

Planning in bridge with thematic actions

DEEPAK KHEMANI

Department of Computer Science & Engineering, Indian Institute of Technology, Madras 600 036, India

E-mail: khemani@iitm.ernet.in

Abstract. The task of planning in a dynamic and an uncertain domain is considerably more challenging than in domains traditionally adopted by AI planning methods. Planning in real situations has to be a knowledge intensive process, particularly since it is not easy to predict all the effects of one's actions. Contract bridge offers a domain in which many of the issues involved in real world problems can be addressed without having to make simplifications in representation. Planning in the game of bridge takes us away from the traditional search-based methods (like the alpha-beta procedure), which are applicable in complete-information games like chess. In this paper we look at how knowledge can be structured to plan for declarer play in bridge. This involves deploying known move combinations, triggered by patterns which are abstracted out of the input, and then assembling the structures into a workable plan. The results demonstrate the viability of the proposed concepts.

Keywords. Knowledge-based planning; games; knowledge structures; uncertain environment.

1. Introduction

Traditionally, games invoke images of game-theoretic methods (Rapoport 1966). These methods essentially search in the game space to move towards the saddle point, which is the best one can achieve against a perfect opponent in a complete information situation. Our emphasis, on the other hand, is entirely on the use of knowledge for problem solving. Planning (Hendler 1990) has been one of the key areas of AI research. However, traditional AI planning methods make some very simplifying assumptions about the world (Marks & Hammond 1988). One such assumption is that the world is completely known and stable. The effect of any proposed actions can be predicted precisely. Thus one assumes that a totally correct plan can be generated, and can be executed perfectly. In practice, the world is neither completely known nor static. As a result planning is imperfect and plans may warrant changes *en route*.

Contract bridge provides an excellent domain for the exploration of knowledge-based methods in an uncertain environment (Khemani & Ramakrishna 1989). We

get an environment with all the advantages of working with games, viz. a discrete finite universe which does not require representational compromises, where symbolic interaction is complete, and where performance can be accurately measured. In addition, being a four-player incomplete-information game, bridge provides opportunities which two-person-zero-sum complete-information games like chess do not. These include the necessity to reason under uncertain conditions, the need for formal communication encoded in moves between players, and a gamut of problems to tackle (as opposed to the same starting board position in Chess and Go). Thus, looking for the right moves in bridge can be classified as a planning problem in a *hostile real world* (Noronha & Sarma 1991). At the same time bridge provides a large enough problem-solving environment which has held humankind spellbound for decades. A more detailed description of the problems in the game is given in appendix A.

We have observed that a search-based strategy is not suitable. Instead we adopt a knowledge-based approach in which the planner works with card combinations for which known techniques can be applied. Such knowledge is essentially heuristic in nature, aggregating the efforts of a deeper analysis involving various cases. Representational economy dictates that heuristic knowledge be expressed at a suitably abstract level. However, this means that the *use* of such knowledge by an autonomous system also requires the ability to move between the domain and the abstract levels. Consequently, the knowledge structures proposed in this paper are composed of three diverse generic components: recognition, analysis, and instantiation.

The granularity of the knowledge structures is such that they apply only to a combination of a few cards, rather than the play of the entire hand. That is, the heuristic knowledge, which is in the form of ‘strong’ structures, suggests partial plans. The play of most hands can be seen as a combination of such known plays. Selecting the different relevant structures, and combining their actions into a coherent plan is the task of a scheduler, which has a flavour of a ‘weak’ method.

Planning then, as described in the sections that follow, is accomplished as a two-stage process – the activation of various knowledge structures suggesting possible partial plans, followed by the assembling of a feasible plan using some of them.

2. Heuristic knowledge

Heuristic knowledge can be said to be knowledge that facilitates solving a class of problems. A body constituting such knowledge would in some sense “know” what to do in a given situation. By its nature, it is often domain-specific in nature, though one can think of abstract enough “rules” which would apply across many domains. Heuristic knowledge is essentially associative in nature, matching patterns at a suitably abstract level and suggesting a known solution method. For example, a person learning to brew tea may work with the following rule “when the water begins to boil, add the tea leaves and switch off the stove”. To solve a problem, heuristic rules like these can be blindly employed without doing any deeper cause-effect analysis. This, however, would be only the first step towards expert behaviour. To be able to explain one’s actions, and to be able to adapt to somewhat different situations, or even to break the rule, would require a deeper analysis. The next section illustrates this idea in more detail using the domain of bridge.

Assembling a heuristic knowledge base is, however, not a trivial task. We claim

that this is mainly because such knowledge cannot function in isolation*. Instead, this knowledge applies to patterns in information which have been abstracted out of raw data, away from irrelevant detail. The process of abstraction may be methodologically quite different, but would have to work in tandem with the higher level reasoning, which we may call meta-semantic reasoning (Khemani 1988). In fact, the process of abstraction will be strongly influenced by the sort of features the meta-semantic heuristics need. On the other hand, solutions provided by the heuristic knowledge would again have to be effected at the domain level thus clearly defining their semantics. In this process the detail that was ignored during abstraction has to be again taken into account as actions have, in the end, to be visible at the domain level. Thus, one needs to incorporate diverse generic skills (Chandrasekaran 1986) in a system which is to function autonomously.

In addition, there is a need to structure the knowledge in a manner that relevant parts are invoked automatically in any situation. It has quite convincingly been shown in the context of understanding (Dyer 1986) that various types of knowledge structures are needed. These may vary from frames (Charnaik 1978) and scripts (Schunk & Abelson 1977) to handle stereotypical activity, to more complex structures (Lehnert 1987) to capture intentions, goals, plans and their myriad interrelations.

We represent planning knowledge for declarer play in bridge in the form of thematic acts (TA) (Khemani 1988). Each TA encapsulates a technique to tackle a recognizable situation, and the overall strategy is invoking one or more TA and weaving their suggested actions together into a single plan of action.

3. Knowledge for planning

Planning knowledge is required to charter a path through the combinatorially exploding number of possible sequences of low-level actions. Tactical knowledge to tackle each well-known situation can be represented at a higher level, somewhat like the macro operators introduced by Korf (1985). However, what is also needed is the ability to view the macro operators at a higher level, recognizing the situations in which they could be applicable and the changes that they may effect. We introduce a knowledge structure called thematic act which operates in an abstraction space. A TA has knowledge of patterns to look for in the raw data, knowledge of move combinations (a partial plan) which might be useful, and knowledge of the overall effect of the moves (which we call *meta-semantic* knowledge). The TA captures the theme behind all the similar move combinations which are required to achieve goals in similar situations. We explore the nature of such thematic knowledge in the context of contract bridge.

3.1 Functional abstraction

In any interesting enough domain the number of different cases would be too many to put in memory individually. Instead the goal should be to try and extract general

*Isolated heuristic systems, for example MYCIN (Shortliffe 1976), have been built, of course, but they are only prescriptive in nature. They capture the associations between symptoms and disease (or the therapy), but the tasks of observing the symptoms and effecting the therapy are left to the human user.

principles which are applicable in many cases. Consider, for example, the concept of a 'finesse' in bridge. It can be illustrated by the following case, where all the cards are from the same suit, and "?" represents an unknown card⁺.

North:AQ		
West:??	East:??	(1)
South:3 2		

South plays the 3. Let us say West plays the 8. If North plays the Q, then the play is said to be a finesse. If the K is in the West hand the North-South side makes an extra trick with the Q. If on the other hand the K is in the East hand, the finesse fails and East can win the trick.

The important feature in the above play is the combination of Ace & Queen which forms a 'tenace' over the West hand. Straightaway we can see that the South cards are insignificant, and could have been any of the small cards. Also we can see that the North hand could have had AKJ, and one could have still finessed the Jack. What is more, if the AKQ and J had already been played, then a combination of 10 & 8 could have formed a tenace over East's 9. So *functionally* a tenace can be defined as follows

North:RC1 RC3		
South:x		(2)

which essentially says that if there is a rank-class-1 (RC1) card along with a rank-class-3 (RC3) card, and there is a small card (x) in the other hand, then one can *plan* to take a finesse against the opponent's RC2 card. The rank-class is different from the rank, in the sense that it clubs together cards of touching ranks held by a side. For example, if both the Ace and the King belong to the same side then both are of rank-class 1, whereas the rank of K is 2. In other words, since both the Ace and the King are held by the same side, the King could also be considered to be a winner.

Thus *the tenace is a functional abstraction* which characterizes the situation when a *finesse* can be played. Each card is viewed functionally by its role, which may vary in different situations, rather than by its identifier, which is constant. In addition one is looking for *combinations* which can be exploited. This is quite different from the accepted meaning of abstraction, for example in ABSTRIPS (Sacerdoti 1974), where the emphasis is only on ignoring unnecessary detail.

3.2 *The level of knowledge*

The knowledge captured in a TA is aimed at its use. It does not provide for a situation-action-effect analysis that could provide an explanation for *why* the recommended actions are useful. Consequently it cannot also cater to extra-thematic situations. Consider, for example, the following TA. The hands are

North:K 9 x		
South:A J xxx		(3)

We need four tricks from this combination, and can afford to lose one (but not two).

⁺ Symbols pertaining to the game and used in the paper are defined in the appendix.

The standard play, which the TA recommends in this situation, is to play the A first, and then a small card towards the K-9, intending to play the 9. However, if West does not follow suit, win with the King, and play the 9 towards the Jack. It may be observed that this play is somewhat counter intuitive as it tends to disrupt the A-J tenace that exists.

The knowledge contained in the TA cannot provide explanations as to why this play is best. That requires a deeper analysis which deals with cases, and is along the following lines. There is a danger when one of the opponents holds Q-10-x-x. In such a situation the above play insures against losing two tricks. If West holds Q-10-x-x he is forced to play either the Q or the 10 lest North wins with the 9. And then the 9 and the Jack between them will restrict West to winning only one trick with the other card. If East holds the Q-10-x-x then West will show out on the second round. North wins, and now playing towards the Jack ensures that East can only win with the Queen. The intuitive play, on the other hand, with this card combination, is to first win the King, and then finesse the Jack. This play is best for winning five tricks, but when West has Q-10-x-x it loses two tricks. The recommended play for four tricks also guards against this eventuality. However, the ability to do the above analysis has no effect on the success of the TA in that situation.

There can be extra-thematic situations where the standard advice has to be ignored. One such situation arises when one is willing to alter one's risk-gain equation. Thus, it may happen, in the above example, that in a larger context (e.g. a tournament) one is willing to risk losing two tricks for the small gain of an extra trick (when only four are required). However, such reasoning is beyond the scope of this work.

4. Planning with the TA

The goal in bridge is to arrive at a plan to make a certain number of tricks. The cards dealt define the starting position. The opponents' cards are not known. However some information may be available from the bidding, and more accrues as play proceeds. The next section looks at the role of this information in decision making. In this section we examine the role of TA in forming the plan.

The basic strategy used in the planning process is Means-Ends-Analysis (MEA) (Newell & Simon 1963). This applies easily to the bridge problem as the goal of making a certain number of tricks can be split into conjuncts of developing tricks in different suits. There are, however, overall constraints imposed due to sharing of common resources. For example, if we need to develop three tricks, and there exist two suits that can provide two each, it may still not be possible to develop both due to 'tempo' restrictions. So, the TA merely suggest partial plans to reduce a part or the whole of the difference, and it is left to an overall scheduler to assemble the best combination.

The TA provides the knowledge to tackle a given card combination in a given suit. It is not necessary, however, that a suit can be optimally played in isolation. For example, in the following card combination

North: A Q J	
South: 6 4 3	(4)

the recommended play is to finesse twice. This requires that the play be started twice from the South position. Each time the lead shifts to the North hand, since the trick

is won there. Thus, some other suit needs to provide an 'entry' into the South hand to finesse again. In this way the plays for different suits need to be interleaved and woven together. Notice also that the play for the above combination can be hierarchically composed using the finesse.

Each suit has some potential for developing a certain number of tricks with some probability. For example combination (3) can be played for 4 tricks with greater insurance than for 5 tricks, and it can be played for 3 tricks with assured results. However, there are different conditions required by each of the plays. The 3-trick play requires available *leeway* and *tempo* to be at least 2. These terms are defined below. The 4-trick play requires them to be at least 1, while the 5-trick play has no such constraints.

The selection of different TA combinations is dictated by some overall constraints. Constraints are of two types. One is the maximum number of tricks that the planner can afford to lose. For example, if you have contracted for 13 tricks, and need only 4 from the combination (3), you still cannot select the 4-trick play because it allows for losing one trick to the opponents. That is, the *leeway* required is one, but that available is zero. Therefore, even though the play suggests itself, the scheduler cannot select it.

A related constraint is that of *tempo*. This is determined by how close the opponents are to setting up tricks for themselves. It is measured in terms of the *number of times* the planner can lose tricks to the opponents before they are in a position to defeat the contract. It is the job of the scheduler to pick up a combination of thematic acts within the overall leeway and tempo constraints.

The *task* of planning is to combine thematic acts for different suit combinations to achieve the *overall* goal of tricks. In doing so one has to be careful that the different partial plans can be combined together to form a feasible line of action. For example, if combination (4) is replicated in all suits,

North:A Q J, A Q J, A Q J, A Q J 10
South:6 4 3, 6 4 3, 6 4 3, 6 4 3 2 (5)

then not more than one finesse could be possible, though the TA would create partial plans for 9 of them! However, if some of the suits were interchanged,

North:A Q J, 6 4 3, A Q J, 6 4 3 2
South:6 4 3, A Q J, 6 4 3, A Q J 10 (6)

then all 9 finesses could be effected, possibly winning all 13 tricks.

Hence by doing a meta-semantic analysis of the plays offered by different suits a suitable plan can be arrived at. The thematic acts themselves are tuned to the best suit play, i.e., one that has highest probability of success. The task now is to combine some of them into forming a plan. In doing so one may discover that a TA cannot be directly applied, and one may have to look for an alternate play which can fit into the plan. The plan generated will have a certain probability of success that is determined by the probability of the assumptions made *en route* being true. By and large the best plan selects itself, since the acts take care to suggest the best individual play. But if the best individual thematic acts cannot be combined to form a plan one may have to look at various options. This is discussed in §5. Also, sometimes the mounting evidence may require meta-thematic analysis for selecting a plan.

It can be seen that there are two distinct aspects to the task of planning. The first,

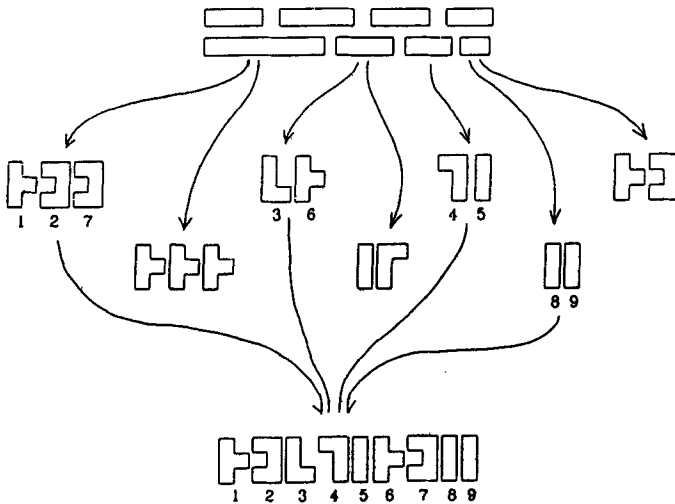


Figure 1. Two-stage planning – A schematic flowchart.

ANNOTATION

The larger problem contains a set of smaller problems for which (partial) solutions are known.

A thematic planner looks at subproblems and the thematic acts suggest partial plans. Numbers below show their position in the final plan.

A general purpose scheduler looks at various thematic actions and assembles together a coherent plan of actions which mesh together.

which involves applying known solutions to recognizable parts of the problem, is more knowledge intensive and has a 'strong' flavour. The TA provide this capacity, and each partial plan is retrieved complete with its own ordering (Berlin 1985). On the other hand, the scheduler which assembles plans is essentially the embodiment of a weak method. Figure 1 schematically illustrates the formation of partial plans and the assembling of a total plan with some of them. Such an architecture is imperative if planning is to be done in any complex domain. An autonomous agent in such a domain will have to cater to many goals which crop up as it functions. Planning, therefore, can no longer be viewed as problem-solving for a well defined goal, but is an ongoing process. A basic requirement then, in such a domain, is the ability to incorporate partial solutions in one's activity. This will also enable opportunistic planning (Hayes-Roth & Hayes-Roth 1979) specifically while executing plans (Birnbaum 1985).

5. Choosing between plans

A given plan succeeds when the implied assumptions it makes are satisfied. One of the strategies in planning is to make as few assumptions as possible. The assumptions made select a subset $S1$ of the all possible situations U when the given plan will succeed (figure 2).




Figure 2. The subset $S1$ of the situation when the plan works in the universe U .




Figure 3. The subset $S1' = S1 \cap U'$ in the constrained universe U' .




Figure 4. Sets $S1'$ and $S2'$ for plans $P1$ and $P2$ in the constrained universe U' .

The aim in planning with thematic acts is to maximize the size of $S1$. However, as play proceeds, one may get more information about the actual situation. This contracts the universe U to a set U' . The subset $S1' = S1 \cap U'$, now is the set of situations when our plan succeeds (figure 3).

Between two competing plans $P1$ and $P2$, the sets $S1'$ and $S2'$ now define the probability of success. The shaded area in figure 4 shows the differing cases for which $P1$ and $P2$ succeed, given U' . It can clearly be seen that the set U' can play a decisive role in selecting the better plan. For example, without any knowledge about world U , plan $P1$ is a better plan if $S1$ is larger than $S2$. However knowing that the universe is constrained to U' plan $P2$ turns out to be better if $S2'$ is larger than $S1'$.

The overall strategy in planning to play a bridge hand is two-fold. One, to try and synthesize plans which make as few assumptions as possible. This is equivalent to saying that the plan should work in as many situations as possible. The second aspect of planning is to try and glean more information about the world. The effort here is to try and reduce the number of situations that have to be catered for. One can see that such a strategy is also likely to prove useful in any incomplete information domain.

6. Implementation

We have implemented a program for declarer play in bridge. The input to the program is the declarer's hand (South), the dummy (North), and the opening lead made by West. Currently the program plays only no trump (NT) contracts. The overall task is broken up into the following modules.

The program begins by abstracting the relevant patterns out of the cards. It counts its tricks, and determines the goal (i.e. the number of tricks that have to be generated), and the constraints. For example, when planning for a 7NT contract there is no leeway, but a 3NT contract allows four tricks for the opponents and the planner often has some tempo. The program then looks for any card combinations that it 'knows' to tackle effectively. This information is provided by the corresponding TA and the pattern also serves to index the TA. The planner constructs partial plans using the procedural knowledge contained in the TA. The current implementation generates one partial plan for each suit and passes it on to the scheduler. If this cannot fit into the overall plan, then the program backtracks and tries another. The ordering of the actions suggested by the TA is done in a heuristic manner by the scheduler. Thus the right play in the case referred to in Berlin (1985) is found directly because of the heuristics employed, rather than by a deeper analysis of interaction of subgoals. The

scheduler orders the actions to link up the different thematic acts. For example, combination (4) requires that the suit be played from the South hand, and it is the job of the scheduler to incorporate the play when the opportunity presents itself. During play, if any assumptions are violated then control is passed back to the planner. This also happens when an unexpected trick comes up. A more detailed description of the program is given in Ramachandran & Khemani (1992).

6.1 Results

The trace left by the program as it plans and plays a few hands is included in appendix B. The trace constitutes of canned messages printed by various sections when they execute. The input to the program are the North and South hands. For our convenience the East and West hands have also been included in the printout.

In hand 1, needing to make all 13 tricks the program spots the tenaces in Hearts. It observes that it can cater to the worst break by finessing 3 times, and forms a plan accordingly. In hand 2 the program relies on the spade break for tricks, and noting the shortage of entries, plans to 'duck' one round of the suit. Hand 3 is a sort of combination of 1 and 2. The feature of using a delayed finesse to solve the communication problem has been encoded as a separate TA. Hands 4 and 5 are the same for North and South (thus leading to the same plan), but different for East and West. In both cases the assumed spade break does not materialize, and the program looks at alternative sources. In hand 4 it avoids the spade finesse because it knows that it is bound to fail, while in 5 it chooses it knowing that it will succeed.

7. Conclusions

In this paper we have looked at a two-stage mechanism for planning, using contract bridge as the domain. In the first stage partial plans are suggested by a theme-based planner. The planner employs knowledge structures called thematic acts which have the flavour of strong AI. Each TA encapsulates the knowledge required for solving a problem in a familiar situation in a well-defined way. To retrieve and use this knowledge in applicable situations, the theme-based planner requires three kinds of processing. First, in functional abstraction the patterns which serve as indices to the TA are extracted. Then, meta-semantic reasoning computes the cause-effect relations of the TA in the situation. Finally, procedural knowledge orchestrates the actual domain level primitive actions. The partial plans thus are combinations of primitive actions, not necessarily complete in themselves, that can be used to solve a part of the problem.

In the second stage, a scheduler combines the actions suggested by the TA into a coherent plan. This may often require that actions of different TA be interleaved. The scheduler is basically an embodiment of a weak method, but one which operates under the constraints imposed by the first stage of thematic planning. In our implementation the strategy of means-ends-analysis has been adopted, as it is well-suited to the domain of planning in bridge.

The results of our implementation have been very encouraging, as can be witnessed in the examples shown. We feel that such a scheme of planning is particularly attractive when planning in a complex domain where (partial) plans for newer goals have to be incorporated into the larger plans of an agent who may be continuously interacting with the world.

Appendix A. Contract bridge: The problems

In this appendix we look at the game of contract bridge, highlighting some of the problems that are of interest. We observe that unlike chess, which is a two-person zero-sum complete-information game, bridge cannot be tackled by elegant mechanisms like the minimax search method. This is because bridge is *not* a complete information game. Then, since one does not know the cards held by the opponents, one cannot project the play into the future to try and discover which strategy is most profitable. Instead one has to rely on some knowledge-intensive method. Bridge can be classified as a two-side incomplete-information game. Further complication is introduced by the fact that each side constitutes two persons. Therefore communication is vital. Not only does one have to convey information, within the rules of the game of course, to the partner but one also needs to intercept opponents messages to learn their intentions. Almost as a corollary, at a more sophisticated level one may even want to send out misleading signals to lead opponents astray.

Contract bridge is played with a regular pack of 52 cards dealt randomly and equally among 4 players. Let us call them North, South, East and West, according to their position on the table. North and South are partners, as are East and West. The cards are ranked in the order Ace (A), King (K), Queen (Q), Jack (J), 10 (T), 9, ... 2 in each suit. Each player plays a card, in clockwise order, and the highest ranking card wins the trick. Thirteen such tricks are played, each time the winner of the precious trick starting play. This constitutes one *deal* or one *hand*.

There are two stages of play in each deal, viz. bidding, followed by the play of cards. The goal in a deal is to maximize *points*. The points essentially depend upon bidding. Bids are made for the number of tricks the side promises to make, given the stated 'trump' suit. Eventually the highest bid is accepted in each deal. This is known as the *contract*. Generally, the higher a side bids the more points it is likely to win, *provided* it can fulfill the contract. That is, if the side can make the number of tricks it has bid for. If it succeeds, it wins some points. Let us call them success-points. If it loses, then the opponents get some points instead, which we can call penalty-points.

A1. Bidding

The straightforward goal in bidding is to bid the *highest* number of tricks one thinks one *can* make. That is, to maximize success-points won. The means used in this process are the following,

- (1) Evaluation of own hand.
- (2) Communication with partner.
- (3) Projection of play.

Of these, the first two are simpler and can possibly be handled by heuristic methods. The third is more difficult, as it would involve constructing plausible distributions (based on the bids heard, and on probability) and then projecting the play.

A more complex goal is to make a *sacrifice bid*. This essentially means overbidding, over an opponent bid, with the hope that the penalty-points loss will be lesser than the opponents' expected success-points gain, thus being an overall gain.

Even more complex goals are to sabotage the opponents communication. This may mean consuming the bidding space (jamming the communications channel), or even making 'false' bids to confuse opponents. In the process, an enterprising planner

may make an 'advance sacrifice' to 'push' the opponents higher than they can manage, or to escape with a lighter penalty.

Considering that all these processes happen when the planner can see only one hand, one observes that bidding is probably the more difficult part of the game.

A2. *Play of the hand*

Once bidding is over, the goal for the play stage has been defined. One side has the contract, and is required to make the bid number of tricks. At this stage one player of the contracting side (called the dummy) exposes the cards to everybody, while the other (called the declarer) plans and executes the play. The opposing side (called defenders) are said to *defend* the contract, trying in fact to defeat the accomplishment of the contract by the declarer.

One can see that the situation at this stage is non-symmetric. The declarer knows the entire strength of his side, and is in total control of the play of the cards. He is also aware of the entire assets of the defence, in terms of material strength, since they have the remaining 26 cards. Each defender knows only his own hand, and cannot see his partner's hand. Therefore the two defenders have to combine their efforts to try and achieve the goal. This necessarily involves (formal) communication between the two. Both can see the dummy also.

Since the cards of all the players cannot be seen, one cannot project moves into the future. Methods like minimax search are therefore ruled out immediately. Instead, the success of a strategy can only be estimated based on the probabilistic distribution of cards, and any information gleaned from the communication taking place. The strategies themselves are derived from knowledge about the various known methods of tackling various card combinations. A constant endeavour is to cater to as many situations as possible, to maximize the chances of success.

The straightforward goal in the play of the hand is to make the number of tricks as stated in the contract. The emphasis is on maximizing the probability of success. If success is assured, then the goal can be revised to increasing the number of tricks won, as some more points can then be gained. If success seems unlikely, then a planner may even choose to minimize losses, i.e. the penalty-points won by the opponents. Like in bidding, the planner may attempt to do better than par, by exploiting the incomplete information that the opponents have. This may introduce complex 'meta-level' goals of protecting information, or sending out misleading signals.

Thus, we see that unlike in games like chess, where a clear-cut strategy of aiming for the minimax value (saddle) points is meaningful, in bridge one has to largely grapple with incomplete information. In the face of such uncertainty, planning in the game of bridge can only be a complex knowledge-intensive activity.

Appendix B. Traces of the program

This appendix contains the traces left by the program as it plans and plays some hands. The trace contains the complete hand, followed by the trace, followed by the record of play. Even in cases where the system replans after some play, the organization of text remains the same. Text in square brackets is comments by the author.

Example B1

[Hand 1][†]


Contract = 7 Notrump

North – Dummy

♠ 752
♥ AQJT5
♦ A43
♣ K3

West

♠
♥ K976
♦ T982
♣ J964



East

♠ JT98643
♥
♦ 765
♣ T87

South

♠ AKQ
♥ 8432
♦ KQ
♣ AQ52

The planning phase...

Have 10 top tricks... Need 3 tricks more.... Tempo available 10. If Hearts break 4–0 we get 3 extra tricks losing 1... can't afford that... first glance over... Tenace over West found in Hearts. Have planned for 3 tricks in Spades, 5 tricks in Hearts, 3 tricks in Diamonds and 3 tricks in Clubs. [Note that it was planned for 5 + 3 + 3 + 3 = 14 tricks. It does not need to decide which one of the winners it will have to discard. The discarding problem has been passed over to the scheduler which will finally throw away a winner.]

The play..

Examining lead... Looks like top of a sequence.

(A * indicates that the corresponding player had the lead.)

No.	West	North	East	South	Won by
1	*J ♦	3 ♦	5 ♦	K ♦	South
2	2 ♦	2 ♣	3 ♣	*A ♦	South
3	6 ♥	T ♥	4 ♣	*2 ♥	North
4	4 ♣	*5 ♦	6 ♣	K ♦	South
5	7 ♥	J ♥	6 ♦	*3 ♥	North
6	8 ♦	*7 ♦	8 ♣	Q ♦	South
7	9 ♥	Q ♥	9 ♣	*4 ♥	North

[†]The symbols ♠, ♥, ♦ and ♣ are used for the suits spades, hearts, diamonds & clubs, respectively.

8	K♥	*A♥	7♦	8♥	North
9	9♦	*5♥	7♣	5♣	North
10	T♦	*4♦	T♣	Q♦	South
11	6♣	K♣	8♣	*2♣	North
12	9♣	*A♦	T♣	Q♣	North
13	J♣	*3♣	J♣	A♣	South

Tricks made = 13

Tricks contracted = 13

Good one! Made the contract!!!

Example B2

[Hand 2]

Contract = 3 Notrump

North – Dummy

♦ A7643

♥ 765

♦ 5

♣ 8754

West

♦ T8

♥ QJT2

♦ JT9

♣ T8

Opening lead – Q♦

East

♦ QJ2

♥ 843

♦ 8763

♣ KJ6

South

♦ K95

♥ AK9

♦ AK42

♣ A32

The planning phase...

Have 7 top tricks. ... Need 2 tricks more. ... Tempo available 4. If Spades break 3–2 we get 2 extra tricks losing 1. If Clubs break 3–3 we get 1 extra trick losing 2. ... first glance over... playing low to low due to lack of entries. Plan workable... Have planned for 4 tricks in Spades, 2 tricks in Hearts, 2 tricks in Diamonds and 1 trick in Clubs.

The play..

(A * indicates that the corresponding player had the lead.)

No.	West	North	East	South	Won by
1	*Q♦	5♦	3♦	A♦	South
2	8♠	3♠	2♠	*5♠	West
3	*J♦	7♥	6♦	K♦	South
4	T♠	4♠	J♠	*K♠	South
5	2♥	A♠	Q♠	*9♠	North
6	9♦	*7♠	3♥	9♥	North
7	T♦	*6♠	4♥	2♦	North
8	T♥	*5♥	8♥	A♥	South
9	Q♥	6♥	6♠	*K♥	South
10	9♠	4♠	J♠	*A♠	South
11	T♠	5♠	8♦	*4♦	East
12	J♥	7♠	*7♦	2♠	East
13	Q♠	8♠	*K♠	3♠	East

Tricks made = 9

Tricks contracted = 9

Good one! Made the contract!!!

Example B3

[Hand 3]

Contract = 7 Notrump

North ÷ Dummy

♦ AQ6543

♥ J7

♦ 84

♣ T65

West

♦ KJ9

♥ Q98

♦ Q96

♣ Q83

Opening lead – K♦

East

♦ T7

♥ T632

♦ T52

♣ J972

South

♦ 82

♥ AK54

♦ AJ73

♣ AK4

The planning phase...

Have 6 top tricks... Need 3 tricks more... Tempo available 4. If Spades break 3-2 we get 3 extra tricks losing 2... first glance over... playing low to low due to lack of

entries... plan workable... Have planned for 5 tricks in Spades, 2 tricks in Hearts, 1 trick in Diamonds and 2 tricks in Clubs.

Examining lead... Looks like top of nothing lead.

The play...

(A * indicates that the corresponding player had the lead.)

No.	West	North	East	South	Won by
1	*K♦	4♦	2♦	3♦	West
2	*3♣	5♣	J♣	A♣	South
3	9♣	3♣	T♣	*8♣	East
4	8♣	6♣	*2♣	K♣	South
5	J♦	Q♦	7♦	*2♦	North
6	K♣	*A♣	2♥	4♥	North
7	8♥	*6♦	3♥	5♥	North
8	6♦	*5♣	6♥	7♦	North
9	9♥	*4♦	T♥	J♦	North
10	Q♥	*7♥	7♣	A♥	South
11	9♦	J♥	5♦	*K♥	South
12	Q♦	8♦	T♦	*A♦	South
13	Q♣	T♣	9♣	*4♣	West

Tricks made = 10

Tricks contracted = 9

Good one! Made the contract!!!

Example B4

[Hand 4]

Contract = 7 Notrump

North - Dummy

♦ AK32
♥ AQT98
♦ J32
♣ Q

West

♦ J975
♥ K3
♦ 7
♣ T7652

Opening lead - J♣

East

♦ 4
♥ J6542
♦ T965
♣ 943

South

♦ QT86
♥ 7
♦ AKQ84
♣ AK8

The planning phase...

Have 11 top tricks.... Need 2 tricks more.... Tempo available 0. If Spades break 3-2 we get 1 extra trick losing 0. If Hearts break 6-1 we get 2 extra tricks losing 2... can't afford that. If Diamonds break 4-1 we get 1 extra trick losing 0... first glance over....

[Note that it does not look for spades as it knows that the break is not favourable.]

Have planned for 4 tricks in Spades, 1 trick in Hearts, 5 tricks in Diamonds and 3 tricks in Clubs.

Examining lead... Looks like top of a sequence.

After trick 3

Plan aborted because of Expectation Failure.

New plan being generated...

Have 8 top tricks.... Need 2 tricks more.... Tempo available 0. If Hearts break 6-1 we get 2 extra tricks losing 2... can't afford that. If Diamonds break 4-1 we get 1 extra trick losing 0... first glance over.... Tenace over East found in Spades. Finesse not possible as opponent has shown out. Tenace over West found in Hearts. Have planned for 1 trick in Spades, 2 tricks in Hearts, 5 tricks in Diamonds and 2 tricks in Clubs.

The play...

(A * indicates that the corresponding player had the lead.)

No.	West	North	East	South	Won by
1	*J ♣	Q ♣	3 ♣	8 ♣	North
2	5 ♣	*A ♣	4 ♣	6 ♣	North
3	7 ♣	*K ♣	4 ♣	8 ♣	North
4	9 ♣	*2 ♣	9 ♣	Q ♣	South
5	2 ♣	9 ♥	2 ♥	*A ♣	South
6	5 ♣	8 ♥	5 ♣	*K ♣	South
7	3 ♥	Q ♥	4 ♥	*7 ♥	North
8	K ♥	*A ♥	5 ♥	T ♣	North
9	7 ♣	*J ♣	6 ♣	4 ♣	North
10	6 ♣	*2 ♣	9 ♣	A ♣	South
11	7 ♣	3 ♣	T ♣	*K ♣	South
12	T ♣	3 ♣	6 ♥	*Q ♣	South
13	J ♣	T ♥	J ♥	*8 ♣	South

Tricks made = 13

Tricks contracted = 13

Good one! Made the contract!!!

Example B5

[Hand 5]

Contract = 7 Notrump

North – Dummy

♠ A K 3 2
♥ A Q T 9 8
♦ J 3 2
♣ Q

West

♠ 4
♥ J 6 5 4 2
♦ 7
♣ T 7 6 5 2

Opening lead – J ♣

East

♠ J 9 5 3
♥ K 3
♦ T 9 6 5
♣ 9 4 3

South

♠ Q T 8 6
♥ 7
♦ A K Q 8 4
♣ A K 8

The planning phase...

Have 11 top tricks... Need 2 tricks more... Tempo available 0. If Spades break 3–2 we get 1 extra trick losing 0. If Hearts break 6–1 we get 2 extra tricks losing 2... can't afford that. If Diamonds break 4–1 we get 1 extra trick losing 0... first glance over... Have planned for 11 tricks. Have planned for 4 tricks in Spades, 1 trick in Hearts, 5 tricks in Diamonds and 3 tricks in Clubs.

The play...

Examining lead... Looks like top of a sequence.

After trick 3

Plan aborted because of Expectation Failure

New plan being generated...

Have 8 top tricks... Need 2 tricks more... Tempo available 0. If Hearts break 6–1 we get 2 extra tricks losing 2... can't afford that. If Diamonds break 4–1 we get 1 extra trick losing 0... first glance over... Have planned for 8 tricks. Tenace over East found in Spades. Have placed the necessary card with opponent. Sure finesse!! Tenace over West found in Hearts. Have planned for 2 tricks in Spades, 1 trick in Hearts, 5 tricks in Diamonds and 2 tricks in Clubs.

[Note that the finesse in hearts is abandoned in preference to the finesse in spades as 10 tricks are available without it.]

The play...

(A * indicates that the corresponding player had the lead.)

No.	West	North	East	South	Won by
1	*J ♣	Q ♣	3 ♣	8 ♣	North
2	4 ♣	*A ♣	5 ♣	6 ♣	North
3	2 ♣	*K ♣	7 ♣	8 ♣	North
4	2 ♥	*A ♥	3 ♥	7 ♥	North
5	5 ♣	*2 ♣	9 ♣	T ♣	South
6	6 ♣	8 ♥	4 ♣	*A ♣	South
7	7 ♣	Q ♥	9 ♣	*K ♣	South
8	7 ♦	3 ♣	J ♣	*Q ♣	South
9	T ♣	J ♦	5 ♦	*4 ♦	North
10	4 ♥	*2 ♦	6 ♦	A ♦	South
11	5 ♥	3 ♦	9 ♦	*K ♦	South
12	6 ♥	T ♥	T ♦	*Q ♦	South
13	J ♥	9 ♥	K ♥	*8 ♦	South

Tricks made = 13

Tricks contracted = 13. Good one! Made the contract!!!

References

- Berlin D L S 1985 SPAN: Integrated problem solving tactics. *Proc. IJCAI-85*p 1047
- Birnbaum L 1985 A short note on opportunistic planning and memory in arguments. *Proc. IJCAI-85* p. 281
- Chandrasekaran B 1986 Generic tasks in knowledge-based reasoning: High-level building blocks for expert system design. *IEEE Expert* 1: 23-30
- Charniak E 1978 On the use of framed knowledge in language comprehension. *Artif. Intell.* 11: 225-265
- Dyer M G 1986 *In-depth understanding* (Cambridge: MIT Press)
- Hayes-Roth B, Hayes-Roth F A 1979 A cognitive model of planning, *Cognitive Sci.* 3: 275-310
- Hendler J, Tate A, Drummond M 1990 AI planning: Systems and techniques. *AI Mag.* Summer: 61-77
- Korf R E 1985 *Learning to solve problems by searching for macro-operators* (Boston: Pitman)
- Khemani D 1988 *Theme based planning in an uncertain environment*. Ph D thesis, Department of Computer Science & Engineering, Indian Institute of Technology
- Khemani D, Ramakrishna R S 1989 Bridge: A benchmark for knowledge based planning. *Commun. Cogn. Artif. Intell. (CC-AI)* 6: 137-151
- Lehnert W G 1987 Knowledge-based natural language understanding. In *Exploring artificial intelligence - Survey talks from the National Conference on Artificial Intelligence - 1986 & 1987* (ed.) H E Shrobe (San Mateo: Morgan Kaufmann)
- Marks M, Hammond K J, Converse T 1988 Planning in an open world: A pluralistic approach. *Proceedings of Workshop on Case-Based Reasoning*
- Newell A, Simon H A 1963 GPS: A program that simulates human thought. In *Computers and thought*. (eds) E A Feigenbaum, J Feldman, (New York: McGraw-Hill)
- Noronha S J, Sarma V V S 1991 Knowledge-based approaches for scheduling problems: A survey. *IEEE Trans. Knowledge Data Eng.* 3: 160-171
- Ramachandran R, Khemani D 1992 Planning declarer play in bridge, Tech. Rep. IITM-CSE-92-02
- Rapoport A 1966 *Two-person game theory* (Ann Arbor: Univ. of Michigan Press)
- Sacerdoti E 1974 Planning in a hierarchy of abstraction spaces. *Artif. Intell.* 5: 115-135
- Schank R C, Abelson R 1977 *Scripts, plans, goals, and understanding* (Hillsdale, NJ: Lawrence Erlbaum)
- Shortliffe E H 1976 *Computer-based medical consultations: MYCIN* (New York: Elsevier)