Control and Cybernetics

VOL. 21 (1992) No. 1

Modelling Negotiation Strategies with Two Interacting Expert Systems[1]

by

Zbig Koperczak

Stan Matwin

Stan Szpakowicz

Department of Computer Science
University of Ottawa
Ottawa, Ontario
Canada K1N 6N5

We evaluate the effect of different negotiating strategies on the convergence of a negotiation process with a clearly defined quantitative aspect. We present an experimental strategy testing environment, designed within the framework of the decision-support system NEGOPLAN. We have extended NEGOPLAN to facilitate measurements of several different characteristics of the negotiation process. The models of a trade union and a management of a company have been expressed in NEGOPLAN and pitted against each other. One of the interesting experimental results is that a strategy consisting of moderate concessions serves interests of the union the best, while either a decidedly firm or a soft stance is beneficial to the management.

[1]This work was partially supported by the Natural Sciences and Engineering Research Council of Canada.

1. Introduction

The experiment that we present in this article was designed within the framework of the NEGOPLAN system (Kersten et al. 1991; Michalowski et al. 1991; Kersten et al. 1990; Kersten and Szpakowicz 1990; Matwin et al. 1989; Kersten et al. 1988). This is a decision-support system that had been used, in particular, to model one side in negotiation–one negotiating *agent*. In the work described here, we have used NEGOPLAN to evaluate the effect of negotiating strategies on the convergence of a negotiation process with a clearly defined quantitative aspect. A model of a trade union and a model of a management of a company have been expressed in NEGOPLAN and pitted against each other.

NEGOPLAN helps model two dimensions of a negotiation problem: the goals of the negotiating parties, and the dynamics of the negotiation process, during which the goals may change. One of the fundamental assumptions that underlie NEGOPLAN states that a decision problem can be hierarchically decomposed into subproblems. In the case of negotiation, there is a principal goal (normally, shared by both sides: to reach a compromise) that can be broken down into subgoals. The decomposition process stops at elementary subproblems, referred to as *facts*, that are directly verifiable or can be treated as common knowledge.

Decomposition results in a tree-like structure that we term *goal representation*. Its leaves are facts; in negotiation, leaves describe the actual issues on the negotiating table. A problem (a goal or a fact) is expressed by a predicate, and the tree can be interpreted as a logical formula in which tree nodes represent AND and OR operators. If each leaf is associated with a truth value (with true treated as "achieved" and *false* as "not achieved"), the principal goal's value is determined. An assignment of truth values that satisfies the goal is referred to as a *goal solution*. We note that a goal solution may only assign a truth value to selected leaves; the remaining leaves represent those issues whose final values are irrelevant, because the goal will be satisfied regardless of their truth values. In the experiment presented in this paper, all the choices (OR nodes) have been removed, as the focus was on the influence of various negotiating strategies on the payoff in the negotiation process.

A fact annotated with its truth value is referred to as a *metafact*. A negotiating position is composed of metafacts. Reactions to positions are modelled by means of *metarules*. A metarule specifies a number of metafacts whose presence in the knowledge base triggers the reaction, as well as tests and actions that must be performed on the parameters of facts before the metarule is applied.

An application causes the addition of new metafacts to the knowledge base.

In actual union-management negotiations, issues are typically grouped in pairs to facilitate the bargaining process; see the notion of "reciprocity of concessions" in Ikle (1964). In our model, we associate a single number with each issue. A hierarchical decomposition of the whole problem only helps describe the knowledge about both sides. Facts is what really matters, and each fact summarizes one issue by associating a symbolic label that names the issue with a number (this is a significant restriction of the normal expressive power of NEGOPLAN). In effect, a problem is seen as a sequence of facts, and each fact as a number: the problem really is a vector of numbers.

The paper has four more sections. First, we define the basic notions of our model. We then describe our case study, the union-management negotiations, and we introduce elements of the NEGOPLAN notation. Next, we discuss the experiments and results. Finally, we draw a few conclusions.

2. Assumptions and Definitions

2.1. Limitations

To enable experiments with strategies, we have slightly extended NEGOPLAN. On the other hand, in order to be able to measure, react to and compare the consequences of different strategies, we have also introduced a few limitations.

A. Predicates in the leaves have a special form that allows us to assign a numeric value to every negotiating issue of either side. As a result, a negotiating position that normally consists of a set of truth-valued facts can be reduced to a vector of numerical values. This simplification seems justified in view of our emphasis on testing negotiating strategies. The predicate `percentage_change(ITEM, AGENT, PCT)` reads "the AGENT should change the ITEM by PCT percentage points". For example,

   ```
   percentage_change(robot_use, management, 25)
   ```

   means that robotization of the company operations should increase by 25 percent. The demand

   ```
   percentage_change(absenteeism, union, -10)
   ```

   means that the union should decrease absenteeism by 10 percent.

B. Each side ranks its issues.

C. There exists a metric that allows us to compare positions, and to measure the distance between positions. The metric converts a vector of issue values into a single value; it uses ranks as weights that determine how the individual issues contribute to the single value.

D. The agents never change values in such a way that the distance between them would grow.

E. We assume that the agents have a compatible set of demands (that is, few if any issues do not appear in the positions of both) but may have different ranges of values.

F. Both agents negotiate the same issues at the same time. This assumption implies that in our model the negotiating parties honour partial agreements: if both sides agree on a point, they will not go back to discuss it later (Ikle 1964, p. 100).

2.2. Definition of a Strategy

The work in the framework of game theory (Axelrod 1987), quite different in spirit from our approach, assumes the existence of an objective, well-defined payoff function for *each* issue, as well as the identity of sets of issues being discussed. In our case, the payoff is only defined for the final result of the negotiation, and sets of issues must intersect rather than coincide. Early attempts to view negotiation as an interaction of two final automata (England 1979) are also much less comprehensive than our approach. Rather than settling for a fixed set of states of a finite automaton, we introduce a modifiable set of positions whose components (metafacts) can also be dynamically modified during the negotiation process. Instead of describing a bargaining behaviour with a finite transition function, we rely on logical inference with unification as the underlying rule–application mechanism

The objective of this work is to provide an experimental testbed where one could test various negotiation strategies and their mutual effects in a two-party negotiation. Therefore, one of the basic notions to define is that of a *negotiating strategy*. A review of the available literature on conflict, negotiation and bargaining (Leng and Walker 1982; Leng and Wheeler 1979; Oscamp 1971; Patchen 1987; Touzard (1977)) has shown that the existing definitions of a strategy are too conceptual and not precise enough to provide a specification that could be implemented in a system such as NEGOPLAN. Consequently, we have come up with our own definition which is, nevertheless, anchored in the existing work. Our starting point was Simon's view of what constitutes a strategy: "The series of decisions which determines behaviour over some stretch of time may be called a strategy..." (Simon, 1976, p. 67). A more specific definition is proposed by (Hamner and Yukl 1977): "A bargaining tactic can be defined as a position to be taken or a maneuver to be made at a specific point in the bargaining process. A bargaining strategy, on the other hand, consist of a series of bargaining tactics to be used throughout the bargaining process."

Negotiation is seen as movement between positions, where a position is a vector of values. The overall objective of both sides is to reduce the distance between their positions, that is, to converge. A threshold exists to tell us that two positions are satisfactorily close. Each side has at its disposal several operations that adjust values according to a strategy and tactic. There may be a move in the direction of the other side's offer, or the value is kept unchanged.

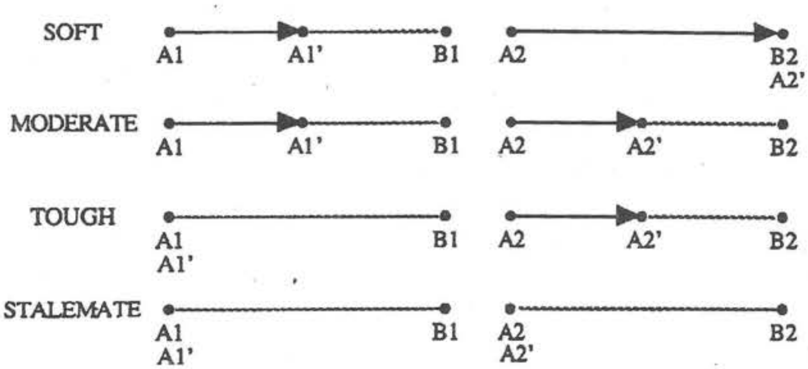

Figure 1. Illustration of concessions in the four basic negotiation tactics

We distinguish four *tactics*, where a tactic is a method of making such numeric changes: soft, moderate, tough, stalemate. The changes are illustrated in Figure 1. A1, A1' are one side's old and new value for the more highly ranked issue, B1 is the other side's value for the same issue. A2, A2' and B2 are the analogous values for the other issue. The picture illustrates movement or concessions. For instance, in the SOFT tactic, a movement toward the other side's position is made on the more highly ranked issue, while the other side's position is adopted completely for the other issue.

Negotiation consists of steps; the negotiation time is measured in cycles of the underlying NEGOPLAN machinery. A side has a tactic at each step, and it stays with this tactic so long as there is movement toward convergence. A new tactic is selected whenever there has been no movement for some time (that is, for a number of steps). A method of selecting new tactics that is consistent over time is a *strategy*. We define a strategy as a *sequence* of tactics. This general definition allows us to represent intuitive negotiating strategies such as, for example, "tit for tat": the next element of our strategy is the same as the previous element of the opponent's strategy.

3. The Case Study: Union-Management

3.1. The Quantitative Aspects of the Models

In the previous experiments with NEGOPLAN we had built qualitative models to simulate negotiations. This was adequate for negotiations with a terrorist or for Camp David talks. It was slightly less appropriate in management–union negotiations where quantitative values had to be represented with qualitative symbols. In this experiment, we represent issues of an agent as

```
percentage_change(ITEM, AGENT, PCT)
```

(the AGENT should change the ITEM by PCT percentage points). Examples are shown in section 2.1. The percentage values are changed by metarules, as shown in section 3.3.

All facts are assigned a truth value, and in most cases this value is used by NEGOPLAN. For example, the facts `company(bankrupt)` and `go_on_strike (union)` may be false: this means the situations they describe are not desirable. As a special modification of NEGOPLAN, the truth value is ignored for `percentage_change`.

We assume that the agents have a compatible set of demands (that is, few if any issues do not appear in the positions of both) but different ranges of values. We have divided the set of all demands into categories: job security, remuneration, benefits and safety. Each category corresponds to a phase of the negotiations. We assume that the participants negotiate issues in pairs. For example, an increase in overtime pay is negotiated together with the issue of absenteeism. The agents negotiate issues in one category until an agreement has been reached on all of them, and then they move on to another category.

Negotiation is over when both agents have reached an agreement on all issues in all categories.

Each agent represents the negotiation goals as a tree. The leaves are the issues which are being discussed. An agent orders the facts in strictly descending order according to the ranking of issues. We assume that a negotiator who is dealing with two issues of unequal importance will be lenient on the less important issue and tough on the more important one.

The notion of distance between two positions is essential for our system. According to assumption C from section 2.1, such a distance can be computed. In section 4 we show in detail how this has been done in our experiment. The distance is defined only for demands of the form

```
percentage_change(ITEM, AGENT, PCT)
```

If an issue is absent from one of the positions, the distance is 1. Otherwise an issue must be present in both positions, usually with different percentage values, and the distance is defined as the normalized Euclidean distance between these numeric values.

The agents have different negotiation strategies and different initial values of the negotiation issues. During negotiation the values of the issues change, and the strategy decides the pace of change. A strategy is, in effect, a sequence of compromise formulæ. We assume four kinds of such formulæ: *soft, moderate, tough, stalemate*, in the order of increasing inflexibility (Hamner and Yukl 1977). See Figure 1 for a graphical illustration. A participant decides on a strategic change from one formula to another if the weighted sum of concessions that has been made since the negotiation started is greater than the analogous sum for the opponent. For example, an agent may start with the moderate compromise formula. If during the negotiation her total concession is substantially higher than that of the opponent, there will be a change of the compromise formula to tough. If the difference still does not decrease, the agent may decide to change the formula to stalemate. Later, if the difference diminishes, the moderate formula may be reused.

The strategy choice affects the strategy parameters used to calculate the numerical values of the issues. We have introduced the notions of *compromise factor, compromise threshold*, and *compromise margin*. The factor is used to calculate the agent's concession on an issue. If it is 1, we give in on this issue (we agree to the opponent's demand). If it is 2, we get the arithmetic average. The greater the factor, the smaller the concession. The threshold determines when a participant will ignore the difference between the issue values, that is, when a decision will be made that the positions are so close on this issue that an agreement is in order. This is important because negotiations can move to the next phase only if the negotiators have agreed on all the points discussed in the given phase. Finally, the margin indicates how flexible is the negotiator.

3.2. Goal Representation

The management's goals are represented by the initial goal tree. Its leaves are the demands presented to the union side. The management's goals are a combination of its "personal" goals and those concerning the company. ('The 'gather_info' subgoal has only a technical significance. It helps collect the necessary information about the outcome of negotiation.)

```
goals(management) <-
      company_goals &
      personal_goals &
      gather_info.
```

The management wants to avoid the company's bankruptcy:

```
personal goals <-
      not company(bankrupt).
```

The management will use upcoming labour negotiation to change the situation by committing the union to higher productivity, by increasing profits and lowering spending. It wants to avoid the union's interrupting the talks and disrupting the company's operations by going on strike:

```
company_goals <-
      increase_profits &
      lower_spending &
      increase_productivity &
      keep_company_operational.
keep company operational <-
      not go_on_strike(union).
```

The company can become more profitable if the number of work hours per week will be increased, part of the work force will be laid off, part of the factory robotized, and some of the operations moved to the areas with lower labour costs. The management sets its demands much higher than the required level because this will give it an extra edge later on during negotiation (Ramberg 1977).

```
increase_profits <-
      percentage_change(work_time, management, 10) &
      percentage_change(layoff, union, 20) &
      percentage_change(robot_use, management, 25) &
      percentage_change(relocate, management, 25).
```

The management wants to increase productivity by raising work quota, reducing material waste, increasing the quality of the product, and reducing absenteeism. Again, the demands are much too high. The management also sees safety improvements and increased workers' participation as a means of increasing productivity.

```
increase_productivity <-
      percentage_change(quota, union, 10) &
      percentage_change(material_waste, union, -10) &
      percentage_change(quality, union, 10) &
      percentage_change(absenteeism, union, -10) &
      percentage_change(safety_improvement, union, 20) &
      percentage_change(improvements_suggestions, union, 20).
```

The management wants to reduce expenditures by not contributing to the employees' dental plan or pension plan, and it does not want the company to subsidize the employees' home loans. It wants to cut costs by reducing the hourly wage and overtime pay. Again, we can interpret this as a tactical move to gain room for manoeuvre during negotiations. The offer of the company contribution to the university plan is generous; we can interpret it as a long term investment in human resources.

```
lower_spending <-
     percentage_change(dental_plan, management, 0) &
     percentage_change(pension_plan, management, 0) &
     percentage_change(child_university_plan, management, 50) &
     percentage_change(subsidized_mortage, management, 0) &
     percentage_change(hourly_wage, management, -10) &
     percentage_change(overtime_pay, management, -20).
```

The union's goals are represented similarly. The complete tree, not shown here because of the size of the screen image, is reproduced in Koperczak et al. (1991). Most issues (that is, leaves) are the same as in the management's tree, although the decomposition is quite different.

Issue	Union	Management
absenteeism	0	−10
child_university_plan	25	50
dental_plan	100	0
hourly_ wage	20	−10
improvements_suggestions	0	20
layoff	0	20
material_waste	0	−10
overtime_pay	50	−20
pension_plan	100	0
quality	0	10
quota	0	10
relocate	0	25
robot_use	0	25
safety_improvement	100	20
share_plan	—	25
subcontracting	−20	—
subsidized_mortgage	3	0
work_time	−10	10

Table 1. Initial positions of the union and the management

3.3. Negotiation

Negotiation starts when the union presents its position to the management. In real life this happens when the contract is up for renewal and the management is presented with a written list of demands. In our experiment, the system representing the union writes its demands, in the form of metafacts, to a file. The demands are represented as elements of the goal solution. The union's position as described by the goal representation contains (apart from a few metafacts with a purely technical role) a number of `percentage_change` metafacts, and two metafacts with significant truth values, namely *false*: `company(bankrupt)` and `go_on_strike(union)`. The management's response contains analogous metafacts. We summarize the positions of both sides in Table 1.

The management then reads the union's position from a file, and begins preparing an answer to this position. In real life, the management does not work in a vacuum but must consider the environmental parameters such as settlements reached in the same or similar industries (the so-called *parity*), the size of the workforce, the average wage, dental insurance rates and the average pension plan contribution. It should also be concerned with the average costs and benefits of robotization, relocation and layoffs. This *objective* information is available to both sides who may use it to calculate the payoffs. (It must be noted that payoffs are only used to show the results of a negotiation but they do not bear on decisions. Also, the management's win need not mean the union's loss, and *vice versa*.) Here is an example of objective information:

```
company(hourly_wage, 7)
company(overtime_pay, 100)
company(overtime_rate, 0.15)
company(work_time,  40)
company(work_weeks, 50)
company(workforce, 1000)
rates(child_university_plan, 150)
rates(dental_plan,  100)
rates(house_price,  150)
rates(new_tools, 100)
rates(pension_plan, 500)
rates(relocate_cost, 200)
rates(relocate_gain, 500)
rates(robot_cost, 100)
```

```
rates(robot_gain,  300)
rates(safety, 25)
rates(safety_cost, 25)
rates(safety_gain, 15)
rates(university_attendance, 0.1)
rates(waste_cost, 100)
rates(num_of_dependents, 3)
phase(management, start)
```

This also includes information about the negatiation phase. At the beginning the negotiation is in the 'start' phase. We assume both side negotiate in phases and each phase takes its name from the group of demands being discussed. For example, when the union's demands on hourly rate and overtime pay are discussed, we say that the negotiations are in the remuneration phase. There are six phases in our labour–management negotiations: start, job security, safety, remuneration, benefits and final. Both sides move to the final phase if negotiations have been successful. In this phase an agreement is signed and the negotiations terminate.

We simulate the bargaining process by performing forward chaining on facts that represent both positions and the objective information. In NEGOPLAN we describe reactions to positions by means of *response metarules* of the following form:

$\texttt{existing\_metafact}_1$ &
....
$\texttt{existing\_metafact}_n$ &
{...............}
==>
$\texttt{new\_metafact}_1$ &
....
$\texttt{new\_metafact}_m$

A metarule is applied if all the existing metafacts have been found in the knowledge base, and the tests and actions contained in the "window" (that is, within curly brackets) have been performed. An action might be the calculation of a new value of an issue. For a thorough discussion of "windows", see Kersten et al. (1990).

As an example, we show a metarule that updates the values of the issues *hourly wage* and *quota*:

```
phase(union,remuneration) ::= true &
threshold(union,THR) ::= true &
tactic(union,TACT) ::= true &
factor(union,FAC) ::= true &
management: percentage_change(hourly_wage,management,M_HRLY_WAGE)
                                                        ::= true &
union: percentage_change(hourly_wage,management,U_HRLY_WAGE) ::= true &
management: percentage_change(quota,union,M_QUOTA) ::= true &
union: percentage_change(quota,union,U_QUOTA) ::= true &
{
     get_rank(percentage_change(hourly_wage,management,_),
                                         U_HRLY_WAGE_RANK),
     get_rank(percentage_change(quota, union,_), U_QUOTA_RANK),
     (  (
           U_QUOTA_RANK < U_HRLY_WAGE_RANK,
     compr_formula(TACT,U_QUOTA,M_QUOTA,U_NEW_QUOTA_TMP,FAC,
                U_HRLY_WAGE,M_ HRLY_WAGE,U_NEW_HRLY_WAGE_TMP)
        );
        ( U_QUOTA_RANK >= U_HRLY_WAGE_RANK,
     compr_formula(TACT,U_HRLY_WAGE,M_HRLY_WAGE,U_NEW_HRLY_WAGE_TMP,
             FAC,U_QUOTA,M_QUOTA,U_NEW_QUOTA_TMP)
        )
     ),
     calculate_distance(U_NEW_HRLY_WAGE_TMP,M_HRLY_WAGE,DDP),
     ( (DDP =< THR,U_NEW_HRLY_WAGE is M_HRLY_WAGE) ;
       (DDP > THR,U_NEW_HRLY_WAGE is U_NEW_HRLY_WAGE_TMP)
     ),
     calculate_distance(U_NEW_QUOTA_TMP,M_QUOTA,DPP),
     ( (DPP =< THR,U_NEW_QUOTA is M_QUOTA) ;
       (DPP > THR,U_NEW_QUOTA is U_NEW_QUOTA_TMP)
     )
}
==>
union: percentage_change(hourly_wage,management,U_HRLY_WAGE) ::= any &
union: percantage_change(hourly_wage,management,U_NEW_HRLY_WAGE)
                                                       ::= true &
union: percentage_change(quota,union,U_QUOTA) ::= any &
union: percentage_change(quota,union,U_NEW_QUOTA) ::= true .
```

The management first deals with job security issues. The phase changes from start to job security when the following rule is fired:

```
phase(management, start) ::= true
==>
phase(management, start) ::= false &
phase(management, job_security) ::= true.
```

The system applies a tactic to choose a response. Four tactics are available: soft, moderate, tough and stalemate. A soft tactic causes the negotiator to give in on the less important issue and propose a compromise on the more important one. The moderate tactic proposes a compromise on both issues, and the tough tactic suggests a compromise on the less important issue and no movement on the more important issue. The stalemate tactic dictates no move on either issue. The extent of compromise (if any) is determined by the compromise factor.

The four issues that belong to the job security phase have been paired up by tying robotization with company relocation, and layoffs with the workers' participation in the company matters. Using the ranking of the issues, the system determines the relative importance of issues within pairs. Let relocation be more important to the management than the robotization of the company, and let SOFT be the current tactic in the current strategy. Our system will choose a metarule that generates the management's new position with a larger concession on the robotization issue and a small concession on the relocation issue. The size of concessions also depends on the management's strategy.

A strategy is a sequence of tactics, for example, "soft–moderate–tough–stalemate–moderate" or "always soft". It is possible that real–life negotiators do not alter their tactics, but we assume that the negotiator moves from one tactic to another depending on the flow of negotiations, in particular on the distance of the current position from the initial one. According to the literature (Rubin et al. 1977, p. 48) a negotiator is answerable to the constituency in that it must be proven that she did not give in too much. We assume that negotiators alter their tactics after analyzing the concessions made since the beginning of negotiations, and that each side is trying not to give in more than the opponent.

We model this by allowing the management to preset changes of tactics before the negotiation begins. For example, assume that the system representing the management starts with a soft tactic, and it finds after a few steps that its total concessions are much higher than the opponent's. The system will then upgrade its tactic to a tougher one, say, moderate. If this is not effective enough, there will be a move to the tough tactic, then to stalemate. The stalemate tactic

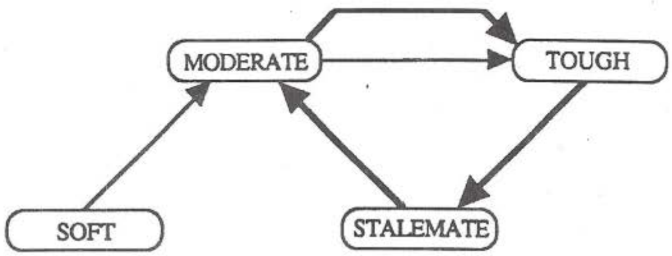

Figure 2. An example of a strategy

should help diminish the gap between the relevant positions, unless the other side is stalemating at the same time. The system stays with this tactic until the gap disappears, that is, the total concessions made by both sides become equal. Afterwards there will be a move back to the moderate tactic, the tough tactic, and so forth if necessary. We now present, as an example, a response rule that changes the management's tactic from soft to moderate:

```
tactic(management, soft) ::= true &
factor(management, FAC) ::= true &
margin(management, MARG) ::= true &
threshold(management, THR) ::= true &
{ compute_toughness(management, MTOUGH),
  compute_toughness(union, UTOUGH),
  NEW_UTOUGH is UTOUGH * MARG,
  MTOUGH > NEW_UTOUGH,
  NEW_FAC is FAC + 1,
  NEW_THR is THR / FAC
}
==>
factor(management, FAC) ::  false &
threshold(management, THR) ::  false &
factor(management, NEW_FAC) ::  true &
threshold(management, NEW_THR) ::= true &
tactic(management, soft) ::  false &
tactic(management, moderate) ::= true.
```

In Figure 2, we show the tactic changes that constitute one possible strategy. The thin lines mark the first two steps. The heavy lines mark three further steps that can be repeated.

The compromise factor, the margin and the threshold are initialized at the beginning of negotiation. The margin indicates the management's willingness of concessions. For example, the margin of 1.1 means that the management is willing to concede up to 10% more than the union, before it will switch to a more assertive tactic. The procedure `compute_toughness` calculates the total concession made by a side since the beginning of the negotiations. This is defined as a weighted distance between the present and initial position: a sum of weighted arithmetic distances between the initial demand on an issue and the present demand. The weighting is done by dividing the arithmetic distance by an issue's relative importance, which tells us how high on the management's agenda is this issue positioned. If an issue is present in only one side's position, and not mentioned by the other side, the arithmetic distance is assumed to be 1. In order to determine the value of the total concession, each side has to remember the opponent's initial position. A side uses its own weights when evaluating the opponent's total concession. This is consistent with actual negotiations, where an issue can be insignificant to one side but very important to the other side. For example, while the union does not really care about the issue of material waste, this is quite important to the management because it affects the company's total costs.

Figure 3 illustrates the formulæ used to compute distances and compare positions. It also shows how *toughness* and *distance* influence the change of tactics. These parameters express changes in the distance and represent the total weighted concession of each side.

4. Experiments and Results

4.1. Assumptions

We have designed experiments to determine how the choice of tactics can affect the outcome of the negotiations. Each tactic is defined as a formula to calculate concessions of a negotiating side. As we said earlier, issues are negotiated in pairs. Each side organizes its set of issues according to their importance. This order is used in the tactic formula.

Let A_{OUR}, A'_{OUR} be the old and new value of an issue, A_{THEIR} – the other side's value of the same issue. Let B_{OUR}, B'_{OUR}, B_{THEIR} be the respective values for the second, less important, issue in the pair. Let *COMP* be the compromise factor. Here are the exact formulæ for all tactics, earlier illustrated in Figure 1.

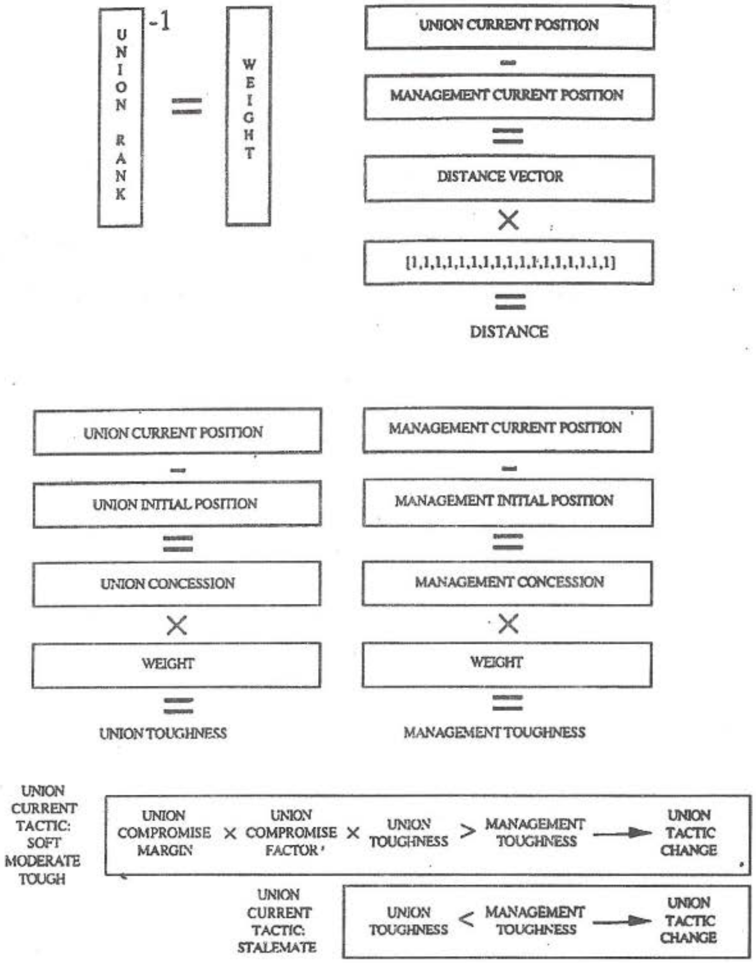

Figure 3. Details of the computation of the TOUGHNESS parameter

soft:

$$A'_{\text{OUR}} = A_{\text{OUR}} + (A_{\text{THEIR}} - A_{\text{OUR}})//COMP, \quad B'_{\text{OUR}} = B_{\text{THEIR}}$$

moderate:

$$A'_{\text{OUR}} = A_{\text{OUR}} + (A_{\text{THEIR}} - A_{\text{OUR}})//COMP,$$

$$B'_{\text{OUR}} = B_{\text{OUR}} + (B_{\text{THEIR}} - B_{\text{OUR}})//COMP$$

tough:

$$A'_{\text{OUR}} = A_{\text{OUR}}, \quad B'_{\text{OUR}} = B_{\text{OUR}} + (B_{\text{THEIR}} - B_{\text{OUR}})//COMP$$

stalemate:

$$A'_{\mathrm{OUR}} = A_{\mathrm{OUR}}, \quad B'_{\mathrm{OUR}} = B_{\mathrm{OUR}},$$

4.2. Experiments

We have conducted a series of experiments to find out how the choice of a particular strategy affects the outcome of the negotiations. We used two NEGOPLAN systems running as Prolog processes. An agent communicated the position by writing it into a file that was next read by the opponent. Both systems were running in batch mode. This mode was chosen because of a large number of experiments and because some of the negotiations would last thirty or more cycles.

The experiments were conducted using Quintus Prolog 3.1 and Xwindows on a Sparc 1+ workstation. We followed the progress of negotiation by displaying in a graphical form the change in the distance between the positions, the change in concessions made by each participant, and the strategy changes. We also gathered data on the number of cycles each negotiation lasted and on the payoffs for both participants at the end of negotiation.

In the first set of experiments we had made the following assumptions:

1) The tactic remains constant throughout the negotiation process.

2) The sides negotiate mostly the same issues (some may be 'local' to one side but a majority of issues is common).

3) Each side has the same set of compromise parameters:

   - the factor of 2,
   - the margin of 1.1,
   - the threshold of 0.01.

Each side could use one of four tactics. This gives sixteen possible experiments. The results are presented in Table 2. The notation NN + means that there was a deadlock after NN steps.

In the second set of experiments we have made the following assumptions:

1) The tactic changes during the negotiation process.

2) The sides negotiate mostly the same issues, and they have the same strategy.

3) Each side has the same set of compromise parameters:

- the factor of 2,
- the margin of 1.1,
- the threshold of 0.01.

Each side could initially use one of four tactics during the negotiation, for 16 possible experiments. An agent changed its tactic depending on the extent of concession made by the opponent, compared with own concession. We assumed that the strategy will have the following pattern:

soft → moderate → tough → stalemate → moderate →

→ tough → stalemate → moderate → ...

There are denumerably many finite sequences over a four–element alphabet, that is, denumerably many strategies. This particular strategy pattern seems very intuitive. Other intuitive paths ought to be equally good in showing that our approach can be used to simulate real–life negotiation.

The tactic changes were described by the set of metarules. In every cycle an agent would compare the concession made since the beginning of negotiations with the total concessions made by the opposite side. The total concession was defined as the weighted sum of concessions made on each issue. A concession made on an issue was the distance between the present and initial position on the issue. The weight of an issue was defined as R^{-1} where R (for Ranking) was the position of the issue on the priority list. If an agent made a substantially bigger concession than the opponent, the tactic would become more rigid.

The substantiality was decided by means of the compromise margin. In the change from soft to moderate we would consider that a side made substantially bigger concession if it was greater then the opponent's concession multiplied by the margin. In going from moderate to tough, and from tough to stalemate the concession had to be bigger then the opponent's concession multiplied by the margin and the compromise factor. An agent would abandon the stalemate strategy only if the opponent's concession became larger.

When the tactic changes, the factor and the threshold change as well. The factor is incremented by 1, and the threshold divided by the previous factor. This would cause slower progress of the negotiation towards the achievement of a compromise.

EXP.#	SIDE	init. tactic	cycles	payoff in $1,000,000	total payoff in $1,000,000
1	union	soft	7	-1.08	-2.88
	management	soft		-1.80	
2	union	moderate	12	-0.56	-0.44
	management	moderate		0.12	
3	union	tough	33 +	-4.05	-4.29
	management	tough		-0.24	
4	union	stalemate	33 +	-4.39	-4.65
	management	stalemate		-0.26	
5	union	soft	12	-1.98	-3.14
	management	moderate		-1.16	
6	union	soft	18	1.21	0.76
	management	tough		-0.45	
7	union	soft	26	-4.39	-3.67
	management	stalemate		0.72	
8	union	moderate	10	1.05	1.29
	management	soft		0.24	
9	union	moderate	19	2.09	2.42
	management	tough		0.33	
10	union	moderate	25	-4.39	-3.67
	management	stalemate		0.72	
11	union	tough	16	3.71	1.43
	management	soft		-2.28	
12	union	tough	25	2.76	-0.29
	management	moderate		-3.05	
13	union	tough	30 +	-3.89	-4.41
	management	stalemate		-0.26	
14	union	stalemate	20	5.88	5.62
	management	soft		-0.26	
15	union	stalemate	26	5.88	5.62
	management	moderate		-0.26	
16	union	stalemate	30 +	-3.89	-4.15
	management	tough		-0.26	

Table 2. Results of an experiment with a constant tactic

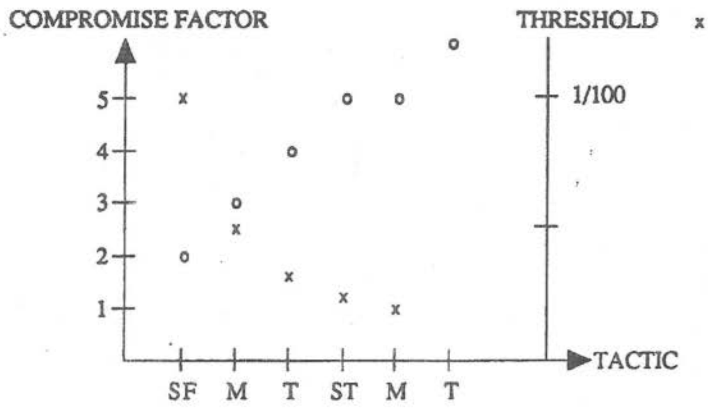

Figure 4. Changes in compromise factors and thresholds

Changes in compromise factors and thresholds are illustrated in Figure 4.

The results of the experiment with changing tactics are presented in Table 3.

The current implementation of NEGOPLAN includes a display facility that shows graphically the progress of negotiation in terms of changing distances and tactics. An example screen image is reproduced in Koperczak et al. (1991).

4.3. Analysis of Results

We now summarize the findings from 32 experimental runs of our system. The discussion is organized according to the parameters of negotiation that were varied between experiments.

The payoff determines what is a negotiating side's total gain (or loss) that results from all the concessions made by both sides during negotiation. In the competitive situation, we assume that each side wants to maximize its final payoff. In the cooperative situation, we assume that the sides together want to maximize the total payoff. In most situations, the payoff does not influence the progress of negotiations in our model, but is useful for the interpretation of the results.

Some of the negotiation runs ended with a deadlock. We distinguish two kinds of deadlock situations. If both sides stalemate, obviously a deadlock results. Deadlock may also arise from the intransigeance of both sides on an issue that they both rank highly. Such intransigeance will disable the tactic–changing mechanism, because the distance between the positions does not change.

An interesting–and realistic–situation is one in which the strategy of the

EXP.#	SIDE	init. tactic	cycles	payoff in $1,000,000	total payoff in $1,000,000
1	union	soft	19	-2.95	-2.95
	management	soft		0.00	
2	union	moderate	13	-0.56	-0.44
	management	moderate		0.12	
3	union	tough	33 +	-4.05	-4.29
	management	tough		-0.24	
4	union	stalemate	14	-0.56	-0.44
	management	stalemate		0.12	
5	union	soft	13	-1.98	-3.14
	management	moderate		-1.16	
6	union	soft	17	1.21	0.76
	management	tough		-0.45	
7	union	soft	14	1.28	0.67
	management	stalemate		-0.61	
8	union	moderate	19	-2.46	-2.01
	management	soft		0.45	
9	union	moderate	20	2.09	2.42
	management	tough		0.33	
10	union	moderate	14	2.24	2.38
	management	stalemate		0.14	
11	union	tough	20 +	-3.51	-3.77
	management	soft		-0.26	
12	union	tough	20 +	-3.51	-3.77
	management	moderate		-0.26	
13	union	tough	40 +	-3.80	-4.02
	management	stalemate		-0.22	
14	union	stalemate	20	-2.93	-2.42
	management	soft		0.51	
15	union	stalemate	20	-2.52	-2.07
	management	moderate		0.45	
16	union	stalemate	20	-1.71	-1.25
	management	tough		0.46	

Table 3. Results of an experiment with changing tactics

other side is unknown (as opposed to a strategy that conforms to a known pattern). If that is the case, a side cannot anticipate the opponent's moves, while it still wants to choose a tactic that would yield the best final payoff. It is, nevertheless, possible to determine a good strategy by analyzing all possible tactics and selecting one that offers the best outcome.

Side	Crite rion	Competitive situation	Strategy class	Initial tactic to use	Initial tactic to avoid
union	time	competition	constant		Tough//Stalemt
mgmt	time	competition	constant		Tough//Stalemt
union	money	competition	constant	Tough//Stalemt	Soft
mgmt	money	competition	constant	Soft//Stalemt	Moder
union	time	competition	changing		Tough
union	money	competition	changing		Tough//Stalemt
union	time	competition	unknown		Tough
mgmt	time	competition	unknown		Stalemt
union	money	competition	unknown	Moder	
mgmt	money	competition	unknown	Tough//Soft	Stalemt
union	money	cooperation	constant	Moder	Soft
mgmt	money	cooperation	constant	Soft	Stalemt
union	money	cooperation	changing	Moder	Tough
mgmt	money	cooperation	changing	Tough//Stalemt	Soft//Moder
union	time	cooperation	unknown		Tough
mgmt	time	cooperation	unknown		Stalemt
union	money	cooperation	unknown	Moder	
mgmt	money	cooperation	unknown	Tough//Soft	

Table 4. Recommendations for the negotiators

Some of the preliminary findings of this analysis are shown in Table 4. Its format is intended as an "advisory table" that would reflect the average *expected* outcome.

In Figure 5 we show that, in one experiment, MODERATE is overall the best strategy for the union, regardless of the initial tactics of the management. This experiment confirms the fact that our method is capable of finding good advice for the negotiator, for example, "be moderate, and you will win most of the time; do not be tough, or you will lose most of the time".

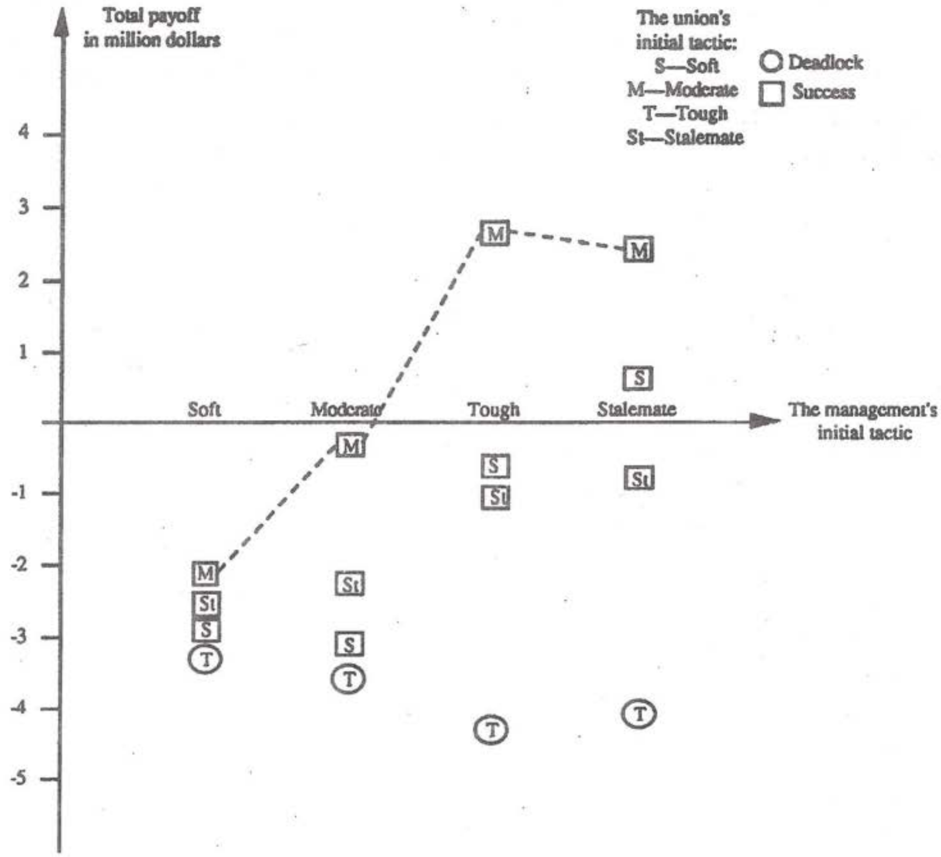

Figure 5. Details of a recommendation

5. Conclusion

We have presented a system that provides a testbed for experiments with negotiation strategies. The user of our system is completely free to define her own negotiation tactics and strategies. Section 4 sums up the findings from the experiments that we have performed. One of the interesting conclusions is that in union-management negotiations the moderate tactic can be universally recommended for the union. At the same time, this tactic is definitely not recommended for the management (this is not as pessimistic a conclusion as it may seem at a first glance, because "not moderate" includes the soft tactic as well).

A distinctive feature of this work is that, unlike in the approaches derived from game theory, the choice of tactics is based on the changes of the distance between the positions, rather than on the changes in the payoff. This allows us to view possible outcomes of a two–party negotiation more generally than in the game theory framework. Instead of just win//lose, we can describe the outcome as win//win, win//lose, and lose//lose.

One could question the pivotal role of ranking of issues of the negotiating sides in our approach. In fact, we do not only assume that such ranking exists, but we require that rankings of both sides be not identical. In our opinion, ranking is an important element of negotiation, and as such it is reflected in the pre-negotiation positioning of the two sides (Bartas 1974).

References

AXELROD R. (1987) The Evolution of Strategies in the Iterated Prisoner's Dilemma, L. Davis (ed.) *Genetic Algorithms and Simulated Annealing*, London, Pitman, 32–41 .

BARTOS O. (1974) *Process and Outcome of Negotiations*. New York, Columbia University Press.

ENGLAND J. L. (1979) Two Bargaining Automata. *J of Conflict Resolution*, **23**, 2, 296–325.

HAMNER W.C., YUKL G.A. (1977) The effectiveness of different offer strategies in bargaining, D. Druckman (ed.) *Negotiations, Social–Psychological Perspectives*. Sage, 137–168.

IKLE F.C. (1964) *How nations negotiate*. NY, Harper and Row.

KERSTEN G.E., MICHALOWSKI W., MATWIN S., SZPAKOWICZ S. (1988) Representing the Negotiation Process with a Rule–based Formalism, *Theory and Decision*, **25**, 225–257.

KERSTEN G.E, MICHALOWSKI W., SZPAKOWICZ S., KOPERCZAK Z. (1991) Restructurable Representations of Negotiation. *Management Science*, **37**, 10, 1269–1290.

KERSTEN G.E., SZPAKOWICZ S. (1990), Rule–based Formalism and Preference Representation: an Extension of NEGOPLAN". *European J of Operational Research*, **45**, 309–323.

KERSTEN G.E., SZPAKOWICZ S., KOPERCZAK Z. (1990), Modelling of Decision Making for Discrete Processes in Dynamic Environments. *Computers and Mathematics with Applications*, **20**, 9, 29–43.

KOPERCZAK Z., MATWIN S., SZPAKOWICZ S. (1991), *Modelling Negotiation Strategies with Two Interactive Expert Systems*. University of Ottawa, Department of Computer Science, TR–91–33.

LENG R.J., WALKER S.G. (1982) Comparing Two Studies of Crisis Bargaining. *J of Conflict Resolution*, **26**, 4, 571–591 .

LENG R.J., WHEELER H.G. (1979) Influence Strategies, Success and War. *J of Conflict Resolution*, **23**, 4, 655–683.

MATWIN S., SZPAKOWICZ S., KOPERCZAK Z., KERSTEN G.E., MICHALOWSKI W. (1989), NEGOPLAN: An Expert System Shell for Negotiation Support. *IEEE Expert*, **4**, 4, 50–62.

MICHALOWSKI W., KERSTEN G.E., KOPERCZAK Z., SZPAKOWICZ S. (1991), Disaster Management with NEGOPLAN. *Expert Systems with Applications*, **2**, 7–120.

OSCAMP S. (1971) Effects of Program Strategies on Cooperation in the Prisoner's Dilemma and Other Mixed–motive Games. *J of Conflict Resolution*, **15**, 225–229.

PATCHEN M. (1987) Strategies for Eliciting Cooperation from an Adversary. *J of Conflict Resolution*, **31**, 1, 164–185.

RAMBERG B. (1977) Tactical Advantages of Opening Positioning Strategies. *J of Conflict Resolution*, **21**, 4, 685–701.

RUBIN R., JEFFREY Z., BROWN B.R. (1977) *The Social Psychology of Bargaining and Negotiation*. Academic Press, NY.

SIMON H.A. (1976) *Administrative Behavior. A Study of Decision–making Process in Administrative Organizations*. New York: Free Press.

TOUZARD H. (1977) *La médiation et la résolution des conflicts*. Presses Universitaires de France, Paris.