to a repeated but finitely long application of the basic strategy), and results in δ that should be given explicitly by the user, since, as ξ, it intrinsically depends on the application domain: the lower are δ and ξ, the longer could be the process; moreover, the lower is ξ, the higher is the risk that the process fails even if agreements exist, because all the ones found have been rejected by lies. An important property of this strategy is that *all the reasoning made in the previous steps* (in terms of reduction of the active region) *remains valid*: a proposal unacceptable for the counterpart will remain such when lowering k (as long as belief about the counterpart doesn't change). This preserves efficiency. Fig. 3 shows a fragment of a negotiation process where the local agent has an utility function to maximize, proposes vertices (as required to guarantee termination), and exploits safe lies.

3 Experiments and perspectives

In order to assess the effectiveness of this approach, we built a Java P2P application fully implementing the present framework, and developed a random generator to produce a great number of negotiation scenarios (consisting of pairs of random polyhedra with controllable probability of intersection, and random utility functions). Experiments show (detailed results in Appendix A.6 and summarized in Table 1) that iAgree is able to deal with quite complex negotiation instances (e.g., 3 variables, 45 vertices per side) in reasonable time, that Best vertex plays a role in efficiency (saving up to 18% in terms of nb. of rounds wrt Closest vertex), and that moderately loose acceptance policies ($\xi \sim 0.4$) may lead to both high success ratios (existing agreements are actually found in $\sim$95% of the cases) and agreement quality (close to optimum of $\sim$75% in terms of utility).

The effectiveness of the approach let many important directions for future work emerge: *(i)* investigating the possibility of handling variables on discrete domains (potentially exploiting the current form of reasoning to perform suitable relaxations)

Average results over 200 instances	$\xi_{loc}=\xi_{rem}=0$	$\xi_{loc}=\xi_{rem}=.4$	$\xi_{loc}=\xi_{rem}=1$	$\xi_{loc}=0, \xi_{rem}=1$	$\xi_{loc}=1, \xi_{rem}=0$
Number of rounds	25	27	34	29	28
Reasoning time (local agent, sec)	27	30	35	30	30
Success ratio	∼60%	∼95%	100%	∼75%	∼85%
Agreements' quality (local agent)	∼65%	∼75%	∼65%	∼83%	∼48%
Rounds saved by BestVtx wrt ClsVtx	0%	6%	7%	18%	4%

Table 1 Some results for 3 vars (#vertices in 16..45, ∼35 on average). 200 random instances solved for each combination of ξ_1, ξ_2 in $\{0, .2, .4, .6, .8, 1\}$. Success ratio = #agr. found/#sat inst. Agr. quality = $(u_{loc}(agr.) - min_{P \in R_{\mathrm{loc}} \cap R_{\mathrm{rem}}} u_{loc}(P))/(max_{P \in R_{\mathrm{loc}} \cap R_{\mathrm{rem}}} u_{loc}(P) - min_{P \in R_{\mathrm{loc}} \cap R_{\mathrm{rem}}} u_{loc}(P))$.

and non-linear constraints and utility functions (cf., e.g., [7]); *(ii)* devising new strategies and heuristics; *(iii)* dealing with deadlines [6] (e.g., by a dynamic handling of δ and ξ); *(iv)* dealing with negotiations among more than two agents, extending the reasoning capabilities of the agent to cope with multiple counterparts; *(v)* dealing with non-convex regions, e.g., *union* of polyhedra, also exploiting local search and hybrid techniques [4, 7].

Acknowledgments. The author thanks co-authors of the preliminary version of iAgree [2]: Marco Cadoli (who introduced, already in [1], projections as a mean to perform reasoning in cooperative negotiations) and Guido Chella (who implemented the first version of the system during his Bachelor thesis).

References

1. Cadoli, M.: Proposal-based negotiation in convex regions. In: Proc. of CIA 2003, *LNCS*, vol. 2782, pp. 93–108. Springer (2003)
2. Cadoli, M., Chella, G., Mancini, T.: iAgree: a system for proposal-based negotiation among intelligent agents. In: Proc. of ECAI 2006. IOS Press (2006). System demonstration.
3. Conry, S.E., Kuwabara, K., Lesser V. R. Meyer, R.A.: Multistage negotiation for distributed constraint satisfaction. IEEE Trans. on Systems, Man and Cybernetics **21**(6), 462–477 (1991)
4. Costantini, S., Tocchio, A., Tsintza, P.: Experimental evaluation of a heuristic approach for P2P negotiation. In: Proc. of RCRA 2007 (2007)
5. Faratin, P., Sierra, C., Jennings, N.R.: Using similarity criteria to make issue trade-offs in automated negotiations. Artif. Intell. **142**(2), 205–237 (2002)
6. Fatima, S.S., Wooldridge, M., Jennings, N.R.: Multi-issue negotiation with deadlines. J. of Artif. Intell. Research **27**, 381–417 (2006)
7. Ito, T., Hattori, H., Klein, M.: Multi-issue negotiation protocol for agents: Exploring nonlinear utility spaces. In: Proc. of IJCAI 2007, pp. 1347–1352. Morgan Kaufmann (2007)
8. Jennings, N., Faratin, P., Lomuscio, A., Parsons, S., Wooldridge, M., Sierra, C.: Automated negotiation, prospects, methods and challenges. Group Dec. and Negot. **10**, 199–215 (2001)
9. Jennings, N.R., Norman, T.J., Faratin, P., O'Brien, P., Odgers, B.: Autonomous agents for business process management. Applied Artif. Intell. **14**(2), 145–189 (2000)
10. Lin, R., Kraus, S., Wilkenfeld, J., Barry, J.: Negotiating with bounded rational agents in environments with incomplete information using an automated agent. Artif. Intell. **172**(6–7), 823–851 (2008)
11. Sycara, K.P., Roth, S., Sadeh, N., Fox, M.: Distributed constrained heuristic search. IEEE Trans. on Systems, Man and Cybernetics **21**(6), 446–461 (1991)
12. Yokoo, M., Katsutoshi, H.: Algorithms for distributed constraint satisfaction: A review. Autonomous Agents and Multi-Agents Systems **3**(2), 185–207 (2000)
13. Zlotkin, G., Rosenschein, J.S.: Mechanisms for automated negotiation in state oriented domains. J. of Artif. Intell. Research **5**, 163–238 (1996)