Applying Adversarial Planning Techniques to Go

Steven Willmott [a] Julian Richardson [b] Alan Bundy [b]
John Levine [c]

[a] *Laboratoire d'Intelligence Artificielle,*
École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland
email: willmott@lia.di.epfl.ch

[b] *Institute of Representation and Reasoning*
Division of Informatics, University of Edinburgh, Scotland
email:{julianr, bundy}@dai.ed.ac.uk

[c] *Artificial Intelligence Applications Institute*
Division of Informatics, University of Edinburgh, Scotland
email: johnl@aiai.ed.ac.uk

Abstract

Approaches to computer game playing based on $\alpha-\beta$ search of the tree of possible move sequences combined with a position evaluation function have been successful for many games, notably Chess. Such approaches are less successful for games with large search spaces and complex positions, such as Go, and we are led to seek alternatives.

One such alternative is to model the goals of the players, and their strategies for achieving these goals. This approach means searching the space of possible goal expansions, typically much smaller than the space of move sequences. Previous attempts to apply these techniques to Go have been unable to provide results for anything other than a high strategic level or very open game positions. In this paper we describe how adversarial hierarchical task network planning can provide a framework for goal-directed game playing in Go which is also applicable both strategic and tactical problems.

Key words: Computer Games, Go, Planning, Adversarial Planning

1 Introduction

Most approaches to computer game playing are based on game tree search and position evaluation functions (*data-driven* approaches). Data-driven approaches are appropriate for games with low branching factors, and for which it is possible to accurately assign values to positions which indicate who is winning. While this approach has been very successful for many games including Chess, it has been less successful when applied to games with high branching factors and complex positions, such as Go or for games with a high degree of uncertainty such as Bridge.

An alternative to the data-driven approach is *goal-driven* search, in which a single agent searches for a way to satisfy its goals in the game. Goal-driven search has been extensively explored in the Artificial Intelligence literature, in particular as Hierarchical Task Network (HTN) planning [42,15]. When multiple agents need to be modelled and can compete against one another, this approach becomes *adversarial planning*. In this paper we describe how adversarial hierarchical task network planning can provide a framework for goal-directed game playing in Go. We consider different types of goal expansion in HTN planning and suggest how to make planning approaches more applicable to tactical problems.

1.1 Paper Overview

In §1.2, we review goal-driven and data-driven approaches to making move choices in game playing. §1.3 presents an overview of previous work on computer Go (§1.3.2) and Go's search space (§1.3.1). §1.4 completes the first section with an outline of previous work in the field of adversarial planning. The main body of the paper includes the following two principal contributions:

- The description of a new adversarial planning framework (§2.2), based on hierarchical task network planning (§2.1). The framework models two agents, each of which attempts to satisfy its own goals while refuting those of its opponent. This behaviour is achieved by backtracking, and the enforcement of a time linearisation during goal decomposition (discussion of the decomposition strategy can be found in §2.3).
- The application of the planning framework to the game of Go (§3). In order to prove the concept without coding large amounts of knowledge, our implementation, GoBI, focuses on life-and-death problems of the type found in Go teaching books (although there is no restriction to *enclosed* problems).

GoBI was systematically tested on examples from volume I of "Graded Go Problems for Beginners" [49], finding the correct answer for 74% of the problems. We analyse both successes and failures in §4.4.

The remainder of the paper examines the advantages and disadvantages of applying a goal-driven approach to computer Go (§5), and compares our work to adversarial planning systems in other domains (§5.4).

Within a specific game, move or action choices often depend upon the state of the game, the phase of the game (e.g. opening, endgame etc.), the future actions of the opponent, the ability of a player to follow up an action appropriately and many other diverse factors. It is these interacting influences on the choice and effect of moves which make games so fascinating for human players and so challenging for machines.

In computer game playing there are two main approaches to making move choices:

- **Data-Driven:** At each step, rules, patterns or heuristics are applied to the game state to suggest useful moves. The resulting set of plausible actions is then evaluated using search in the tree of moves. Each move is played out in a world model followed by the possible responses of the opponent. The search continues until the leaves of the tree are reached. [1] These leaf nodes are then evaluated and used to select one of the original plausible actions as the one which leads to the most desirable (by some measure) set of leaf states.
- **Goal-Driven:** During play, a goal-driven system keeps a number of abstract goals [2] in an agenda. The goals in the agenda represent the things the system would like to achieve in the short, medium and long term. To choose a move, goals are expanded into plans (which are conjunctions of goals at lower levels of abstraction) and eventually into concrete moves (actions in the world). Repeated decompositions form a plan for achieving the goal.

In a data-driven search tree, each node represents a possible game position and has one branch for every move suggested in that position. In contrast, each node in a goal-driven search tree represents a *plan* for achieving the top level goal with some parts still sketchy (abstract goals) and others fixed (concrete actions), and each node in the search tree has one branch for each way the system suggests to further refine the plan.

Which approach (goal-driven or data-driven) is most advantageous is heavily dependent upon the domain, in particular on the size of the data-driven and goal-driven search trees. In Bridge, for example, the locations of the cards are not in general known during play, which leads to a large space of possible card plays and therefore a prohibitively large data-driven search tree. [3] Smith et. al.

[1] Which nodes are the "leaves" can be variously defined by a depth cut off point, quiescence, or further domain dependent heuristics.
[2] Abstract goals are aims which cannot in general be achieved by a single primitive action in the world.
[3] A data-driven search can still be practicable, but only by dramatically limiting

(in [38]) show that a goal-driven approach can very successfully play Bridge, and work described in [18] demonstrates that a relatively small number of operators is sufficient to describe all the relevant plays.

1.3 The Game of Go

The game of Go is considered by many to be the next great challenge for computational game playing systems. It presents new, significant and different challenges to Chess which has been long been considered the "task par excellence" for Artificial Intelligence and computer game playing [3]. A good introduction to the game can be found in [4].

1.3.1 Size of the Go Search Space

The search space of Go is both wider and deeper than that of Chess; there are estimated to be about 10^{170} states (cf. Chess $\approx 10^{50}$), games last approximately 300 moves (cf. Chess ≈ 80) and the branching factor at each turn is on average 235 states (cf. Chess ≈ 35). The number of possible games of Go is estimated at 10^{700} compared to $\approx 10^{120}$ for chess [5]. It is often hard to evaluate the relative strength of Go positions during play, since guiding principles such as value of material which are very useful in chess are often misleading for Go. The difficulty of Go is as much due to this difficulty in state evaluation as to the large raw search space. Useful complexity results have also been shown:

- [24] shows that determining the eventual winner for an arbitrary Go position on an $n \times n$ board is PSPACE-hard. [4]
- [33] extends this result to show that deciding whether or not black can win from an arbitrary position is EXPTIME-complete (taking into account situations arising from Ko positions).

Given these difficulties, the brute-force game tree search which has been so effective for Chess will potentially have much greater difficulty with Go.

1.3.2 Approaches to Computer Go

Although Go has received far less attention than Chess in terms of research, there have been many varied approaches to computer Go:

---

the part of the move space which is searched. For example, Ginsberg's GIB [20] adopts a sampling approach, aiming to sample a representative part of the search tree, and choosing its moves on this basis.

[4] [41] gives proofs for a similar result for a class of generalisations of Chess to $n \times n$ boards.

- Hybrid approaches such as Go 4++ [6], Many of Faces of Go [16] and Handtalk [6] are by far the most successful at playing the complete game to date. These systems are mainly data-driven (Go 4++ for example, works almost exclusively on pattern recognition) but their long period of development (10-15 years for Many Faces of Go) has seen the addition of many other types of reasoning and specialist modules.
- Non-symbolic techniques have been used to learn/evolve controllers and rules based upon patterns of stones for use during play. The techniques applied include Genetic Programming [10], Genetic Algorithms [32,22,12], and Neural Networks [13]. These approaches have so far been less successful than the hybrid programs but have the advantage that Go knowledge does not need to be added by hand.
- Cazenave's GoGol [8] applies learning techniques to good effect. An off-line program uses introspection to prove theorems about Go tactics which can be used to generate knowledge useful in pruning Go search trees. This knowledge is then added into GoGol's playing procedures. GoGol was the top ranked non-commercial program [9], finishing 6th out of 40 participants in the IJCAI'97 international computer Go tournament [17].
- Progress has also been made by focusing on specific subproblems in Go. Wolf's GoTools [47,48] uses very deep search to analyse closed "life and death" game positions and can often solve very complex problems. [26] applies combinatorial game theory to endgame situations, which enables precise calculation of the values of sub games and hence perfect play, but is intractable except for the last few moves of the endgame.
- There have also been several applications of planning techniques to Go. These systems are discussed separately in §1.4.

The top Go programs are now able to challenge good amateur players, but there remains substantial room for improvement. The advantages mentioned in §5.1, earlier work on Go planners [37,23], and the success of the goal-driven approach in other domains (notably Bridge [40]) suggest that a goal-driven approach may be useful. It also has much psychological validity, since protocol analysis indicates that Go players consider few candidate moves, and concentrate on their own and on their opponents' purposes [36]. Finally, even in data-driven approaches to computer Go, it is still necessary to consider high-level goals, for example in order to decide whether or not a satisfactory result has been achieved in life and death problems (e.g. some strings may be allowed to die but others must live) [48].

1.4 Applications of Adversarial Planning

Domains in which agents compete against one another are termed *adversarial*. The first attempts to use goal-driven reasoning in adversarial domains include work by Pitrat [29] and Carbonell [7]. The former was extended by Wilkins in [43] to produce the PARADISE system for Chess. Using goals to guide search, PARADISE was able to find moves which involved looking as far ahead as 20 ply in certain situations, a depth well beyond its contemporary search-intensive competitors. More recent work studies battlefield management [1], command and control [50], and Bridge [39,40]. The work on Bridge is perhaps the most successful application of goal-driven reasoning to games to date, presenting a system for Bridge declarer play (TIGNUM2) good enough to beat the current top commercial computer player [39]. [19] describes the FINESSE program for Bridge which also applies a plan-based approach. Extensive testing has shown that FINESSE can find optimal plans for single-suit play, and correctly calculate their probability of success.

Go has been used as an application domain for adversarial planning systems:

- The INTERIM.2 Go program [30] was probably the first Go program to use the idea of goals to guide move choice.
- This initial interest was followed by two Go planners due to Sander and Davies [37] and Lehner [23] which both addressed only very open positions.
- More recently, Hu [21] also concentrates on high level strategy looking at the possible use of multipurpose goals.
- Finally the GOBELIN system developed by Ricaud [31] also has aspects of a goal-driven approach. GOBELIN uses an abstract representation of the Go game state to form plans before mapping these back into the ground state.

A common feature of all [5] of the goal-driven Go programs to date is that although some are able to report good results for high level strategic planning [37,23] or during opening play [31], none adequately addresses the tactical level of Go play. There has been little success for this type of approach in the middle or end-game where tactical and strategic considerations become closely linked. Contributions of this paper include offering an explanation for this lack of progress, based upon the type of decomposition strategy used during planning, and presenting an approach which makes goal-driven reasoning applicable at both the strategic and tactical levels.

2 An Adversarial Planning Architecture

This section describes an adversarial planning architecture which models goal-driven reasoning for adversarial domains. [6] The goal-driven approach and use

[5] Except INTERIM.2 which applies data-driven search for local tactical lookahead.
[6] More details (including a precise breakdown of the planning algorithm) can be found in [45].

of abstract plans is motivated by work on hierarchical task network (HTN) planning. HTN systems were first used in NOAH [34] and INTERPLAN [42] and have since been extensively studied in the AI planning field. Erol et. al. [15] give a complete definition for an HTN scheme and present UCMP, which is a provably sound and complete HTN planner and provides a good template for this type of system.

2.1 Principles of HTN Planning

HTN planning is based on three types of object: *Goals*, *Operators* and *Plan Schemas*. Operators are actions which can be performed in the world (such as flicking a switch, taking a step). Goals are more abstract and express aims in the world such as "Go to the Moon", "Become Prime Minister". Schemas (also called Task Networks or Methods), specify the subgoals which must be achieved in order to satisfy the goal. For example, the following schema expresses the fact that G can be achieved by satisfying the conjunction of subgoals G_1, G_2 and G_3:

$$G \Rightarrow G_1 + G_2 + G_3$$

The G_i should be at a lower level of abstraction than G, and can generally be satisfied in any order. Operators are at the lowest level of abstraction.

Given these three types of object, HTN planning starts with an initial world state and a set of goals which form the initial abstract plan. The plan is then refined step by step by expanding the goals within it. Goals are expanded by selecting a schema whose antecedent (the G above) matches the chosen goal, and replacing the instance of G in the current plan by the subgoals (the G_i above) listed in the consequent of the schema.[7] As the planning process continues, interactions, incompatibilities and conflicts may arise between combinations of goals. These "interactions" in the plan must be resolved, which can result in backtracking and (in partial order planners) ordering constraints between goals.

The process is complete when all goals have been expanded into sets of operators and all arising interactions have been resolved. The sequence of operators thereby generated should, upon execution in the initial world state, lead to the achievement of the planner's goals in the world.

The characteristics of an HTN planner are principally determined by its goal expansion strategy. At any stage a plan can be *totally ordered*, which means that every step is time ordered with respect to every other step or *partially*

[7] Some planners may also instantiate variables as part of this expansion process which adds extra complexity [14].

ordered, in which case some steps do not have ordering constraints between them. A partial order planner allows goals to remain unordered until ordering becomes necessary whereas a total order planner chooses a linearisation over the plan steps to ensure that the steps are always totally ordered. Planners may also use two different inference mechanisms: backward chaining and forward chaining. Backward chaining planners begin with the goal state and use their knowledge to work backwards towards the start state, forward chaining planners do the opposite — starting with the initial state and applying sequences of operators to reach the goal. [8] Most standard HTN planners in use are backward chaining partial order planners. [9]

The extension of this idea into adversarial domains is non-trivial since plans are no longer sequences of actions but trees of contingencies which take into account the actions of opponents. The interactions in the plan are considerably more complex and serious since the goals of opponents in the world are often conflicting and the planning agents are non cooperative. HTN planning for adversarial domains is therefore computationally considerably more complex than HTN planning in standard domains.

2.2 Adversarial Planning Framework

The adversarial planner presented in this paper models two agents (named Alpha and Beta) which represent two players (*adversaries*) in a game. [10] To solve a problem in the domain, each agent is given a set of input goals to achieve. We assume that the input goals of the two agents are mutually exclusive, so that a successful plan for one agent to satisfy its input goals acts as a refutation of the other agent's input goals. [11] Alpha and Beta's tasks are to achieve their own goals while preventing the other agent from achieving theirs.

Each agent keeps an open agenda of goals (descended from the input goals by expansion steps) which represents its current plan of action. Using these plans the agents attempt to find sequences of moves which satisfy their input goals. As the agents take turns to try and achieve their goals, their interaction in the world can be used to find a plan for the problem.

[8] There is a distinction between *forward* and *backward chaining* and *data* and *goal driven*: forward and backward chaining describe ways of traversing the search space, whereas data and goal driven relate to the motivations behind move choice.
[9] In terminology we follow [25].
[10] The framework can be generalised to include more than two players.
[11] We could usefully relax the restriction that the input goals of the two agents be mutually exclusive, by instead requiring only that when one agent satisfies its input goals, the other agent's goals become less desirable.

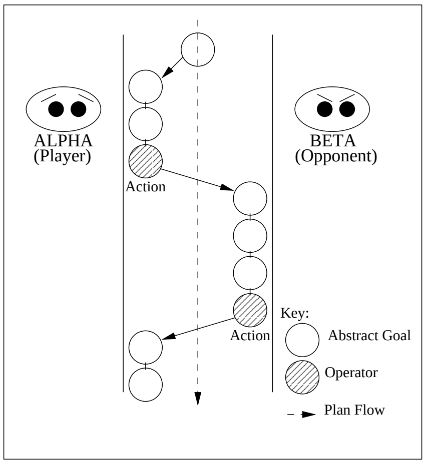

Fig. 1. Planning steps alternating between two agents.

Both agents (Alpha and Beta) share the same knowledge base of plan schemas. For any given problem the two agents are likely to use different parts of the knowledge base (one using attacking plans, the other using defensive plans for example), although overlaps are not excluded (some attacks are powerful defences). The system allows the two agents to take control of the reasoning apparatus in turns (Alpha takes control first). Once an agent has control, it expands some of its abstract goals until it is able to decide upon a concrete action. The chosen action is then performed in a world model [12] before control is passed to the other agent. Figure 1 shows the flow of control during the reasoning process. An agent may need to expand several abstract goals before being able to decide upon an action in the world. During this "active" period, it uses its own agenda of goals and has control of the shared reasoning apparatus. Once an action is chosen, control passes to the other agent.

At any one time an agent has a plan which consists of actions already taken (square boxes in figure 2) and goals at various levels of abstraction (circles in figure 2). The actions (squares) are represented in the world model, the abstract goals (circles) are held in the agenda.

A planning step involves selecting an abstract goal (such as X in figure 2) and expanding it. A plan schema is selected for X which expresses how X could be achieved using a conjunction of subgoals at a lower level of abstraction. For example, in figure 2, X is replaced in the plan by the two subgoals X1 and X2. Once expansion has reached the lowest level of abstract goals, these lowest level goals need to be shown to be already true or replaced by actions which make them true.

Once one of the agents (Alpha say) has achieved all of its goals (been able

[12] A world model is not a standard feature of HTN planners — see §2.3.2 and §5.4.3 for more explanation of its use.

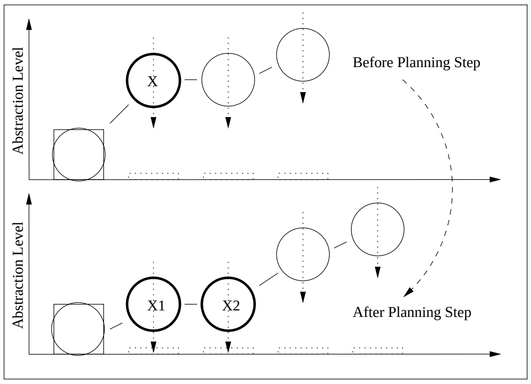

Fig. 2. Plan refinement: abstract goals are expanded to and replaced by sets of subgoals at lower levels of abstraction.

to perform actions in the world model which make them true), it knows that it must have satisfied its top level goals (since all its subgoals are strictly descended from these). The opposing agent is made aware of this fact and may force backtracking.[13] Backtracking is only forced by an agent when the other agent achieves *all of its input goals*. Thus one agent can allow the other to achieve any number of subgoals while pursuing its own agenda. This trading of subgoal achievement ("while my opponent is achieving X, I will achieve Y which seems more important to me") is a vital part of many games. Agents are allowed to backtrack to any of their previous goal or expansion choices but only to their *own* decisions. Neither agent may force the other to change plans directly.

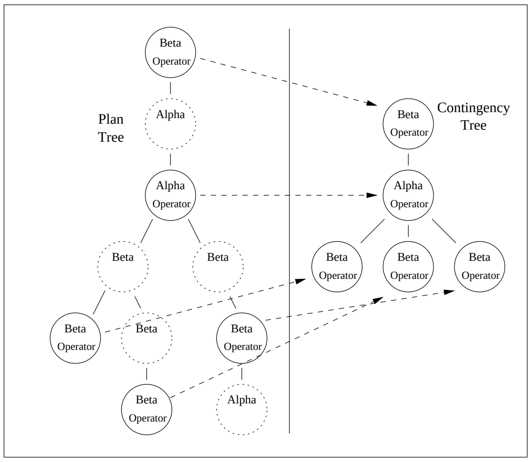

Fig. 3. The plan tree on the left is reduced to the contingency tree on the right by dropping the abstract reasoning nodes.

The backtracking activity explores the various interacting plans Alpha and Beta have for the situation and creates a plan tree as shown on the left of

[13] This rule expresses the fact that Alpha and Beta's goals are in some sense mutually exclusive (and the agents must prevent each other from satisfying their respective goals).

figure 3. Each choice made by an agent creates a new branch. Underlying the plan tree is the *contingency tree* which is found by removing all the abstract goal decomposition steps in the plan tree to leave only the operators/actions (shown on the right in figure 3). Moves in the contingency tree are *directly descended* from the goals of the two agents, and the tree structure naturally reflects the interactions between the two adversaries. Taking any branch, the moves chosen near the leaf (at the end of a move sequence) descend from the same goal as those near the root of the tree, and therefore serve the same purpose.

Since all search is based upon the knowledge available, the knowledge base can be used to control the extent (time taken) of a particular search. Restricting the amount of knowledge (plan schemas) taken into account by the planner leads to less complete plans/counterplans but reduces the amount of time taken during search.

The final contingency tree acts as a form of proof that the first move is a good step towards achieving Alpha's goals. Hence it supports the choice of the first move in the tree. [14] In general, the final contingency tree contains only a small subset of the move tree which would be generated by considering all the available options at each turn as in a data-driven approach. (See §5.1.)

2.3 Goal Decomposition Schemes and a World Model

To better describe how the planner works, this section concentrates on its goal decomposition scheme. The method of goal decomposition defines how the planner applies its knowledge to problems and defines the fundamental characteristics of the planner.

2.3.1 Standard HTN Goal Decomposition

The standard scheme for HTN decomposition (widely used to good effect in non-adversarial domains) is as shown in Algorithm 1. [15]

At each refinement step a goal is chosen and decomposed to reduce the overall level of abstraction in the plan. In general all goals are kept at roughly equal levels of abstraction. The average abstraction level of the goals in the plan

[14] The tree itself can also be used to respond to any of the opponent's moves which are represented in it, but replanning may be required if other moves are made. The question of how long (for how many moves) such a plan remains valid is not addressed here.
[15] This algorithm is a simplified version; for more detailed coverage, see [15].

procedure *Refine-Plan*

  while *Not Empty(Agenda)* **do**

    1. choose a goal $NextGoal$ from $Agenda$ to decompose

    for *goal* $NextGoal$ **do**

      2. select a decomposition schema S

      3. remove goal $NextGoal$ from $Agenda$

      4. insert schema S into $Agenda$

Algorithm 1: At each step one goal is chosen and decomposed. ($Agenda$ is the list of open goals held by each agent.)

then gradually decreases as the goals are decomposed, eventually reaching the stage where all goals are concrete actions in the world (primitives) and the plan can be executed. This approach is a purely backward chaining strategy in the space of plans. The system starts with the top level goal state and systematically fleshes out the plan using plan knowledge. The algorithm does not specify any particular subgoal ordering, and therefore is suitable for both total and partial planning. A total order planner would additionally make execution order choices in step 1.

At each level of abstraction (after each decomposition) the relationships and interactions between goals can be used to guide the choice of schemas. In a partial order setting, interactions between goals would lead the planner to constrain available orderings.

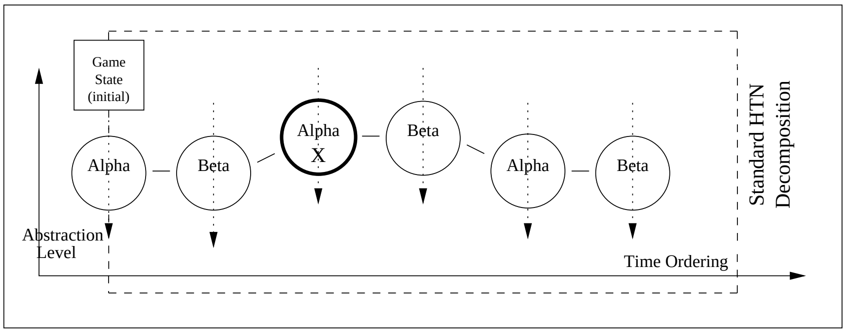

Fig. 4. Goal decomposition in "standard" HTN planning. The dashed arrows represent the decomposition of goals (eventually) into concrete actions.

The dashed box in Figure 4 represents the items the planner can reason about when decomposing the goal marked X. In this standard HTN model, the planner can reason about all the goals (and hence their interactions) at their current level of abstraction and about the world in its *initial* state. The effects of actions subsequent to the initial state and how goals might interact must be modelled explicitly by the plan knowledge.

2.3.2 Modified HTN Goal Decomposition

In the adversarial planning architecture presented here, goals are expanded *in time order* using a linearisation. The linearisation is used to make goals concrete as soon as possible and model the effect of the resulting actions in a "world model". This linearisation behaviour contributes to the planner's ability to reason simultaneously and consistently about both abstract (strategic) and concrete (tactical) aspects of the game. As noted above, a world model is not a standard feature of HTN planning systems and goal expansion is quite different from the standard version described above. Goal expansion in the modified approach is shown in Algorithm 2.

```
procedure Refine-Plan-(Modified)
  while Not Empty(Agenda) do
    1. choose a goal NextGoal from Agenda to decompose
    while NextGoal has not resulted in a primitive action do
      2. select a decomposition schema S for NextGoal (taking into
      account the world model)
      3. remove goal NextGoal from Agenda
      if S describes a primitive action then
        4. perform the action in the WorldModel
      else
        choose a goal DescendentGoal from schema S
        5. let NextGoal = DecendentGoal
        6. insert schema (S – DecendentGoal) into Agenda
```

Algorithm 2: This expansion algorithm is similar to to the standard expansion algorithm (1) but additionally ensures that one of the subgoals descended from *NextGoal* produces an action in the world before another *NextGoal* is chosen.

There are three key differences between the two decomposition strategies:

- In the modified algorithm, steps 1, 4 and 5 contribute to the linearisation of the plan by forcing the planner to choose a time order for goal expansion (and execution). Once a goal has been chosen for expansion the planner works on its descendent subgoals until at least one of these can be used to choose an action in the world.
- When actions are chosen, their effect is played out in a world model. The information in the world model can subsequently be applied to the choice of plan schema (step 2).
- The planner is performing backwards chaining in the *space of plans* (starting with the top level goals and decomposing) but with an element of forward chaining in the *space of moves*. That is, the initial state and subsequent

modelled states clearly show which options are available for the achievement of low level goals. Plans and world state come together to choose the move.

This approach necessarily leads to a total-order planner. At each stage there can be goals of all abstraction levels in the plan, but there must also be a linearisation, and goals are decomposed in time order. Once a primitive action has been performed, the subgoals (generally at various levels of abstraction) which remain on the agenda are retained and may be expanded next time the agent gains control of the planning apparatus.

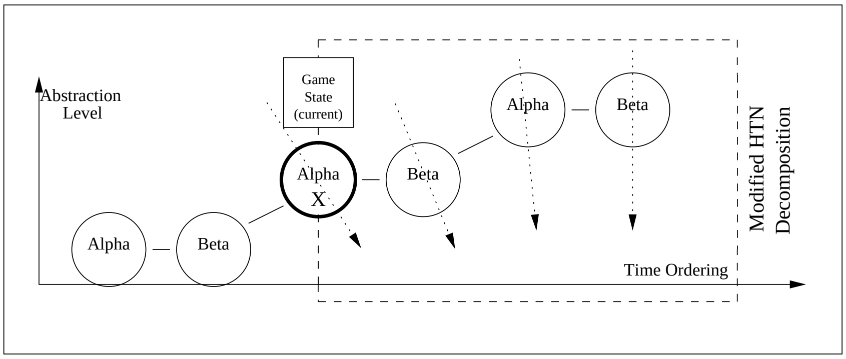

Fig. 5. Goal decomposition in modified HTN planning. The dashed arrows represent the decomposition of goals (eventually) into concrete actions.

Figure 5 illustrates goal decomposition for the modified HTN model. Again, the dashed box represents the items the planner can reason about when decomposing the goal marked X. The decomposition progresses left to right in time order; here the planner can reason about the outstanding abstract goals and about the *current* state of the world. [16] The first two primitive actions (one chosen by Alpha and one by Beta) are now represented in the world model and their effect on the state of play clearly seen.

The planning process can be seen as guiding search using goals, similar to the way PARADISE applied chess knowledge (§5.4.1). This process helps to choose the moves, but the complex interactions between goals are partly modelled in the world model. As the planner traverses the space of plans (explores the different options for attack/defence that might apply), it also traverses the space of moves. For comparisons with other planning systems see §5.4.

3 An Adversarial Go Planner

The planning architecture was instantiated as a Go reasoning system called GOBI both to test the architecture in a challenging domain and to investigate

[16] Note that the dashed box can be enlarged again upon backtracking.

the usefulness of the goal-driven approach for Go. GoBI consists of a set of knowledge modules which plug into the planning architecture. The knowledge modules provide the Go domain knowledge, plan schemas and goal types which the reasoner can use to solve problems.

3.1 An Example Go Plan

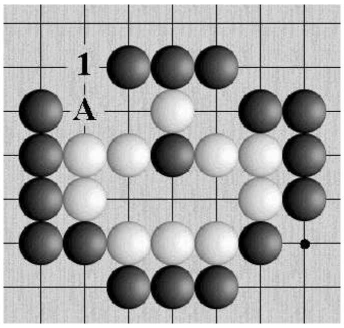

Fig. 6. Example A: Black to play and kill the white group.

Figure 6 shows a problem from Volume I of "Graded Go Problems for Beginners" [49]. The aim is for black to move first and kill the white group of stones. The task is specified to GoBI as two abstract goals: the goal *kill-group* for agent Alpha (playing black) and the goal *save-group* for agent Beta (playing white). [17] Agent Alpha takes control first, decomposing the *kill-group* goal using one of the available plan schemas. An abstract plan for killing this group might be a conjunction of the following subgoals: [18]

- *surround-group* — stop the group from running and connecting.
- *squeeze-space* — reduce the space the group has to live.
- *prevent-eye-formation* — block any attempt by the group to make eyes. [19]

One of these subgoals is then expanded further to the next level and so on until at the lowest level in the hierarchy a move such as *play at A* is chosen to satisfy a simple goal such as *prevent-escape-at-1* (figure 6).

Alpha plays the move onto the board in the world model which gives the new world state for Beta to work with. Alpha still has a set of goals at various levels of abstraction remaining in its agenda. These remaining goals represent the plan on how to follow the first move, i.e which other subgoals/actions need to be achieved to make the plan complete. To validate that this first

[17] The goals need not be directly opposing.
[18] This abstract plan is quite intuitive. It is not obvious how a data-driven system would represent the equivalent of such a plan.
[19] An eye in Go is an enclosed space where the opponent may not play — a group with two eyes is unconditionally alive.

move is good (in this case playing at A would not be), Alpha must eventually show that all these goals can be achieved no matter what Beta does. These remaining goals are kept by Alpha until after Beta's turn.

Beta now begins by expanding its single goal *save-group* in the context of the new board position (after Alpha playing at A in Figure 6). A possible plan schema for this goal is:

- *make-eye-space.*
- *make-eyes* (try to form two eyes).

After Beta's move is played into the world model, control is returned to Alpha which then tries to satisfy the rest of its goals. The interleaving of goal expansions by the two agents continues until one is able to satisfy all of its aims (and thus implicitly its main aim). The opposing agent is informed of this fact and then backtracks to explore any alternative options it has which might produce a better outcome for itself. In this way a contingency tree is generated which either proves or refutes the validity of the first move (Alpha's).

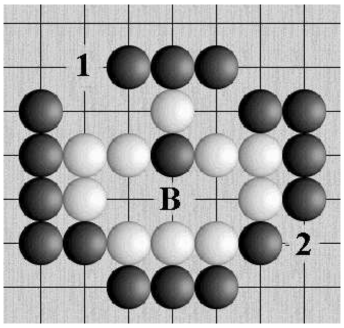

Fig. 7. Example A: GoBI plays at B and this kills the group.

For this example (figure 6) GoBI returns the move at B in figure 7, which kills the group. Among the defences tried by Beta are trying to run out at 1 and counter-attacking by playing at 2 (which puts the single black stone in the bottom right hand corner under threat). Since all the moves tried by both agents must be *part of plan of action*, the number of possible moves searched is very small compared to the number of available moves (which is quite large even in this small problem).

3.2 Representing Go Knowledge

The planning architecture and Go knowledge modules which make up GoBI are all written in Common Lisp. Around 1400 lines of code make up the plan knowledge (i.e., schemas and goal types) GoBI has. Writing a full-board Go-playing program is a significant exercise in knowledge engineering, so to enable us to add enough depth in knowledge to do useful testing in a short time, GoBI's knowledge is focused on the area of killing and saving groups.[20] The knowledge base is made up of 45 goals at five different levels of abstraction. The average number of applicable plan schemas per goal is approximately two (thus the knowledge has a relatively low branching factor). The two highest level goals available in GoBI are *kill-group* and *save-group*, which emphasises the focus on life and death problems.

The following example plan schemas taken from GoBI's knowledge base illustrate the structure of the knowledge the program holds (the *save-group* goal was also mentioned in the previous example). Note that the plan knowledge is not complete (for example making eyes is not the only way of following up a counter attack); more work is needed to extend the knowledge base.

```
GOAL: save-group,
LEVEL = 5,

 Schema 1 - Find Eyes:
  *make-eye-space,
  *make-eyes.

 Schema 2 - Escape Group:
  *running-sequence,
  *secure-escape.

 Schema 3 - Counter Attack String:
  *locate-vulnerable-string,
  *kill-string,
  *make-eyes.
```

In turn, the *make-eye-space* goal from `Schema 1` has two alternative schemas:

```
GOAL: make-eye-space,
LEVEL = 4,

 Schema 1 - Ambitious Extend:
  *large-extending move,
  *consolidate-space.

 Schema 2 - Creep extension:
  *creeping-extending-move,   //A single
  *consolidate-space.         //step extension.
```

The structure of the knowledge shown here is intuitive for Go and is very

[20] This does not mean GoBI is limited to enclosed problems (see §4.5).

different from the kind of heuristic information used in data-driven approaches. The knowledge represented can be seen as an AND-OR tree with the AND component represented in the conjunction within the plan schema and the OR component represented in the choice between alternative schemas. Goals within plans are not evaluated at every step. Simple goals are established by finding a satisfying move or using a test. For higher level goals, truth is inferred from the fact that all the goals descended from them were achieved.

3.3 The Planning Process

This section provides some more details of how the Go instantiation of the planning architecture works.

3.3.1 Choosing Top Level Goals

To be able to plan the planner needs to have goals to achieve. In GoBI these goals are given by hand. In order to extend GoBI to whole-board play, one of the most important tasks is to define goals and abstract plans above the current highest level. The planner may keep these goals open (and still unexpanded) from one turn to another and indeed for long periods of the game. Since shorter term tactical goals for choosing individual moves are also expressed in terms of goals (these are the goals GoBI already has) this should provide for a smooth transition between the tactical and strategic levels of Go.

The current top level goals in GoBI are *save-group* and *kill-group* which is quite limiting. Potential additional higher level goals include, for example: *make-moyo*, *prevent-invasion*, *extend-influence*, *extend-territory*, *hold-corner*, right up to *win-game*.

3.3.2 Planning Steps

As described in §2 during the reasoning process a planning step is made up of two parts:

(1) Choosing a current goal to replace in the plan and
(2) Choosing an appropriate plan schema to apply to the goal.

In GoBI, the choice of which goal to expand next is unordered, but after choosing a goal from the agenda, GoBI uses a linearisation and continues working on the same goal down through several levels of expansion. Goal expansion follows the algorithm presented in §2.3.2. Once a goal has been chosen in step 2 of the algorithm, one of its associated set of plan schemas must be selected for use. The plan schemas in GoBI are tried in a fixed order designed to try the most promising goal decompositions first. Plan preconditions can also be used to screen out schemas not thought to be suitable for the current situation. If one schema leads to a dead-end, backtracking allows GoBI to try the others available for the situation.

The expansion of goals into plan schemas eventually leads to goals at the lowest level of abstraction. These goals need to be checked for satisfaction and used to choose moves. Some example lowest level goals are:

- *fill-a-liberty*,
- *play-a-hane-move* (near *here*),
- *play-a-connecting-move* (between *string1* and *string2*),
- *play-a-placement-move*,
- *play-a-blocking-move* (near *here*),
- *play-an-extending-move* (near *here*).

The plans stored in the Go modules are not preprogrammed solutions and are expanded *in the context of the current game state.* The development of a plan is influenced by schema preconditions and by the choices in the world for making the lower level goals true. The failure of goals early in the plan forces choice of alternative sub plans for making these goals true.

3.3.3 Using a World Model to take Opportunities

Although the work here focuses on goal-driven approaches to Go, it is clear that human players mix both types of reasoning. Patterns are thought to play a large part in human Go play. Thus, there is a need to consider how the two approaches can be mixed successfully.

One of the advantages of using a world model in the planner is that the changing situation of the state during planning is reflected in the world model. The changing world state may highlight interesting opportunities (or obvious problems) which arise as a result of the plan actions but were not expected effects. The architecture described in §2 was extended to include plan *critics* which have access to the world model and watch for important situations. The critics are able to insert goals into the agendas of the agents for the planner to consider in parallel with the current track of reasoning. Thus, the planner is made aware of opportunities or obvious problems during planning and can react to them. [21]

[21] As noted in §2.3.2 the use of a world model also provides for forward chaining in the space of moves. The forward chaining however is only to model which moves are available for goal satisfaction. The critics here represents situations where the game state prompts *explicit* motivation for a move or new goal.

Two critics were added to GoBI which detect groups under immediate threat of capture and insert (optional [22]) group-saving/group-attacking goals into the agents' agendas. The planner may have a good plan of attack but have plan knowledge falling short of suggesting the final killing move and hence the system would normally decide that the plan had failed and try another approach. [23] With the critic switched on, the strong positions from the good attack can be checked to see if there are opportunities for killing a group and the planner told about them. The application of critics increases the size of the search space, but makes the system more robust in situations where the *a priori* plan knowledge is insufficient.

4 Testing and Evaluating GoBI

Yoshinori's four-volume series [49] provides an excellent source of problems for testing Go programs. The books provide problems designed to teach human players how to play Go. Each Go problem gives a board configuration and asks the player to choose the move or moves to achieve a particular aim (such as defending a group). The configurations of stones are usually situations which could arise during normal play but taken out of a full game context. Most problems have a unique solution which is useful for evaluation.

Since setting up each test problem was time-consuming, we chose (test set I) a representative sample (85 problems, approximately one third) of the problems from volume I of [49]. Problems were chosen for the test set to concentrate on the more complex life and death problems, excluding the sections on simple capturing problems, Atari, Ko, opening, endgame and Seki. Some problems from the Ko and Seki sections were included out of interest.

A second set of tests using problems from volume II [49] and from "Life and Death" [11] was also conducted. All the tests were limited to 30 seconds of runtime on a SUN UltraSparc.

4.1 Test Results

The system successfully solved 74% [24] of the examples in test set I, which is a significant achievement given its limited knowledge and the complexity of some of the problems. GoBI has also solved considerably harder problems from

[22] Optional in this context means precisely that satisfaction of the top level goals is not contingent on also achieving the additional optional goal.
[23] All moves have to be part of a subgoal to be suggested.
[24] 67% when critics were disabled.

volume II [49] and [11]. These further tests were not comprehensive however, and used a range of hand picked examples, so the most useful indicator of performance remains the performance on problems from volume I [49].

In 98% of the correct solutions, GoBI considered the most significant defences or opponent responses in its plan (and refuted them). This statistic is encouraging since not only were the answers correct — so was the reasoning behind them. Most of the failures were due to incomplete plan knowledge. Several problems relied on groups connecting and escaping to live, for example, for which GoBI currently has no plans. Another weak area was in problems which required a lot of forcing moves (e.g. ladders). GoBI has no special way of handling these and so is not able to take advantage of the problem simplification they provide (they are planned for in the normal way).

Strengthening GoBI's defensive knowledge led to an improvement in attacking plans and vice versa, reflecting the fact that the better opponent model is more likely to find refutations for poor attacking plans. For some of the more difficult problems in [49] and [11] reasoning was less complete (fewer of the possible counter strategies were explored). On these harder problems, as GoBI's plan knowledge ran out there was a trend of increasing reliance on the critics (data-driven aspect) for finding solutions. This trend is encouraging since it shows that data-driven and goal-driven approaches can be successfully combined in this way (but less encouraging for our knowledge engineering!).

4.2 The Cost of Move Choice

The cost of move choice can be divided into three components:

(1) *Abstract Planning Cycles* are very cheap since they consist purely of two cycles of matching items in a list (choose a goal, choose one of its plans). The process can be more expensive if complex preconditions are used to select between plans.
(2) *Checking the satisfaction of low level goals* is also inexpensive, since the goals are very focused (*are these two strings connected? could black run out here?*). [25] Checking is made easier by the fact that it can be supported by various representations of the game state — strings, groups, influence etc.
(3) *Using low level goals to generate moves* is the most expensive part of the process, although the cost is kept down by the fact that the goals by this stage are focused and limited to a small area (the *here* in the section above). In GoBI, selection is done using simple rules which define the type

[25] Again, some checks can be more costly such as checking for simple eyes (GoBI currently has a very simple eye recognition module).

of move which could satisfy a goal.[26] An example set of rules is that a *connecting-move* must:
(a) be a liberty of (next to) *string1* and
(b) be a liberty of (next to) *string2*.

Although the workload is divided between these three parts, it is clear that this move choice method is considerably more expensive for choosing a single move to try in the move tree than most data-driven methods. The potential gain of using this method is in *reducing the number of candidate moves in the move tree which need to be tried*, and *simplifying the application of evaluation functions* by restricting both the parts of the board on which they operate, and allowing them to consider only how well a position satisfies a given goal instead of trying to somehow produce an estimate of the absolute value of this position.

For the examples in test set I (see §4) there were on average 5.9 planning cycles per move chosen (including the cycles which choose moves). This average was relatively constant over the whole range of problems. To analyse the cost of using planning knowledge in GoBI we converted the planner to run without its plan knowledge and perform depth first search with $\alpha-\beta$ cutoffs. The $\alpha-\beta$ search running in the planning framework used the same board data structures, checks for move legality, backtracking mechanism and evaluation functions as GoBI. That is the search used available goals to test for termination of search and for leaf nodes. The analysis showed that given GoBI's current knowledge, each move tried in the game tree involves an average overhead of 2.5μ, where μ is the average cost of a move taken by $\alpha-\beta$ in this setup. Consequently, the cost of choosing a move in GoBI is about 3.5 times as much as a move made by the search algorithm. We note that although keeping the other features apart from the knowledge constant makes for good comparison (since modules such as the board manager could be implemented much more efficiently) general comparison with $\alpha-\beta$ search is dependent upon the efficiency of using GoBI's goals as an evaluation function.

Regarding the cost of move choice there are two important tradeoffs in the planning process:

- If goal selection and expansion mechanisms get more complex (i.e., through the extensive use of complex preconditions) their time cost will increase, but their applicability will be restricted, reducing search during planning.
- The simpler and more specific the lowest level goals are, the easier it is to establish their truth and choose actions with them but the more planning is needed to decide upon them.

[26] Move selection could just as well have been done with local pattern matching. The planning framework poses no restriction on how this relationship between abstract goals and concrete actions is established (in general this is domain dependent).

Many data-driven game playing systems are based on some form of $\alpha-\beta$ search which makes it a good scheme for comparison with GoBI. The aim of this section is to discuss the number of nodes in the search tree GoBI visits compared to the number a standard $\alpha-\beta$ search might visit. Unfortunately, the search spaces for the test problems quickly become very large so empirical testing of an $\alpha-\beta$ algorithm is inaccurate and time consuming. We therefore rely on a theoretical estimate which is derived in Appendix A.

Figure 8 compares estimates of the size of the search spaces for each of the 63 examples from test set I which GoBI answered correctly (see §4.1). The results have been grouped together into sets of five to make the graph more readable.

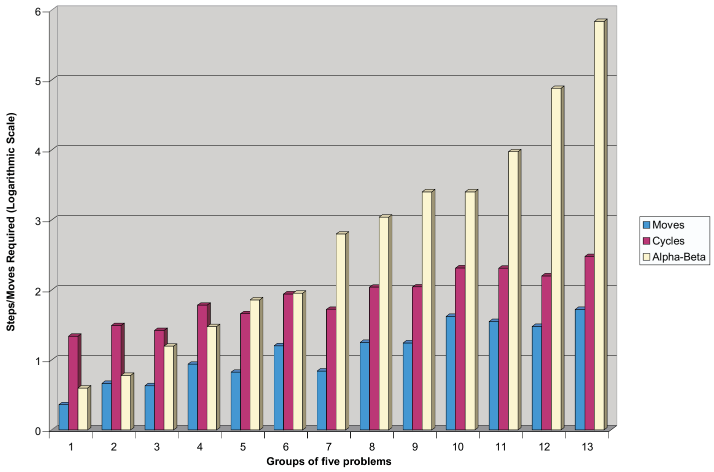

Fig. 8. The number of moves and planning cycles taken by GoBI, plus the estimated $\alpha-\beta$ search space size for the 63 correctly answered problems from test set I.

For each group the graph shows: the average number of moves taken by GoBI (moves), the average number of planning cycles required (cycles) and the size of the estimated $\alpha-\beta$ search space (alpha-beta). To group the results, the 63 cases were ordered in ascending order of estimated $\alpha-\beta$ search space (primary criterion) followed by ascending order of number of moves taken by GoBI (secondary criterion). The values were then grouped together in fives according to this order (the first five, the next five etc.). The full table of these results is available in Appendix B.

Since the graph's vertical axis is on a $\log_{10}$ scale, the graph clearly shows the cost of $\alpha-\beta$ increases rapidly as the problems become more open (larger search spaces). GoBI on the other hand is able to dramatically reduce this rate of increasing cost both in the number of nodes visited and in the number of knowledge cycles. For small search spaces, GoBI requires more planning cycles than $\alpha-\beta$ visits nodes in the search tree. This result is to be expected since even for simple problems GoBI needs to go through its reasoning process. Since the planning cycles are much cheaper than making moves in the world this extra cost is not a significant problem. More important is the gain in using the plan knowledge as the problems become larger. As a commentary to the results in the graph:

- Already for problems in Group 5 GoBI requires less planning cycles and less moves in the move tree. The problems in this group have search spaces in the order of 90 nodes[27] and GoBI solves them in 3 – 10 moves.
- Typical results from group 7 include trying 10 moves, with 46 planning cycles for estimated search space 630.
- Typical for group 10 is 28 moves tried with 125 planning cycles (the estimated $\alpha-\beta$ search space has 2520 moves).
- By groups 12 and 13, the difference is already several orders of magnitude. GoBI typically takes around 33 moves to solve problems with 10 or 11 points to play on (and hence estimated $\alpha-\beta$ search spaces of size greater than a hundred thousand and a million nodes).

The average number of moves tried by GoBI averaged over all 85 problems in the test set is 31 moves (for an average of 174 planning cycles per problem). Averaging over only the 63 correctly answered problems (those shown in Figure 8) this becomes 24 (with average 144 planning cycles per problem). Giving the $\alpha-\beta$ average is not very useful since it is dominated by the largest terms.

For the problems that GoBI failed on, 9 used an average number of moves before giving up, 8 resulted in no moves at all (or very few) and the remaining 5 timed out after 30 seconds without having found a solution and having tried between 72 and 410 moves. (A full listing of results is given in Appendix B.)

4.4 Examples from the Test Sets

To provide a better idea of the kind of problem GoBI was tested on this section gives six commented examples from the test sets. In each case the * symbols in the figure mark the correct first move (black plays first in each case).

[27] See Appendix A for explanation of these estimates.

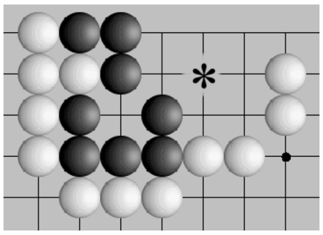

Figure 9: Example B, black to play and save the group (number 30, test set I).

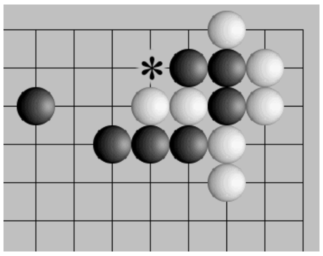

Figure 10: Example C, black to play and kill the partly enclosed white string (hand-picked example).

The problem in figure 9 was solved in 22 moves and 107 cycles (estimated $\alpha - \beta$ search space ≈ 22680). The move at * ensures that there is enough room for black to subsequently make a second eye. A determined attack from white can still make saving the group difficult and can go up to 7 moves deep.

Example C (figure 10) is harder than it appears. A standard attack on the white string would be to play one point to the left of the * (this attack is called a *net move*). For this problem when GoBI tries the net move, the opponent counter-attacks and kills the black string closest to the edge of the board. The successful counterattack forces GoBI to consider another plan. By playing at *, GoBI forces the white string to the left, trapping it against the black stone on the far left. GoBI's move also strengthens the vulnerable black string at the same time (though this is not the primary objective). The plan takes 67 moves and 244 cycles to construct (estimated $\alpha - \beta$ search space ≈ 113400).

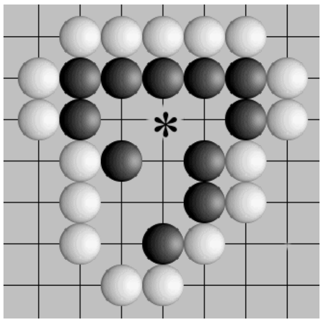

Figure 11: Example D, black to play and save the group (number 26, test set I).

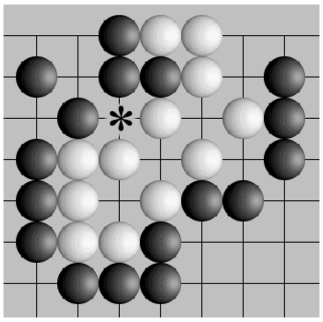

Figure 12: Example E, black to play and kill the white group (number 48, test set I).

GoBI successfully solves the problem in figure 11 using 21 moves and 146 planning cycles (estimated $\alpha - \beta$ search space ≈ 630). The move at * is correct in this case since it is the only way to ensure that black can make two eyes for its group.

Figure 12 shows a harder problem with several interacting groupings of stones. GoBi takes 60 moves and 294 planning cycles to solve this problem (estimated $\alpha-\beta$ search space $\approx$ 1 247 400 since there may well be captures during search). The move at * breaks the linkage between two parts and must be followed up by at least one more telling move to ensure that the white group eventually dies.

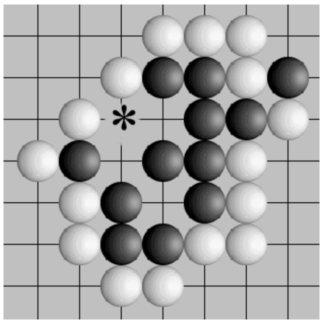

Figure 13: Example F, black to play and save the group (number 73, test set I).

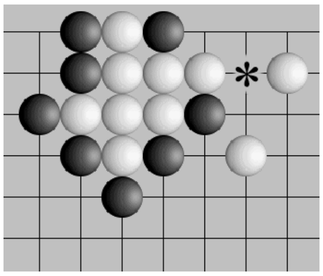

Figure 14: Example G, black to play and kill (from volume II [49]).

GoBi fails for both examples F and G. In example F (figure 13) it takes 15 moves and 68 planning cycles before giving up (the size of the $\alpha-\beta$ search space is hard to estimate since the problems involves a possible breakout through the diagonal line of white stones on the top left). This problem is difficult since it relies on a Ko threat being applied.

Failure for example G (figure 14) is due to lack of knowledge. GoBi sees all the white stones as a single group and tries a plan to encircle them but gives up without trying any moves. GoBi has no plans for breaking up groups before attacking them. This problem also raises the point that the top-level goal here needs to be redefined to allow GoBi to settle for taking just the larger part of the group. (For completeness GoBi takes 0 moves and 5 planning cycles, again search space estimation is difficult).

To give an idea of how representative these problems are compared to the rest of the test sets: Problems B and D are about average, problems C and E and F are quite hard (C in particular), problem G is easy (although GoBi still fails on it).

4.5 Open Examples

Although currently limited by its knowledge base, GoBi is more general than specialist life-and-death programs. Knowledge for other parts of the game can be freely added and GoBI is also able to handle more open positions, such as example H in figure 15 (which for example GoTools would have trouble with [48]). One would expect GoBI to compare well against existing systems in open (though still tactical) positions where the number of possible move sequences increases rapidly. The plan knowledge in the system acts to focus on only the relevant parts of the move tree. This effect can be clearly seen in figure 8 (§4.3) where the cost in moves increases much more slowly than for $\alpha-\beta$.

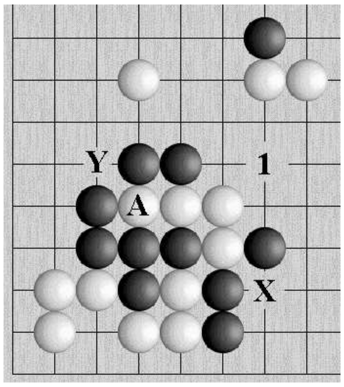

Fig. 15. Example H: GoBI solves this more open position, killing the white string marked A by playing a net move at 1. Adding white stones at X or Y for example (or various other places) would cause GoBI to realise that the string can no longer be killed directly.

5 Evaluation of the Goal-Driven Approach to Go

As with any method for applying knowledge to searching large spaces, the success of the goal-driven approach depends upon features of the domain of application. This section examines the advantages and disadvantages of applying this approach to Go play.

5.1 Advantages of the Goal-Driven Approach

The goal-driven approach which is presented here has some clear advantages for Go and other similar games. Together with some of the previous work on the usefulness of planning at the strategic level of Go, GoBI shows that this approach can be used for reasoning at all levels. The hope is that it will eventually be possible to move transparently between tactical and strategic levels of play and providing a unifying framework for Go.

5.1.1 *Representation and Communication of Domain Knowledge*

Go knowledge in books and folklore is often expressed in a form appropriate for encoding as decompositions of abstract goals into other goals at different levels of abstraction. As reported in [35], there is a rich vocabulary which Go players can use to express their reasons for making a move. There are many popular proverbs, which convey strategies at various levels of abstraction, for example "death lies in the hane" is more tactical, whereas "don't push along the fifth line" is quite strategic. It may be easier to add this kind of knowledge to a goal-driven system than to a data-driven system which requires the delicate balancing of heuristics.

Adding knowledge in the form of plans is made very much easier by the existence of natural levels of abstraction when describing the game state. Game positions are often considered not only in terms of stones but in terms of macro structures such as strings groups and areas of influence.

By following the trace of the goal decompositions one can see *why* the Go player is trying to do something — its aims and plans. Access to aims and plans is not only helpful for adding knowledge and debugging, but could be useful in the context of a Go tutoring aid. Some of the current commercial Go systems (Many Faces of Go for example) have teaching mechanisms but it is not clear whether these mechanisms are based directly on the reasoning process of the computer player, or are instead based on a high-level post-hoc analysis (which therefore will not accurately reflect the reasoning which led to this choice of move).

5.1.2 *Search Properties*

Using a goal-driven approach leads to a very different search from a data-driven approach. Some key advantages are:

- There is less need for global evaluation functions: Full board evaluation reduces to checking if lower level goals can be satisfied. Evaluation functions may still be used in a limited way to carry out judgements which cannot be made on a move-by-move level, for example using an influence function to judge whether or not a group can successfully run, or using a fast $\alpha-\beta$ search to determine the life or death of enclosed groups. When an evaluation function is employed in determining whether or not a goal has been satisfied, the goal can be used to provide a focus for the evaluation function to a limited area of the board. The use of goals to focus evaluation makes sense when thinking about how humans play go — often making instinctive judgements (evaluations) of a local situation but rarely considering (evaluating) the whole board.
- Quiescence (a problem for many game-playing programs) is defined auto-

matically, thus avoiding problems such as the horizon effect in search. Since each agent has an agenda of goals, a situation can be defined as "open" (unsettled) until all goals are achieved.
- What seem like heuristically bad moves (e.g. sacrifices) are not discriminated against because the only consideration is their value to the plan.

While these advantages are important, perhaps the overriding advantage often associated with goal-driven approaches is their resilience in domains with large search spaces and branching factors. As the figures in §4.3 show, the planner dramatically reduces the number of nodes searched in the move tree. The advantage is still significant even if the extra cost of planning is taken into account.

5.2 Disadvantages of Goal-Driven Search

Obviously the goal-driven approach is not always the best choice and has its own difficulties:

- The goal-driven approach requires significant effort to encode strategies as goal decompositions. By contrast, in the data-driven approach, good play can be achieved using even relatively simple evaluation functions if the game tree can be searched deeply enough.
- There are some types of knowledge which are hard to express in a goal/plan oriented framework, such as knowledge which is not reliant on understanding the motivation behind a move (patterns for example). It seems clear that good Go play requires both pattern (data-driven) and abstract plan (goal-driven) knowledge which is what leads us to try and integrate the two approaches (see §3.3.3).
- For games with low branching factors, shallow search trees or where near exhaustive search is possible data-driven approaches have a strong advantage. It is only when searching most of the move tree is infeasible and large amounts of knowledge are needed to prune the tree that goal-driven approaches become more useful. [28]

The main problem with GoBi itself is still lack of knowledge: much more is needed before GoBi could be used as a tool to play the complete game of Go. Unfortunately adding knowledge takes a long time since it must be hand coded; most of the top programs have had knowledge added over periods of years. In this respect the goal-driven approach (and many current approaches)

[28] This point was illustrated by PARADISE [43] in the early eighties which despite being able to solve Chess problems requiring search up to 20 ply deep (far beyond other Chess programs of the time), still saw its knowledge based approach outstripped by the ever increasing efficiency of fast search techniques.

are at a disadvantage to learning or non-symbolic approaches which can use automatic training to improve.

5.3 Scaling Issues

As with all knowledge based approaches, the search space is determined by the knowledge in the system. The key to making knowledge based approaches work is good organisation of that knowledge. The size of the search space is determined by the number of goal schemas, and by how much they overlap. A small number of goal schemas, which are mutually exclusive, lead to a relatively small search space. Conversely, a very large number of goal schemas which overlap in many ways lead to a large search space. In practice, the size of the search space lies between these two extremes.

Schemas which do not overlap, for example schemas which address different parts of the game, can be added without adversely affecting performance. Adding more detailed plans needs to be done carefully, since it can lead to overlaps and redundant search. The levels of abstraction used in GoBI (and in many AI planners) are key in avoiding redundancy since knowledge can be appropriately structured and similar plans can be grouped together and expressed in a compact form. Badly structured or overlapping knowledge can cause similar problems to occur in data-driven approaches where pattern matching is used to suggest moves.

In Wilkins' Paradise [43] work, as in our work, there was a worry that extending the system by adding new knowledge would have adverse affects on system performance. New production rules were added to Paradise in order to allow it to solve problems outside the initial development set [43, p.193]. The new problems were correctly solved, and system performance was not degraded. In fact, the extended system now found *better* plans for some of its previously solved positions. This result is heartening for our own work.

5.4 Comparisons With Other Approaches to Adversarial Planning

One of the advantages of building a generic planning framework (§2.2) and then separately instantiating this to play Go (§3) is that the characteristics of the planner are clearly laid out. This separation improves understanding of the planner's behaviour and facilitates comparison with systems from other domains.

The most useful comparisons can be made using a classification based up on the expansion scheme used during planning, since this is the defining feature of any HTN planning system. The planning systems can be classified as follows:

- **Total order decomposition:** PARADISE, TIGNUM2, GOBI, the battlefield management system in [1].
- **Partial order decomposition:** Work in [21], [23], [37] and the INTERIM.2, and GOBELIN programs.

In this classification and leaving aside the application domain, GOBI is most similar to PARADISE and TIGNUM2. In fact these planners not only use a form of total order expansion but also incorporate aspects of forward chaining (§2.3.2).

5.4.1 PARADISE

The idea of using goals to guide search originally comes from [29] and was then incorporated into PARADISE. Although PARADISE is based on rule-based expert systems and in this sense is quite far removed from HTN planning, its search approach leads to a type of goal guided search at two levels of abstraction. Each goal is expanded by a rule based knowledge source in the context of its final place in the plan. This guided search model is conceptually very similar to the modified HTN expansion scheme described in §2.3.2. There are many differences between PARADISE and the planner used in GOBI, but perhaps the most significant at a conceptual level are that:

- PARADISE only used two levels of abstraction — the ground space (moves) and one abstract level (goals or concepts). With its current knowledge GOBI uses five levels of abstraction.
- PARADISE's starting plans are formed from a static evaluation of the initial game state and global static evaluations may be applied at various points throughout the search. In GOBI, goals can have associated information about what information to apply from the environment for expansion (in the form of preconditions). This precondition mechanism gives a full range of flexibility (from very focused detailed evaluations of a local area for low level goals to very sketchy high level evaluations of the whole board for high level goals).
- PARADISE can analyse failed plans and use the extracted information in constructing a new plan.
- PARADISE uses different models for attacking and defending players. While the attacker uses the plan suggested by static analysis to focus on only a few candidate move choices, the defender tries all "reasonable" moves. In GOBI, attack and defence use the same plan schemas. Indeed, both players may be attacking simultaneously.

Above all, the planning architecture used by GOBI provides a simple generic framework which is simpler to understand than the mixture of different elements which come together to make up PARADISE.

5.4.2 TIGNUM2

TIGNUM2 [40,39] also has a decomposition model very similar to that used in GOBI (and is explicitly based upon HTN planning). The propositional logic formulas used by TIGNUM2 to represent the state of play directly correspond to the use of a world model in GOBI. The most significant difference between GOBI's planner and the one used in TIGNUM2 is in the expressiveness of the planning language. TIGNUM2's planner is kept total order by restricting *plan knowledge* to being totally ordered, preventing subgoals descended from different higher level goals from being interleaved. [29] In TIGNUM2, given the decomposition tree shown in Figure 16, goals descended from G2 could not be interleaved with goals descended from goal G1. Furthermore the goals descended from G1 must be fulfilled in the order given in the schema — G11, then G12, then G13.

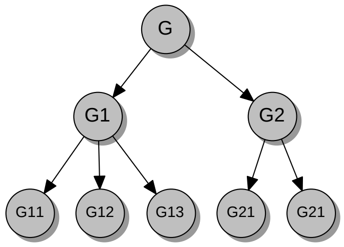

Fig. 16. Goal decomposition tree: the top level goal G decomposes into two subgoals G1 and G2, each of which is again decomposed.

There is no such restriction in the planning framework presented in this paper: subgoals can be interleaved and all linearisations of subgoals are valid [30]. This flexibility preserves the expressiveness of standard HTN planning (which is semi-decidable, see [14]) while TIGNUM2's planner is strictly less expressive (EXPSPACE-hard and in DOUBLE-EXPTIME, [27]). The extra flexibility in GOBI's planner is crucial for games such as Go where interleaving goals is a natural part of play. In Go, capturing a group for example may involve several tasks (surrounding it, preventing eyes etc.) which need to be achieved in pseudo-parallel. That is, concentrating on thoroughly surrounding a group will result in the opponent making eyes and saving the group (and vice versa). Steps need to be taken towards achieving both subgoals while stopping the opponent countering either (having enough spare moves to do so).

[29] This definition of total order planning is slightly different to the one used in [25] and that adopted in this paper.
[30] Constraints can be placed on orderings, but these constraints form part of the explicit plan knowledge.

5.4.3 Using Total Order Expansion in Adversarial Domains

For many practical problems the total order decomposition systems identified in §5.4 above have significant advantages over those using partial order decompositions. This observation runs contrary to conventional wisdom in the AI planning community. The least commitment aspect of partial order decompositions was always thought to bring significant advantages. Recent work in [27] discusses the assumptions underlying the use of partial order, backwards chaining strategies for HTN planners and suggests that they are not always valid (and in particular do not apply for Bridge play and process planning). In the context of an adversarial planning domain we can elaborate on these reasons (and include observations noted in [46]).

As is noted in [14] "handling interactions among non-primitive tasks [goals] is the most difficult part of HTN planning". In adversarial domains these difficulties are compounded. More precisely:

- *The interactions are significant:* Goal expansion is usually heavily dependent on the game state. The simplifying assumption often taken in AI planning research that expansion choices (*how* to do something) and ordering choices (*when* to do it) are independent [2] is less valid in adversarial domains.
- *More of the interactions need to be taken into account:* In an adversarial domain, the adversaries have control over some of the actions, therefore not only must there be at least one valid ordering, the opponents must not be able to order their actions in such a way as to make the plan fail. The planner must explore many more of the options and orderings possible to ensure they all lead to a desirable outcome.
- *The interactions are difficult to model:* The interactions between moves in a game are generally *designed* to be complex. So complex, in fact, as to fool and entertain human players. Effects are often cumulative, subtle and only obvious over long sequences of moves. This complexity makes interactions inherently hard to model and even harder to model in the abstract where much of the detail is stripped away. In his 1979 paper [44] Wilkins observes that defining abstract spaces for Chess is very difficult because "small details are very important".

Games tend to be built around the effects of complex interactions between the moves of the players. It becomes increasingly difficult to model all these interactions adequately at higher levels of abstraction. If important details (interactions) are not somehow modelled in the abstract space and taken into account during planning, plans will repeatedly fail at the lowest level and cause repeated backtracking.

The use of a total order scheme which is able to model interactions between moves and goals directly in a world model is very important in avoiding some of these problems. It also means that an aspect of forward chaining can be introduced at the move level to reduce the number of options available (since the starting state is always known even though the goal state may not be perfectly described). These differences appear to be critical in enabling total order planners to outperform standard HTN planners in tight tactical situations.

6 Conclusions

In this paper we have presented an adversarial planning architecture capable of reasoning about games, and an application of this architecture to Go. The planning architecture and Go reasoner reported here represent an advance on previous work for goal-driven planning in Go. The system:

- Has a clear separation of domain knowledge from the abstract planning architecture and a clear model of the opponent in the game.
- Can reason at multiple levels of abstraction simultaneously.
- Can address complex tactical situations as well as high level strategic problems.
- Can provide support for the integration of data-driven and goal-driven approaches.

We presented the advantages that a goal-driven approach could have for Go. GoBI as a prototype is certainly no match for current systems which play the full game of Go, but it does represent a step towards understanding how goal-driven approaches can be applied to Go, even at a tactical level.

We believe that GoBI demonstrates that an approach based on adversarial planning can be effective for Go. The two main limitations of GoBI as a Go-playing program are:

- The lack of top-level goals. We have indicated how GoBI could be extended with high-level persistent goals which would enable it reason about the entire game.
- The small knowledge base. The knowledge base needs to be very significantly extended in order to make a realistic Go program.

Go has several strong domain features which make goal-driven approaches applicable: very large search spaces, clear layers of abstraction in domain descriptions (stones, strings, groups etc), and a wealth of knowledge similar in structure to abstract plans. GoBI represents a further step towards turning this theoretical possibility into a reality. The work described in this paper again shows that Go is an excellent test bed for Artificial Intelligence research. There has been very little work on adversarial planning in recent years — the challenge of Go really motivated this work.

We have outlined why GoBI outperforms previous Go planners at tactical play. The total order decomposition scheme and use of a world model were key in making progress. Taken together with previous work on Chess and Bridge this suggests that the techniques we have described in this paper should be strongly considered when applying planning techniques in future game playing systems.

Acknowledgements

The first author is very grateful for financial support from the EPSRC under grant number 96417312. The second and third authors are supported by EPSRC grants GR/L/11724 and GR/M/45030. The fourth author is partly funded by DARPA/AFRL contract F30602-95-1-0022. Special thanks also go to Martin Müller for suggestions regarding test data, and to Neil Firth for numerous instructive games of Go. We are grateful to the anonymous referees whose comments have resulted in numerous improvements to previous versions of this paper.

A Theoretical Estimate of $\alpha-\beta$ Search Space

As stated in §4.3 estimating the size of $\alpha-\beta$ search spaces for given problems empirically is difficult since search spaces and search times rapidly increase. The performance of $\alpha-\beta$ is also heavily dependent on the evaluation function used. This difficulty forces us to rely on theoretical estimates. A number of papers published by UCLA in the late 70s and early 80s give theoretical analyses of the $\alpha-\beta$ algorithm's performance. Many of these results are summarised in [28]. However, the results generally rely on a constant branching factor (which in tight tactical Go problems is not valid) and sampling leaf node values from a continuous distribution. For a theoretical analysis closer to the type of problem GoBI solves we need to go a little outside these results.

The following assumptions allow a simple treatment:

- Assumption 1: The number of available moves is reduced by one each round and both players play moves from the same finite set of moves. These assumptions do not always hold for Go since captures can increase the number of moves available, some moves may only be legal for one player and players may choose to pass, however they are acceptable for a simple treatment.

- Assumption 2: There is only one solution (i.e., one correct move to play first which leads to a win being possible in all of the subtrees following this move). This assumption also does not hold for general Go play but is valid in the type of problems in the test sets.

Considering a search with n options at the top level, assumption 1 gives the total number of leaf nodes as $n!$ where leaf nodes are defined as the nodes at depth n (or equally $n-1$ in the search). We now estimate the number of leaf nodes visited, on average, by an $\alpha-\beta$ search algorithm which is able to evaluate leaf states as *win* or *lose*, does not evaluate intermediate (non-leaf) nodes, and stops once it finds the correct first move (as GoBI does).

Over all orderings of moves and solutions for a given problem one would expect the algorithm on average to choose the correct move having tried half the possible moves (contingent on assumption 2). Following a move selection, *all* the opponents responses need to be tried to ensure the move leads to a win everywhere, thus making the number of leaf nodes visited equal to:

$$\frac{n}{2} \times (n-1) \times \frac{n-2}{2} \times (n-3) \times ... \tag{A.1}$$

$$\text{which gives:} (\frac{1}{2})^{\frac{n}{2}} n! \tag{A.2}$$

(The odd n are those multiplied by the factor half.) It is interesting to note that the pruning power $(2^{\frac{n}{2}})$ of $\alpha-\beta$ is not simply a linear factor of the search space size (See the table in Figure A.1). For an accurate comparison in the number of *moves* tried by the search, we also need to count the number of moves tried in reaching the leaf nodes. The number of moves required becomes:

$$(\frac{1}{2})^{\frac{n}{2}} n! + \frac{(\frac{1}{2})^{\frac{n-1}{2}} n!}{(n-(n-2))!} + \frac{(\frac{1}{2})^{\frac{n-2}{2}} n!}{(n-(n-3))!} + ... \tag{A.3}$$

Note again that this figure is valid in the limit, but generally the power of the pruning factor is rounded down to the nearest integer, reflecting the fact it alternates on the levels with odd n. In the limit the new value in Equation A.3 is around double the value given in Equation A.2. Thus the number of nodes visited is dominated over n by the number of leaf nodes visited, $(\frac{1}{2})^{\frac{n}{2}} n!$.

Analysis of the number of moves tried becomes considerably more complex if the algorithm can evaluate intermediate (non-leaf) nodes since the performance of $\alpha-\beta$ algorithms is generally sensitive to the evaluation function used. We have not taken this possibility into account in our analysis.

The result obtained for the pruning power of $\alpha-\beta$ on this type of problem (a binary set of state values, one correct move, stopping on finding this move

Options	Raw	Pruning Factor	E[Leaf nodes]
3	6	2	3
4	24	4	6
5	120	4	30
6	720	8	90
7	5040	8	630
8	40320	16	2520
9	362 880	16	22680
10	3 628 800	32	113400
11	39 916 800	32	1 247 400
12	479 001 600	64	7 484 400
13	6 227 020 8000	64	97 297 200

Fig. A.1. Calculating the expected number of leaf nodes searched. The column headed *Options* is the number of moves considered in the initial position, *Raw* is the size of the unpruned search space, *Pruning Factor* is the factor by which $\alpha-\beta$ search reduces the number of leaf nodes visited, and *E[Leaf nodes]* is the expected number of leaf nodes visited.

and having a steadily decreasing branching factor), seems to concur with the UCLA results [28] which also predict low order exponential pruning power for their problems.

A.1 Estimating the search space for Go Problems

To estimate the search space that would be required by $\alpha-\beta$ search for the Go problems in the test set each problem has a search area mask defined for it. The mask defines which points on the board need to be tried as moves by the search in order to find a satisfactory solution (meaning one which chooses the correct move but also encompasses checking all significant defences). Figure A.2 shows a mask defined for one of the problems from test set 1. The * characters mark the points which define the limited search space.

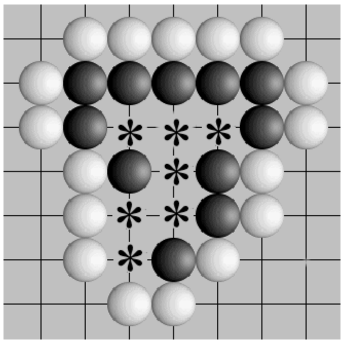

Fig. A.2. A mask for a Go problem from test set I. The * symbols indicate the area that needs to be considered for search.

For our estimates all the masks were defined by hand. Once the mask has been defined for a problem, the number of points defined in it gives the number of moves available in the search space (by Assumption A1 in the previous section). Some masks also included occupied points where stones were captured as part of attack or defence. Together with the formula in Equation A.2 (or Figure A.1), this gives the estimate of the expected search space for that problem. We use Equation A.2 rather than Equation A.3 since the number of leaf nodes is the dominating term.

This estimation of search space size for a Go problem is clearly very conservative since the mask and the search depth are in fact complex to compute automatically. Arguably finding this information forms part of the search problem itself. The advantage of using the hand derived mask and search depth however, is that they give a clearly defined and understandable way of generating estimates. Without these restrictions, search space estimates which could be made almost arbitrarily large and be of little use for comparison.

B Test Results

For completeness we include a listing of the results for the 85 problems tested in test set I (results are from tests using critics). In the following table, the column headings are: the problem number (*N*), the problem number given in [49] (*Ref*), the number of planning cycles GoBI takes to plan the problem (*P*), the number of moves GoBI tried (*M*), whether GoBI succeeded in finding a correct problem solution (*Res*, C = correct answer, F = fail), and the estimated number of moves in the $\alpha-\beta$ search tree for this problem (*E*). For details of how these examples were selected from [49] see §4.1.

N	Ref	P	M	Res	E	N	Ref	P	M	Res	E
1	41	42	8	C	6	44	145	702	126	C	2520
2	42	20	3	C	30	45	147	49	7	C	30
3	43	17	2	C	6	46	148	1002	410	F	30
4	44	49	7	C	6	47	149	285	44	C	113400
5	45	30	4	C	6	48	150	294	60	C	1 247 400
6	99	184	35	C	2520	49	151	20	3	C	30
7	100	74	13	C	90	50	152	22	3	F	30
8	101	75	19	C	90	51	153	47	7	F	90
9	102	70	14	C	30	52	154	9	1	F	90
10	103	176	39	C	113400	53	155	20	3	C	90
11	104	25	4	C	2520	54	156	61	12	C	2520
12	105	136	22	F	90	55	157	120	20	C	1 247 400
13	106	41	4	C	3	56	158	33	5	C	2520
14	107	38	10	C	6	57	159	62	22	C	2520
15	109	132	18	C	2520	58	160	1002	155	F	2520

N	Ref	P	M	Res	E	N	Ref	P	M	Res	E
16	110a	46	10	C	630	59	161	32	7	F	1 247 400
17	110b	1002	164	F	2520	60	162	1002	151	F	113400
18	111	724	138	C	2520	61	163	122	18	C	2520
19	146	210	35	F	630	62	164	23	2	F	3
20	121	17	2	C	3	63	165	434	54	C	2520
21	122	42	7	C	30	64	166	609	110	C	2520
22	123	106	12	C	30	65	167	125	28	C	2520
23	124	29	4	C	90	66	168b	21	2	C	6
24	125	116	20	C	2520	67	169b	31	3	C	6
25	126	54	6	C	30	68	170a	20	2	C	3
26	127	146	21	C	630	69	172	17	6	F	3
27	128	93	16	C	90	70	173	11	1	F	6
28	129	279	59	C	630	71	108	4	1	F	3
29	130	77	13	C	90	72	176	23	5	F	1 247 400
30	131	107	22	C	22680	73	177	68	15	F	113400
31	132	20	2	C	30	74	178	90	19	F	630
32	133	129	20	C	90	75	222	12	2	F	90
33	134	55	8	C	90	76	223	118	23	C	22680
34	135	36	5	C	630	77	224	20	2	F	6
35	136	90	13	C	630	78	225	59	12	C	22680
36	137	18	5	C	6	79	226	767	118	C	630
37	138	54	6	C	630	80	227	518	53	C	22680
38	139	65	18	C	113400	81	228	103	17	F	90
39	140	93	9	C	630	82	230	56	10	C	90
40	141	50	6	C	630	83	231	12	1	F	30
41	142	113	14	C	30	84	232	35	5	C	6
42	143	87	10	C	22680	85	233	1002	72	F	2520
43	144	814	147	C	1 247 400						

References

[1] C. Applegate, C. Elsaesser, and D. Sanborn. An Architecture for Adversarial Planning. *IEEE Transactions on Systems, Man and Cybernetics*, 20(1):186–294, 1990.

[2] A. Barrett and D. S. Weld. Partial-Order Planning; Evaluating Possible Efficiency Gains. *Artificial Intelligence*, 67(1):71–112, 1994.

[3] H. J. Berliner. A Chronology of Computer Chess and its Literature. *Artificial Intelligence*, 10:201–214, 1978.

[4] R. Bozulich. *Second Book of Go*. The Ishi Press Inc, 1987.

[5] D. J. H. Brown and S. Dowsey. The Challenge of Go. *New Scientist*, 81:303–305, 1979.

[6] J. Burmeister and J. Wiles. An Introduction to the Computer Go Field and Associated
Internet Resources. Technical report, The University of Queensland, January 1997. Available online at: http://www/psy.uq.edu.au/˜jay/go/go_page.html.

[7] J. G. Carbonell. Counterplanning: A Strategy Based Model of Adversarial Planning in Real World Situations. *Artificial Intelligence*, 16(1):295–329, 1981.

[8] T. Cazenave. *Système d'Apprentisage par Auto-Observation. Application au Jeu de Go*. PhD thesis, L'Université Paris 6, 1996.

[9] T. Cazenave. Metaprogramming Forced Moves. In Henri Prade, editor, *Proceedings of the 13th European Conference on Artificial Intelligence (ECAI'98)*, pages 645–649. John Wiley and Sons, 1998.

[10] S. F. Da Silva. Go and Genetic Programming, Playing Go with Filter Functions. Master's thesis, Leiden University, Netherlands., 1996. Available online at: http://www.wi.leidenuniv.nl/MScThesis/dasilva.html.

[11] J. Davies. *Life and Death*. The Ishi Press Inc, 1978.

[12] P. Donnelly, P. Corr, and D. Crookes. Evolving Go Playing Strategy in Neural Networks. *AISB Workshop in Evolutionary Computing*, 1994.

[13] M. Enzberger. The Integration of a Priori Knowledge into a Go Playing Neural Network. Technical report, University Munich., 1996.

[14] K. Erol, D. Nau, and J. Hendler. HTN Planning: Complexity and Expressivity. *Proceedings of AAAI'94*, July 1994.

[15] K. Erol, D. Nau, and J. Hendler. UMCP: A Sound and Complete Planning Procedure for Hierarchical Task-Network Planning. *Proceedings of AIPS-94*, June 1994.

[16] D. Fotland. Knowledge Representation in The Many Faces of Go. Technical report, American Go Association, 1993. Available online at: ftp://bsdserver.ucsf.edu/Go/comp/mfg.Z.

[17] D. Fotland and A. Yoshikawa. The 3rd FOST Cup World-Open Computer-Go Championship. *ICCA Journal*, 20(4):276–278, 1997.

[18] I. Frank. *Search and Planning under Incomplete Information: a study using Bridge card play*. PhD thesis, Department of Artificial Intelligence, University of Edinburgh, 1996.

[19] I. Frank, D. Basin, and A. Bundy. An Adaptation of Proof-Planning to Declarer Play in Bridge. In *Proceedings of the 10th European Conference on Artificial Intelligence (ECAI'92)*, pages 72–76, Vienna, Austria, 1992. Longer Version available from Edinburgh as DAI Research Paper No. 575.

[20] M. L. Ginsberg. GIB: Steps Toward an Expert-Level Bridge-Playing Program. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI'97)*, 1999.

[21] S. Hu. *Multipurpose Adversary Planning in The Game of Go*. PhD thesis, George Mason University, 1995.

[22] T. Kojima, K. Ueda, and S. Nagano. An Evolutionary Algorithm Extended by Ecological Analogy and its Application to the Game of Go. *Proceedings of the 15th International Joint Conference on Artificial Intelligence (IJCAI'97)*, pages 684–689, 1997.

[23] P. Lehner. Strategic Planning in Go. In M. A. Bramer, editor, *Computer Game Playing: Theory and Practice*, pages 167–176. Ellis Horwood, 1983.

[24] D. Lichtenstein and M. Sipser. Go is Polynomial-Space Hard. *Journal of ACM*, 27(2):393–401, April 1980.

[25] S. Minton, J. Bresina, and M. Drummond. Total-Order and Partial-Order Planning: A Comparative Analysis. *Journal of Artificial Intelligence Research*, 2:227–262, 1994.

[26] M. Müller. *Computer Go as a Sum of Local Games: An Application of Combinatorial Game Theory*. PhD thesis, Swiss Federal Institute of Technology, Zurich, 1995.

[27] D. S. Nau, S. J. J. Smith, and K. Erol. Control Strategies in HTN Planning: Theory Versus Practice. In *Proceedings of the Innovative Applications of Artificial Intelligence Conference (in conjunction with AAAI'98)*, pages 1127–1133. AAAI Press, 1998.

[28] J. Pearl. The Solution for the Branching Factor of the $\alpha-\beta$ Pruning Algorithm and its Optimality. Research Paper UCLA-ENG-CSL-8019, University of California, Los Angeles (UCLA), June 1981.

[29] J. Pitrat. A Chess Combination Program Which Uses Plans. *Artificial Intelligence*, 8(1):275–321, 1977.

[30] W. Reitman and B. Wilcox. The Structure and Performance of the INTRIM.2 Go Program. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI'79)*, pages 711–719. 1979.

[31] P. Ricaud. A Model of Strategy for the Game of Go Using Abstraction Mechanisms. *Proceedings of the 15th International Joint Conference on Artificial Intelligence (IJCAI'97)*, pages 678–683, 1997.

[32] N. Richards, D. Moriarty, and R. Miikkulainen. Evolving Neural Networks to Play Go. Technical report, The University of Texas at Austin., 1997.

[33] J. M. Robson. The Complexity of GO. Technical Report TR-CS-82-14, Department of Computer Science, The Australian National University, Canberra 0200 ACT, Australia, October 1982. Also published IFIP: International Federation of Information Processing 1983.

[34] E. D. Sacerdoti. *A Structure for Plans and Behaviour*. Elsevier, 1977.

[35] Y. Saito and A. Yoshikawa. Do Go Players Think in Words? — Interim Report of the Analysis of Go Player's Protocols. In Hitoshi Matsubara, editor, *Proceedings of the Second Game Programming Workshop*. Computer Shogi Association, 1995.

[36] Y. Saito and A. Yoshikawa. An Analysis of Strong Go-Players' Protocols. In Hitoshi Matsubara, editor, *Proceedings of the Third Game Programming Workshop*. Computer Shogi Association, 1996.

[37] P. T. Sander and D. J. M. Davies. A Strategic Approach to the Game of Go. In M. A. Bramer, editor, *Computer Game Playing: Theory and Practice*, pages 152–166. Ellis Horwood, 1983.

[38] J.J. Smith and D.S. Nau. Strategic Planning for Imperfect Information Games. In *Games: Planning and Learning, Papers from the 1993 Fall Symposium*, pages 84–91, AAAI Press, 1993.

[39] S. J. J. Smith, D. S. Nau, and T. A. Throop. A Planning Approach to Declarer Play in Contract Bridge. *Computational Intelligence*, 12(1), 1996.

[40] S. J. J. Smith, D. S. Nau, and T. A. Throop. Total-Order Multi-Agent Task-Network Planning for Contract Bridge. *Proceedings of AAAI'96*, pages 108–113, 1996.

[41] J. A. Storer. On the Complexity of Chess. *Journal of Computer and System Sciences*, 27(1):77–100, August 1983.

[42] A. Tate. Generating Project Networks. In *Proceedings of the 5th International Joint Conference on Artificial Intelligence (IJCAI'77)*. 1977.

[43] D. E. Wilkins. Using Patterns and Plans in Chess. *Artificial Intelligence*, 14(1):165–203, 1980.

[44] David E. Wilkins. Using Plans in Chess. In *Proceedings of the 6th International Joint Conference on Artificial Intelligence (IJCAI'79)*, pages 960–967, Tokyo, Japan, 1979.

[45] S. N. Willmott. Adversarial Planning and the Game of Go. Master's thesis, Department of Artificial Intelligence, University of Edinburgh, September 1997.

[46] S. N. Willmott, A. Bundy, J. M. Levine, and J. D. C. Richardson. Adversarial Planning in Complex Domains . Technical report, Department of Artificial Intelligence, University of Edinburgh, January 1998. Research Paper Number 889.

[47] T. Wolf. The Program GoTools and its Computer-Generated Tsumego Database. In *Proceedings of the First Game Programming Workshop*, pages 84–96. Computer Shogi Association, 1994.

[48] T. Wolf. About Problems in Generalizing a Tsumego Program to Open Positions. In Hitoshi Matsubara, editor, *Proceedings of the Third Game Programming Workshop*. Computer Shogi Association, 1996.

[49] K. Yoshinori. *Graded Go Problems for beginners (Volumes I - IV)*. The Ishi Press Inc, 1985.

[50] P. R. Young and P. Lehner. Applications of a Theory of Automated Adversarial Planning to Command and Control. *IEEE Transactions on Systems, Man and Cybernetics*, 16(6):186–294, 1990.