Planning With Thematic Actions

Deepak Khemani
Department of Computer Science & Engineering
Indian Institute of Technology
Madras - 600036, India
email: khemani@iitm.ernet.in

Abstract

The task of planning in a dynamic and uncertain domain is considerably more challenging than in domains traditionally adopted by classical 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. However, many knowledge based implementations are susceptible to brittleness. 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 a more flexible use of knowledge structures for planning in bridge.

Introduction

Some of the criticism that classical planning (Fikes 71, Sacerdoti 77, Tate 77, Chapman 87, Wilkins 84) faces is that it makes some very simpifying assumptions. For example, it is assumed that the world is completely known and stable, and that the effect of any proposed actions can be predicted precisely (Marks, Hammond, & Converse 90) Moreover, it has been shown that the planning problem is NP-complete (Gupta 91), so that even for the toy worlds where these methods are tried the exponential amounts of time required is prohibitive. In other words, the problem of finding a plan *ab initio* from a start state is likely to be hard.

An alternative approach, which we may call knowledge based, endeavours to harness knowledge to charter a path through the combinatorially increasing number of possibilities. Various formalisms for expressing the required knowledge have been tried. Starting with rules (Wilkins 80), efficient retrieval considerations led to more organized structures like goals, plans (Schank & Abelson 77), (Wilensky 83), MOPs (Riesbeck & Schank 89), and cases (Hammond 89). We shall confine our interest here to case based planning. The main idea is to retrieve complete plans from memory, which contains the cumulative experience internalized into easily accessible structures.

In this paper we explore a hybrid approach. The motivation is to augment the heuristic power of the knowledge based methods with the flexibility of the weak methods. We feel that planning often involves a set of subproblems for which solutions can be retrieved from memory in the form of partial plans, which can then be suitably combined using a general scheduler. For instance in the domain of CHEF (Hammond 89), instead of adapting a beef-and-green-beans plan to produce beef-and-broccoli, we would advocate retrieving the knowledge of cooking beef and broccoli, and integrating it into a general stir-fry dish recipe. However, the domain of our present work is the game of contract bridge, as described in the following section.

Bridge

Contract bridge provides an excellent domain for the exploration of knowledge based methods in an uncertain environment (Khemani & Ramakrishna 89) (Pell 93). We get an environment with all the advantages of working with games, viz. a discrete finite universe which does not require compromises in representation, where symbolic interaction is complete, and where performance can be accurately measured. In addition, being a four-player incomplete-information game, bridge provides challenges 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 situations to tackle (as opposed to the same starting board position in chess and go). At the same time bridge provides a complex enough problem solving environment which has held humankind spellbound for decades. A more detailed description of the problems in the game is given in (Khemani & Ramakrishna 89).

In this paper we look at the problem of declarer play in no trumps. This avoids the more complex communication problems in bidding and defence. 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 (Chandrasekaran 86) 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, 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.

Retrieving Partial Plans

It can be reasonably claimed that a bridge hand once played will not be repeated. Nevertheless, one finds that human performance improves with learning. The kind of knowledge that players seem to employ is a repertoire of *standard plays*, for example finesses and squeezes, plans for which are reused in different combinations in new situations. We call the chunk of knowledge which captures the theme behind the actions of a standard play a *thematic act* (TA). A TA has knowledge of patterns to look for in the raw data, knowledge of move combinations (a partial plan), and knowledge of the overall effect of the moves (which we call *meta-semantic* knowledge). We explore the the nature of such thematic knowledge in the context of contract bridge.

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 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 : A Q
West : ?? East : ??
South : 3 2 (1)

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 then 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. Straight away 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 A K J, and one could have still finessed the Jack. What is more, if the A K Q 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 small card 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 equates 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 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 74), where the emphasis is only on ignoring unnecessary detail.

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 *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 x x x (3)

We need four tricks from this combination, and can afford to lose one (but not two). 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 one of the Q or 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 when one is willing to alter one's risk-gain equation. Thus, it may happen, in the above example, that in a larger context (eg. 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.

Planning With The TAs

The goal in bridge is to arrive at a plan to make a certain number of tricks. The dealt 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. We examine the role of TAs in forming the plan.

The basic strategy used in the planning process is Means Ends Analysis (MEA) (Newell 63). 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 TAs merely suggest partial plans to reduce a part or 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. 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 twice the play be started from the South position. Each time the lead shifts to the North hand, since the trick is won there. Some other suit needs to provide an 'entry' into the South hand to finesse again. The partial plan is not complete in itself, and may need to be meshed with other partial plans. 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.

Constraints

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 allows neither leeway nor tempo.

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 available 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 the TAs which is within the overall leeway and tempo constraints.

The Scheduler

The *task* of planning is to combine TAs 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 TAs 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 TAs 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.

Combining Strong and Weak Methods

The overall planning activity can be thus seen as following. The task is to order the play of two sets of thirteen cards. A partial order is imposed by a TA on the play of a suit. However, the partial plan for a suit may be dependent on the existence of a suitable partial plan for another suit. This happens for example in combination 4, where two entries in one hand are needed for if the partial plan is to be incorporated. This meshing of different partial plans is done by the scheduler, whose job is to essentially *merge* the partial plans for the four suits.

It can be seen that there are two distinct aspects to the task of planning. The first, which involves retrieving the TAs, is more knowledge intensive. Each partial plan is retrieved complete with its own ordering (Berlin 85). 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 assembling 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, 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 79) specifically while executing plans (Birnbaum 85).

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 contracts.

The program begins by abstracting the relevant patterns out of the cards. It counts its tricks, and computes the goal (i.e. the number of tricks that have to be generated), and the constraints, leeway and tempo. The program then looks for 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 then 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 moves from the partial plans is done in a heuristic manner by the scheduler. Thus the right play in the case referred to in (Berlin 85) is found directly because of the heuristics employed, rather than by a deeper analysis of interaction of subgoals. While assembling a plan, the scheduler picks moves from the partial plans using a few heuristics. 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 92).

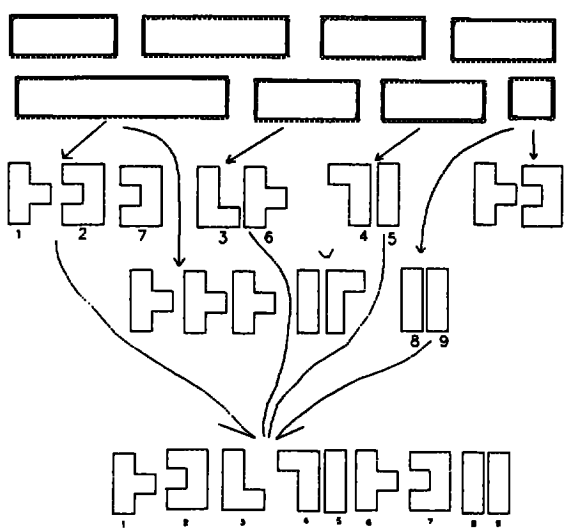

Fig 1 Two Stage Planning — A Schematic Flowchart

ANNOTATION

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

A Thematic Planner looks at subproblems and the TAs suggest partial plans. Numbers below show their position in the final plan.

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

Results

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

In Hand 1 the spade suit is identified as the source of tricks. The feature of using a delayed finesse to solve the communication problem has been encoded as a separate TA. In Hand 2 the assumed 3-2 break in spades does not materialize, and the program is forced to replan. Knowing that the spade finesse will fail, it plans to finesse in hearts. In different run when the East-West spade and heart suits were interchanged, the program *did* prefer the spade finesse knowing it will succeed.

Concluding Remarks

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 knowledge structures called thematic acts (TAs). In the second stage, a scheduler combines the actions suggested by the TAs into a coherent plan. This may often require that actions of different TAs be interleaved. The scheduler is basically an embodiment of a weak method, but one which operates within the order 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

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 are comments by the author.*]

Example 1

CONTRACT = 3 NOTRUMP

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

NORTH-DUMMY
♠AQ6543
♥J7
♦84
♣T65

WEST
♠KJ9
♥Q98
♦Q96
♣Q83

Opening lead-
K♦

EAST
♠T7
♥T632
♦T52
♣J972

SOUTH
♠82
♥AK54
♦AJ73
♣AK4

Example 1

5 tricks in Spades, 2 tricks in Hearts, 1 tricks in Diamonds and 2 tricks in Clubs.
Examining lead ... Looks like a honour sequence.
The play ... A * indicates that the corresponding player had the lead

S.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
****

Example 2

CONTRACT = 7 NOTRUMP

The planning phase ... Have 11 top tricks .. Need 2 tricks more ... Tempo available 0 If Spades break 3 -2 we get 1 extra tricks 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 tricks losing 0 .. first glance over Have planned for 4 tricks in Spades, 1 tricks 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

NORTH-DUMMY
♠AK32
♥AQT98
♦J32
♣Q

WEST
♠J975
♥K3
♦7
♣T7652

Opening lead-
J♣

EAST
♠4
♥J6542
♦T965
♣943

SOUTH
♠QT86
♥7
♦AKQ84
♣AK8

Example 2

tricks losing 0 .. first glance over [*Note that it does not look for spades as it knows that the break is not favourable*] Tenace over East found in Spades Finesse not possible as opponent has shown out Tenace over West found in Hearts Have planned for 1 [*more*] tricks 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

S.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

****

References

Berlin, D.L.S. 1985. SPAN: Integrated Problem Solving Tactics. In Proceedings IJCAI-85. pg 1047.

Birnbaum, L. 1985. A Short Note on Opportunistic Planning and Memory in Arguments. In Proceedings IJCAI-85. pg 281.

Chandrasekaran, B. 1986. Generic Tasks in Knowledge-Based Reasoning: High-Level Building Blocks for Expert System Design. *IEEE Expert, vol. 1.* pg 23.

Chapman, D. 1987. Planning for Conjunctive Goals, *Artificial Intelligence, vol. 32.* pg. 333.

Fikes, R., and Nilsson, N. 1971. STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving. *Artificial Intelligence, vol. 2.* pg. 189.

Gupta, N., and Dana, S.N. 1991. Complexity Results for Blocks-World Planning. In Proceedings AAAI91. pg 629.

Hammond, K. 1989. *Case-based Planning: Viewing Planning as a Memory Task*, Perspectives in Artificial Intelligence, Academic Press, Boston, MA.

Hayes-Roth, B., and Hayes-Roth, F. 1979. A Cognitive Model of Planning, *Cognitive Science, vol. 3.* pg 275.

Khemani D., and Ramakrishna R.S. 1989. Bridge: A Benchmark for Knowledge Based Planning, *CC-AI, Communication and Cognition, vol. 6, No. 2/3.* pg 137.

Marks M., Hammond K.J., and Converse, T. 1988. Planning in an Open World: A Pluralistic Approach. In Proceedings of Workshop on Case-Based Reasoning.

Newell, A., and Simon, H.A. Eds. 1963. GPS: A Program that Simulates Human Thought, in : Computers and Thought, Feigenbaum, E.A., and Feldman, J. McGraw-Hill, New York.

Barney, Pell. 1993. Pragmatic Reasoning in Bridge, Technical Report No. 299, Computer Laboratory, University of Cambridge.

Ramachandran, R., and Khemani, D. 1992. Planning Declarer Play in Bridge, Tech. Report IITM-CSE-92-02, Dept. of CS&E, IIT Madras.

Riesbeck, C., and Schank, R. 1989. *Inside Case-based Reasoning*, Lawrence Erlbaum Associates, Hillsdale, New Jersey.

Sacerdoti, E. 1974. Planning in a Hierarchy of Abstraction Spaces, *Artificial Intelligence, vol. 5*, pg 115.

Sacerdoti, E. 1977. *A Structure for Plans and Behavior*, Elsevier, North Holland.

Schank, R.C., and Abelson, R. 1977. *Scripts, Plans, Goals, and Understanding*, Lawrence Erlbaum Associates, Hillsdale, New Jersey.

Tate, A. 1977. Generating Project Networks. In Proceedings IJCAI-77, Cambridge, Massachusetts, pg. 888.

Wilkins, D.E. 1980. Using Patterns and Plans in Chess, *Artificial Intelligence, vol. 14.* pg. 165.

Wilkins, D.E. 1984. Domain-Independent Planning: Representation and Plan Generation, *Artificial Intelligence, vol. 22.* pg. 269.