Summary

This thesis demonstrates by example how dynamic stochastic control methods may usefully be applied to tackle problems of tree searching and computer game-playing. Underlying this work is the constant tension between what is optimal *in theory* and what is implementable *in practice*. Most of the games studied are solved (under certain conditions) in the sense that a policy is derived which is both optimal and implementable. We examine the various reasons why the general problem of devising an algorithm to play a perfect information game in real time cannot have such a solution, and consider how to respond to this difficulty.

Chapter 1 defines the nature of the problem by introducing the concept of a game tree and explaining the concept of selectivity in game tree search. It then reviews the most important game tree search methods.

Chapter 2 explains what a Markov chain model of computer game-playing might include. It then introduces a much simpler one-player search game and establishes optimal policy under certain conditions. It also contains analysis of a discrete fuel control problem, which was developed as an offshoot of this research.

Chapter 3 details a stochastic satisficing tree search model, and explains the relationship with the standard bandit models. The optimal policy is established, and some important extensions presented.

Chapter 4 considers the difficulties of extending this model to a two-player game tree.

Chapter 5 highlights the connection between the problem of time- and search-control. The chapter then provides an overview of the approaches taken to the problem of time control, before proving an optimal time- and search-control policy for a simple two-player game.

Chapter 6 is the most wide ranging in its scope and is approachable to the artificial intelligence researcher less familiar with the language of dynamic stochastic control. It presents a new selective search algorithm, PCN*, and highlights problems with some of the existing game-playing models.