Application of Tabu Search to Deterministic and Stochastic Optimization Problems

Özgür Gürtuna

A Thesis

in

The John Molson School of Business

Presented in Partial Fulfillment of the Requirements

for the Degree of Doctor of Philosophy at

Concordia University

Montreal, Quebec, Canada

May 2006

Your file Votre référence
ISBN: 978-0-494-16302-3
Our file Notre référence
ISBN: 978-0-494-16302-3

Abstract

Application of Tabu Search to Deterministic and Stochastic Optimization Problems

Özgür Gürtuna

Concordia University, 2006

During the past two decades, advances in computer science and operations research have resulted in many new optimization methods for tackling complex decision-making problems. One such method, tabu search, forms the basis of this thesis. Tabu search is a very versatile optimization heuristic that can be used for solving many different types of optimization problems.

Another research area, real options, has also gained considerable momentum during the last two decades. Real options analysis is emerging as a robust and powerful method for tackling decision-making problems under uncertainty.

Although the theoretical foundations of real options are well-established and significant progress has been made in the theory side, applications are lagging behind.

A strong emphasis on practical applications and a multidisciplinary approach form the basic rationale of this thesis. The fundamental concepts and ideas behind tabu search and real options are investigated in order to provide a concise overview of the theory supporting both of these two fields. This theoretical overview feeds into the design and development of algorithms that are used to solve three different problems.

The first problem examined is a deterministic one: finding the optimal servicing tours that minimize energy and/or duration of missions for servicing satellites around Earth's orbit. Due to the nature of the space environment, this problem is modeled as a time-dependent, moving-target optimization problem. Two solution methods are developed: an exhaustive method for smaller problem instances, and a method based on tabu search for larger ones.

The second and third problems are related to decision-making under uncertainty. In the second problem, tabu search and real options are investigated together within the context of a stochastic optimization problem: option valuation. By merging tabu search and monte carlo simulation, a new method for studying options, Tabu Search Monte Carlo (TSMC) method, is developed. The theoretical underpinnings of the TSMC method and the flow of the algorithm are explained. Its performance is compared to other existing methods for financial option valuation.

In the third, and final, problem, TSMC method is used to determine the conditions of feasibility for hybrid electric vehicles and fuel cell vehicles. There are many uncertainties related to the technologies and markets associated with new generation passenger vehicles. These uncertainties are analyzed in order to determine the conditions in which new generation vehicles can compete with established technologies.

anneme ve babama

to my parents

Acknowledgments

This thesis is the culmination of a wonderful intellectual journey. Although a Ph.D. is, by its very nature, a solitary undertaking, my journey wouldn't have been possible without the generous support and understanding of the following individuals.

I would like to thank my wife, Simone, for her continuous support and encouragement throughout my Ph.D. Her brilliant mind, her loving care and her infinite patience were the foundation of my work.

I would also like to thank my Ph.D. committee and in my particular my supervisors, Professors Jean-Marie Bourjolly and Pierre Lasserre, who enabled me to realize this dream. Professor Jean-Marie Bourjolly has been an excellent mentor to me, and with his very unique style, he guided me during this journey. His wise smile is forever etched in my memory. Professor Pierre Lasserre has been very supportive right from the start of my studies, and he helped me immensely to improve the quality of my work. I am forever indebted for his patience, his friendly guidance and many hours spent discussing some of the fundamental ideas in this thesis. I am also grateful for the generous and constant support of Professor Dennis Kira.

Over the years, I have drawn very valuable inspiration and encouragement from my fellow Ph.D. students. I was especially fortunate to benefit from the wisdom and great friendship of Alejandro Karam, Souheyl Touhami and Daniel Wong.

Heather Thomson was the first person I met at Concordia University, and she has supported every step of my work by her incredible energy, her enthusiasm and her amazing abilities to navigate the administrative hurdles of the joint Ph.D. program.

Chapter 3 of this thesis is based on my joint work with Professor Jean-Marie Bourjolly and Aleksander Lyngvi. It was an honour for me to work with them and I am very grateful for their contributions.

Finally, I would like to express my immense gratitude to my parents without whose numerous sacrifices I wouldn't be able to write these words.

Table of Contents

List of Figures

List of Tables

Chapter 1
Introduction

During the past two decades, advances in computer science and operations research have resulted in many new optimization methods for tackling complex decision-making problems. One such method, *Tabu Search,* forms the basis of this thesis. Tabu search is a very versatile optimization heuristic that can be used for solving many different types of optimization problems.

Another research area, real options, has also gained considerable momentum during the last two decades. Real options is emerging as a robust and powerful method for tackling decision-making problems under uncertainty. Although the theoretical foundations of real options are solid and significant progress has been made in the theory side, applications are lagging behind. The advanced mathematics behind real options intimidates practitioners, especially managers, and slows down the adoption rate of this method in the corporate world. There seems to be a gap between theory and practice.

1.1 Rationale of the Thesis

A strong emphasis on practical applications and a multidisciplinary approach form the basic rationale of this thesis. The fundamental concepts and ideas behind tabu search and real options are investigated in order to provide a concise overview of the theory supporting both of these two fields. This theoretical overview feeds

into the design and development of algorithms to solve real-life decision-making problems.

In order to help bridge the gap between theory and practice, two user-friendly software tools were developed, *Orbital Optimizer* and *Options Optimizer.* Although these tools were designed for very different types of problems, they share a common "optimization engine" based on tabu search.

These software tools are tested extensively for accuracy and consistency by using real-life data as inputs. At the end of this process, a set of practical tools are obtained to solve both deterministic and stochastic decision-making problems.

The process of developing the software tools and performing the analyses is not a linear one. Many series of iterations are needed to give the tools their final form. These tools are then used to conduct the analysis for each of the problems considered in this thesis.

1.2 Research Objectives

This thesis aims to achieve the following objectives:

i) To extend the application of tabu search to a new physical environment: space. It is hoped that the techniques and the overall approach demonstrated in this thesis for modeling the space environment can be useful for leveraging our know-how and experience in Earth-based problems to space-based ones.

ii) To bridge the gap between theory and practice in real options by using tabu search and Monte Carlo simulation together. The Tabu Search Monte Carlo

(TSMC) method developed in this thesis is an intuitive method that can be used to solve complex, multidimensional option problems. The software implementation of TSMC presented in this thesis has the added advantage of being an interactive and visual tool, enhancing its usability (please see Appendix A.3 for screenshots of TSMC software implementations).

iii) To demonstrate the benefits of a multidisciplinary approach for solving complex problems. This thesis is based on the intersection of different disciplines, including operations research and computational finance. It is hoped that the thesis results will strengthen the case for cross-fertilization across different disciplines.

1.3 Research Method and Strategy

The research method of the thesis consists of the following steps.

An extensive literature survey lays the groundwork for the applications by identifying the key concepts of tabu search and real options discussed in the literature.

The literature survey is followed by the definition of the deterministic and stochastic problems that are tackled as part of the thesis work. After the problems are defined, algorithm development starts. The desired inputs for and outputs from the optimization algorithms are determined, the overall flow of the algorithms are planned and simple instances of the problems are studied to gain a deeper understanding of the nature of the problem and the interdependencies between different problem elements.

Once the flow of an algorithm is planned, the programming stage begins. Using the object-oriented C++ environment, the computer code is designed and developed for the algorithms. During this iterative development process, different components of the software tools are tested, and software parameters are calibrated using real-life data. Benchmarking the performance of the code by using test problem instances in the literature helps identify the strengths and weaknesses of the proposed approach.

Data gathering and filtering are the next steps. Using high-quality and representative data is a strong requirement for obtaining the desired results from the optimization tool.

Once the data and the software tools are in place, the analysis begins. By running the optimization tools numerous times, a large set of results are obtained. These results are analyzed using statistical and graphical tools in order to identify possible solutions.

At the end, recommendations are drawn based on the obtained solutions, areas of further work are identified, as well as the lessons-learned to improve the performance of the algorithms.

1.4 Thesis Overview

Following the introductory chapter, a literature survey is presented which covers the fundamental concepts of tabu search and real options.

The first problem presented is the application of tabu search to a deterministic problem. This application demonstrates the functioning of tabu search, and how it can be implemented in different types of problems. The advantage of tabu search for solving computationally complex problems is also shown. The deterministic problem is related to the servicing of satellites on Earth's orbit, a time-dependent, moving-target optimization problem.

Then, tabu search and real options are investigated together within the context of a stochastic optimization problem: option valuation. By merging tabu search and Monte Carlo simulation, a new method for studying options, Tabu Search Monte Carlo (TSMC) method, is developed. The theoretical underpinnings of the TSMC method and the flow of the algorithm are explained. Its performance is compared to other existing methods for option valuation.

The application of TSMC to a real-life decision-making problem forms the next chapter. This is a stochastic problem that investigates the conditions of feasibility for hybrid electric vehicles and fuel cell vehicles in the future. In this problem, the technological uncertainty of future passenger vehicles and the market uncertainty of competing fuel alternatives (hydrogen and oil) create a complex landscape. By using the TSMC method, these multiple sources of uncertainty are analyzed.

The final chapter of the thesis includes a general discussion of the research topics addressed in the thesis and an overview of main findings, identification of future research areas and conclusions.

Chapter 2
Literature Review

The foundation of this thesis is built on several fields of research. Fundamental ideas, definitions, concepts and methods from Operations Research, orbital mechanics, computational finance and real options are explored and integrated in order to tackle different types of optimization problems. In this section, the relevant literature from these fields will be presented so as to provide a general overview and to establish the foundations of the multidisciplinary approach that is used in the subsequent chapters. The chapters on the applications of tabu search also include literature reviews, with a specific focus on the connections between these chapters and the literature.

2.1 Computational Complexity

Throughout the thesis, certain everyday words such as "problem" and "solution" are used extensively. In order to avoid any confusion, it is important to give the operational definitions of these key concepts.

Garey and Johnson (1979) define a problem as *"a general question to be answered, usually possessing several parameters, or free variables, whose values are left unspecified"*. They also state that a problem is described by:

(a) a general description of all problem elements

(b) a statement of the properties that the answer, or *solution*, is required to satisfy.

A problem instance is related to the notion of a problem, but it also has a specific meaning. A problem instance is created when particular values are specified for all the problem parameters (Garey and Johnson, 1979). For instance, "finding the shortest path on a graph" is a problem, whereas "finding the shortest path on a complete graph with 4 vertices" is a problem instance (a complete graph is a graph in which each pair of vertices is connected by an edge).

The relationship between a problem and a problem instance is analogous to the relationship between a class and an object in object-oriented programming languages. Once a class is defined, it can be cast as an object by specifying the values associated with the parameters of the class.

An algorithm is a step-by-step procedure (or a "recipe") for solving all instances of a problem. In most cases, algorithms are computer programs written in a specific programming language (Garey and Johnson, 1979). All of the algorithms designed and developed for this thesis were written in C++.

In general, algorithms are designed to be as efficient as possible, so that computing resources are used in the best possible way. Arguably, time is the primary resource when it comes to computing, and therefore computer scientists refer to "time complexity" when they want to assess the efficiency of an algorithm with respect to its time requirements. Given the "size" of a problem instance, the time complexity function for an algorithm gives the time requirement of the algorithm by calculating the largest amount of time needed to solve a problem instance of that size (Garey and Johnson, 1979). In other words, time complexity measures, in the worst-case scenario, the amount of time it takes for the algorithm to solve the

given problem instance. The size of the problem instance is defined as the amount of input data (generally expressed as the number of bits) required to encode the parameters of the problem instance.

Another way to evaluate the computational complexity of an algorithm is to determine the number of elementary operations (e.g., addition, comparison, etc.) that are needed *at most* to solve an instance of size n (Reinelt, 1994). In other words, how many elementary operations are needed, in the worst-case, to solve a particular instance.

The time complexity function for an algorithm gives the largest amount of time needed by an algorithm to solve a problem instance of a specific size. This function is expressed using the "big oh notation": $O(n)$ (O stands for "on the order of"). An algorithm has polynomial time complexity, if its worst-case complexity is $O(p(n))$ for some polynomial function p. Polynomial algorithms are very desirable, since they can be used to solve large problem instances in a reasonable amount of time. However, not all algorithms have polynomial complexity, and for large problem instances the time requirement for polynomial vs. non-polynomial algorithms can be very different.

In order to demonstrate the concept of time complexity, consider solving a multidimensional American option valuation problem[1] with an exact method, such as the finite difference algorithm. This algorithm works by creating m time steps, and then divides the price path of each of the n assets into p price levels. In this way a multidimensional finite difference grid is created with $O(mp^n)$ nodes. The

[1] See Chapter 4 for a formal definition of this problem.

time complexity of this algorithm is also in the same order as $O(mp^n)$ (Broadie and Detemple, 2004). The memory and time requirements for solving this problem can be reduced by storing only the price information of two adjacent time steps (instead of the whole grid). This reduces the memory requirement to $O(p^n)$. However, for all practical purposes, this efficiency gain is negligible, since both the memory and computation time requirements are exponential in the number of assets, n.

Most exponential time algorithms are variations on exhaustive search (Garey and Johnson, 1979), and the finite difference algorithm is also an exhaustive method. The advantage of such algorithms is guaranteed optimality, since all possible solutions are visited. For instance, in the case of the finite difference algorithm, starting from the end nodes, all possible price nodes are evaluated recursively until time zero is reached. As such, exhaustive search algorithms provide a high-level of accuracy.

If a problem can be solved with a polynomial time algorithm, then it is considered as an "easy" problem. If a problem is so hard that no polynomial time algorithm can possibly solve it, then it is called "intractable" (Garey and Johnson, 1979).

Intractable problems belong to a large class of problems (called NP-complete) for which there are no known polynomial time algorithms and almost every expert believes that none of these problems can be solved with polynomial time algorithms. As it will be discussed in section 2.3, the Traveling Salesman Problem (TSP) is one of the most famous classes of problems that belong to this category.

As discussed in the American option valuation example, dynamic programming based exact algorithms also exhibit more than polynomial complexity.

Faced with an NP-complete problem, a justified solution strategy is to develop an algorithm that is not guaranteed to find the global optimum, but one that can find a close-to-optimal solution in a reasonable amount of computing time. Such algorithms are called "heuristics" and they are discussed in detail in section 2.4.1.

2.2 Optimization

Many problems can be structured as decision problems or as optimization problems. A decision problem only seeks a yes-or-no answer, whereas an optimization problem requires a more detailed answer. For instance, as a decision problem, given the map of Paris, a visitor may ask "Is there a tour on this map that covers the Eiffel Tower, the Louvre Museum and the Notre Dame Cathedral in less than 5 kilometers?". If the same visitor asks "What is the length of the shortest tour that connects these three sites?", then this becomes an optimization problem. In general, if there is an efficient algorithm for answering the decision problem, then with a little more work, the optimization problem can be solved as well (Wilf, 1994).

When the problem elements and solution properties can be defined by linear functions, the problem can be solved with linear programming. However, if the problem elements and solution properties are discrete (e.g., where the solution is a set (or sequence) of integers), then the problem of finding optimal solutions to such problems is known as combinatorial optimization (Reeves, 1993).

An interesting characteristic of combinatorial problems is the presence of local and global optima. While there may only be a single true optimum (the global optimum), there are many other solutions that are only "locally" optimal (Reeves, 1993).

The concept of a neighborhood is a very useful tool in combinatorial optimization, and it can also help clarify the difference between local and global optima. The neighborhood of solution S, $N(S)$, is composed of solutions that can be reached from solution S, by a simple operation. For example, if solution S is the following sequence, $\{0, 1, 0, 0\}$, and if the operation is a binary switch, then solution S' $= \{1, 1, 0, 0\}$ can be reached from solution S by substituting 1 as the first element in the sequence. Thus, solution S' is in the neighborhood of solution S.

When a *greedy algorithm* is implemented, the objective is to reach to the global optimum in a step-by-step fashion. Such an algorithm proceeds from the best solution found in a neighborhood (the local optimum) to the next neighborhood, again searching for the local optimum. At the end of the search process, the best local optimum found is also the best solution obtained. However, greedy algorithms are not usually the best method for solving combinatorial problems. Instead of picking the single best solution at each step, another approach can be to pick another solution, in the hope that at a later step we may find even a better solution and get closer to the global optimum (Wilf, 1994).

One major disadvantage of greedy algorithms is that they are prone to getting trapped in local optima (this is also called *cycling*). Since a greedy algorithm proceeds from one local optimum to the next, if there are no better solutions found

in the next neighborhood, then the algorithm will revert back to the last local optimum found and search the same neighborhoods over and over again without any further improvements. In order to avoid this cycling, the algorithm has to include other rules that guide the search process which can enable accepting temporary deteriorations in the solution in exchange for the option of finding better solutions in the long-run.

2.3 Traveling Salesman Problem

The Traveling Salesman Problem (TSP) is a classical problem in the NP-complete class of problems. Given a collection of cities and the cost of travel between each pair of them, the objective is to find the shortest way of visiting all of the cities and returning to the starting point (hence completing a tour). Although the problem statement is very simple, solving the TSP for large problem instances is a very challenging task due to the NP-complete nature of the problem.

The formal definition of the problem is as follows (Garey and Johnson, 1979):

Given a finite set of cities $C = \{c_1, c_2, ..., c_m\}$ and the distance $d(c_i, c_j)$ between each pair of cities c_i, c_j, minimize the total distance of a tour that starts at city $c_{\pi(1)}$ and ends at the last city $c_{\pi(m)}$. Specifically, minimize:

$$\left[\sum_{i=1}^{m-1} d\left(c_{\pi(i)}, c_{\pi(i+1)}\right)\right] + d(c_{\pi(m)}, c_{\pi(1)})$$

The solution to this combinatorial optimization problem is expressed as a sequence of the optimal tour $\langle c_{\pi(1)}, c_{\pi(2)}, ..., c_{\pi(m)} \rangle$.

Although the terms "city" and "distance" seem to indicate a very specific problem, the TSP is amazingly rich, and there are many real-life applications of this problem. The primary domain of application is the transportation sector, where many variants of the TSP are used to design and operate transportation networks. Other areas of application include telecommunication network design, various manufacturing problems (e.g., circuit board design) and genomics.

A related problem to TSP is the Vehicle Routing Problem (VRP). In a VRP, there are multiple salesmen (or vehicles) and each city has to be visited by exactly one salesman. In this sense, one can think of the TSP as a special case of the VRP, in which there is a single salesman. In the VRP, every salesman starts from the same city (i.e., the depot), and must return to the depot at the end of his/her assigned tour. There are many variants of the VRP as well, such as imposing a capacity constraint on the salesman (for delivering and/or receiving items) and finding the minimum number of salesmen required for a given VRP.

2.4 Tabu Search

When faced with a combinatorial optimization problem, one of the first instincts is to use exhaustive enumeration: evaluating all possible solutions, and then picking the best one. Although in theory this approach can yield the optimal result, for many (if not all) combinatorial optimization problems, its application is severely limited due to computational complexity. Therefore, especially when we need to solve an NP-complete problem, finding a high-quality solution in a reasonable amount of

time becomes a priority. When the problem size is relatively small, exact methods can be used to find optimal solutions. When larger instances need to be solved, metaheuristics such as tabu search are ideal for finding near-optimal solutions.

2.4.1 Fundamental Definitions and Concepts

Tabu search is a metaheuristic that has been developed to tackle NP-complete problems while avoiding the pitfalls of greedy algorithms. The linguistic origin of the word heuristic is from Greek, *heuriskein*, which means to find (Blum and Roli, 2001). The technical definition of a heuristic is as follows (Reeves, 1993):

"A heuristic is a technique which seeks good (i.e., near-optimal) solutions at a reasonable computational cost without being able to guarantee either feasibility or optimality."

Broadly speaking, heuristics can be divided into two groups: problem-specific heuristics and metaheuristics. Problem-specific heuristics are designed to solve a particular problem, and they are tailored to the characteristics of the problem at hand. Metaheuristics, on the other hand, are like templates, they can be implemented for solving very different problem types without a need for extensive customization.

A formal definition of metaheuristics is given by Osman and Laporte (1996): *"an iterative generation process which guides a subordinate heuristic by combining intelligently different concepts for exploring and exploiting the search space, learning strategies are used to structure information in order to find efficiently near-optimal solution."*

Tabu search, first proposed by Glover (1986), is a method that imposes certain restrictions to guide the search process. These restrictions can have different forms, but they mainly operate by "forbidding" certain search alternatives (hence the name "tabu" - or "taboo"), and also by altering the probabilities of selection of search alternatives based on the history of the search process (Laguna, 1994).

The most common tabu restrictions are put in place to prevent cycling, thereby avoiding a common shortcoming of greedy algorithms: being trapped in local optima. Even though a move may deteriorate an existing solution, or render it infeasible, it can still be accepted in a tabu search algorithm. Therefore, a tabu search algorithm can move away from local optima and keep the option of moving towards the global optimum alive.

Flexible memory is at the heart of the tabu search method. Flexible memory works by storing information about the search process, such as the frequency and recency of certain moves, as the search algorithm is running. This information is used to guide the search process during the subsequent moves, by intensifying the search around promising solution spaces, or diversifying it away from not so promising ones.

During the search process, it is very likely that some backtracking will occur and the same solutions will be revisited. In order to avoid the backtracking to happen too soon (and in the worst-case, transform into cycling), a tabu list is used. The tabu list constitutes the short-term memory component of the algorithm, and it records recently implemented moves. For a given number of iterations, the moves in the tabu list are forbidden. After each iteration, the tenure of the moves in the

tabu list are decremented. When the tenure of a move reaches 0, it is taken out of the tabu list.

An evaluation function is a standard feature in tabu search algorithms. This function is used to assess the performance of each move. If a move provides the best solution for a given neighborhood and if it is not in the tabu list, it is implemented. This way, the algorithm proceeds from one best solution to the next *in each neighborhood*, while recording the best overall solution obtained during the whole search process. Note that, as mentioned above, the best solution of a neighborhood may be worse than the preceding neighborhood's best solution. It can also be an infeasible solution. In both cases, it is accepted as the current solution. This is a major difference between tabu search and a greedy algorithm.

Sometimes, a move that is in the tabu list can provide the best overall solution. In this case, an aspiration criterion is employed to accept the move even though it has tabu status.

Once a predetermined criterion is met, such as the number of iterations, or the rate of progress in obtaining better solutions, the search is terminated. At the end of the search, not only the best overall solution is recorded, but also all the relevant information about the search process, including the recency and frequency information of various moves.

Consider the following example that demonstrates the basic principles of tabu search. A blind man finds himself on a large garden shaped as a chessboard. He does not know which cell he is on, and his objective is to find a heat source that is located on another cell, which can be anywhere in the garden. At each itera-

tion, he can take a single step from the cell he is on, and move to another cell (including the possibility of diagonal moves, he has a total of 8 moves that are possible, hence his "neighborhood" is composed of 8 moves). He has a method for evaluating his search performance based on the change in temperature from one cell to the other (this method is his "evaluation function"). He can also track his search process by recording his past moves, and selecting future moves that have the best potential. For instance, if he moved "north" on the chessboard garden, he can exclude moving "south" in the next iteration, so that he does not backtrack (recency information). Similarly, if he has taken a number of steps going "north" without any apparent increase in temperature, he can direct his search in another direction (frequency information). This way, by avoiding backtracking, and by picking moves that have a higher chance to find the heat source, he can direct his search. The frequency and recency information that he collects forms his flexible memory.

2.4.2 Applications

Glover and Laguna (1997) provide one of the most comprehensive overviews of the tabu search method, and they also discuss various application areas. The implementation guide of Laguna (1994) is another source that provides the basics of tabu search in a step-by-step fashion. Another general guide is that of Gendreau and Potvin (2003), which presents the basics of tabu search as well as guidelines for implementation.

Tabu search is very widely used for tackling traveling salesman problems (TSPs) and vehicle routing problems (VRPs) and the variants of these problems.

Malek et al. (1989), Knox (1994), and Fiechter (1994) used tabu search for finding near-optimal solutions to various size TSPs. Malek et al. (1989) and Fiechter (1994) have also demonstrated parallel implementations of tabu search.

Cordeau and Laporte (2002) present various tabu search algorithms that have been developed to solve VRPs. They provide the overviews of various tabu search algorithms, including Taburoute (Gendreau, Hertz and Laporte, 1994), Taillard's decomposition method (Taillard, 1993), Osman's algorithm (Osman, 1993) and Unified Tabu Search Algorithm (Cordeau, Gendreau, Laporte, 1997). Cordeau and Laporte (2002) report that the algorithm of Taillard (1993) remains as one of the most effective VRP implementations of tabu search.

Probabilistic diversification and intensification methods have also been used to increase the performance of tabu search in VRPs. Rochat and Taillard (1995)'s Adaptive Memory Procedure can be used during the search or as a post-optimizer to increase the quality of solutions. Their approach is not specific to VRPs and it has been implemented in other problem types as well (see for instance the political districting problem tackled by Bozkaya, Erkut and Laporte (2003)).

In addition to TSPs and VRPs, tabu search has been implemented in many other problem types including planning and scheduling, telecommunications and inventory management. For a review of tabu search applications see Glover and Laguna (1997).

2.4.3 Why Tabu Search?

For this thesis work, tabu search is chosen as the main optimization method. In theory, one can also choose other types of metaheuristics, such as simulated annealing, genetic algorithms or ant colony optimization in order to solve problems similar to those presented in this thesis. However, tabu search has a proven track record in solving TSPs (see Reeves (1993); Blum and Roli (2001)), and therefore, it is an ideal starting point for solving space-based TSPs. Application of tabu search for solving stochastic optimization problems is not as common, with very few metaheuristics examples in the computational finance literature. Although the choice of tabu search for Chapters 4 and 5 is not as conventional, the versatility and flexibility of tabu search render it a suitable candidate for these two applications as well.

2.5 Real Options

2.5.1 Fundamental Definitions and Concepts

Real Options are a set of methods that can be used to evaluate decision alternatives under uncertainty. The intellectual origin of real options is the options theory in finance. Options are financial instruments that give owners *the right but not the obligation* to exercise their rights to buy or sell financial assets in the future. This is also the fundamental idea behind real options: future investment decisions can be formulated as decisions to obtain and manage certain real options. For instance,

Characteristic	Type of Option
Right to buy or sell	A *call option* gives the option-holder the right to buy the underlying asset, whereas a *put option* gives the right to sell it.
Early exercise features	A *European option* can only be exercised at maturity. An *American option* can be exercised at any time before (or at) maturity.
Underlying asset	There are many different underlying assets for options contracts, such as stocks of companies, stock market indices, and currencies.
Combinations	Combining different features of options have resulted in many "exotic options", such as compound options (options on options), rainbow (or basket) options (options on two or more risky assets) and Bermudan options (options that can be exercised at only certain points before maturity)

Table 2.1: Types of Options

acquiring a license for a particular technology brings the option of commercializing that technology if the market conditions are favorable.

The exercise price and the lifespan (or maturity) of financial options are predetermined at the time of contract creation.

As can be seen in Table 2.1, there are many different types of options.

In 1997, the Nobel Prize in Economic Sciences was awarded to the pioneers of options theory, Robert C. Merton and Myron S. Scholes (Fischer Black, who did some of the original work in option theory, died in 1995). The official citation of the prize underlines the importance of their contribution:

"Black, Merton and Scholes thus laid the foundation for the rapid growth of markets for derivatives in the last ten years. Their method has more general applicability, however, and has created new areas of research - inside as well as outside of financial economics. A similar method may be used to value insurance contracts and guarantees, or ***the flexibility of physical investment projects.***"

Indeed, financial markets are not the only places where options can be useful. The uncertainty associated with future conditions, be it economic, natural or political, is a common feature of our everyday lives. Decision making under uncertainty is a universal challenge, and "options thinking" can be applied to many real-life problems as well.

Real options is the general term that is used for structuring, analyzing and solving such problems by modeling them as options, and applying the tools and techniques that have been originally developed for financial options. By modeling their future investments as options, managers can measure the value of flexibility, they can hedge their bets and control the overall risks present in their portfolio of projects without being ultra-conservative.

2.5.2 Historical Overview

Although the field of real options has been experiencing significant growth during the past two decades, the main ideas and concepts behind real options have a very long history. Perhaps one of the first documented accounts of real options can be found as early as 6th century B.C. in the story about Thales. In *Politics*, Aristotle tells the story of Thales (Aristotle, 1998):

"The story goes that when they found fault with Thales for his poverty, supposing that philosophy is useless, he learned from his astronomy that there would be a large crop of olives. Then, while it was still winter, he obtained a little money and made deposits on all the olive presses in Miletus and in Chios. Since no one bid against him, he rented them cheaply. When the right time came, suddenly

many tried to get the presses at once, and he rented them out on whatever terms he wished, and so made a great deal of money. In this way he proved that philosophers can easily be wealthy if they desire, but this is not what they are interested in."

The story of Thales demonstrates almost all of the core concepts of real options. As a proactive decision-maker, he obtained the right, but not the obligation of using a physical asset (the olive press) by paying the option premium (the deposit). Faced with the uncertainty of the quality and abundance of the harvest, he created a strategy for maximizing his potential for profit while minimizing his downside risk (all he had to lose was his deposit). The story is also a very good example for explaining the nature of uncertainty and risk; in this case, Thales actually benefited from the uncertainty. By controlling the downside portion of his risk, and leaving the door open for beneficial developments in the future, he obtained a great amount of flexibility to manage his overall risk.

Although the intellectual origins of real options can be traced back for many centuries, only in 1900 was a viable method proposed to calculate the value of an option. In this year, Louis Bachelier defended his thesis entitled "Théorie de la Spéculation" in which he derived the price a financial option (Bachelier, 1900). Bachelier's work laid the foundations of many modern concepts and techniques in finance and economics, including the Brownian Motion.

The 1970s brought a revival in the field of option theory: in 1973, Black and Scholes (1973) published the now-famous paper on European option valuation and Merton (1973) generalized their model to include dividend-paying stocks. My-

ers (1977) is cited as one of the first instances where the term "real options" was used and in the following years research on real options started to gain momentum. In 1990s, many books and articles were published that developed the theory behind real options, and extended its application to new areas. The theoretical contributions of many researchers, including Dixit and Pindyck (1994) and Trigeorgis (1996), formed the cornerstones of real options methods.

2.5.3 Real Options in Natural Resources

One of the first application areas of real options was natural resources. Since spot and futures prices from commodity markets were available, researchers were able to calculate critical inputs required for real options analysis from market data, such as volatility and convenience yields.

Using the real options framework, researchers evaluated mining and oil exploration investments. Tourinho (1979) used an options-based method to determine the value of a reserve of a natural resource when the future price of the resource is uncertain. Brennan and Schwartz (1985) evaluated the option to shut down or abandon a copper mine. Paddock, Siegel and Smith (1988) examined the valuation of offshore petroleum leases and gave one of the first examples of a "hitting boundary" in real options literature (also called the "early exercise boundary"). The early exercise boundary plots the critical values of the underlying asset that trigger early exercise of the option. For a call option, above the boundary, the option is in-the-money and under it, the option is out-of-the-money. As will be discussed

in Chapter 4 in detail, the early exercise boundary is one of the central features of the TSMC method.

2.5.4 Real Options in Research and Development

Another major application area of real options is research and development (R&D) projects. Although for many years managers had the intuition that strategic technology investments were a major driver for growth, traditional capital budgeting techniques, such as the net present value (NPV) method, are inherently biased against such investments. Most R&D efforts are long-term, highly uncertain ventures. In most cases, no matter how high the future benefits of an R&D project may be, the present value of such benefits decreases considerably when analysts use high discount rates to account for the risk. This shortcoming, or "myopia" of traditional methods consistently favor shorter-term, low-risk projects over long-term, high-risk projects. Consequently, if one were to use NPV as the main criterion for selecting projects, then the product/service portfolios of most companies would lack breakthrough innovations and radically new products and services.

One of the earlier works in this application area was that of Mitchell and Hamilton (1988). These researchers suggested the following conceptual framework for modeling R&D projects as real options:

- an R&D project can be modeled as a call option on the future stream of benefits that the project will bring (the underlying asset)
- the price of this option is analogous to the cost of the R&D project

- the exercise price is analogous to the cost of commercializing the results of the project

This framework takes into account the managerial flexibility present in the R&D project: at the end of the project, if the technical results are not satisfactory and/or the potential market has deteriorated, then the organization can decide not to commercialize the project results. In other words, the organization has the right, but not the obligation to commercialize project results.

Although the option analogy for R&D projects establishes a promising framework, quantifying some of the model parameters are considerably more difficult than in the case of natural resources. In the R&D case, the underlying asset is not traded in any organized marketplace (therefore there are no spot and future prices that are available). Furthermore, the decisions of the organization has direct effects on the value of the underlying asset (which is not the case for most commodities which are traded by a very high number of parties). Therefore, the decisions of the organization during the lifetime of the project have to be explicitly modeled.

Applications of the real options method to R&D project selection can be found in Boer (1999), Lint (1998), Luehrman (1998), Schwartz and Moon (2000), Paxson (2003) and Shishko (2003).

2.5.5 Modeling and Solving Real Options Problems

There are a wide variety of methods that can be used for modeling and solving real options problems. These methods can be divided into three main groups:

(i) Analytical approaches

(ii) Numerical methods

(iii) Simulation-based methods

These methods are discussed in detail in Chapter 4.

Chapter 3
On-Orbit Servicing: A Time Dependent, Moving Target Traveling Salesman Problem

3.1 Introduction

The robotic capability of maintaining and repairing space assets, on-orbit servicing (OOS), has the potential of changing the way spacecraft are designed, manufactured and operated. Although OOS is still in its infancy, interest in this area is growing, and a number of technology demonstration missions are planned to test the technical and economic feasibility of this concept.

The most common OOS mission concept envisions an orbital "depot", where consumables and spare parts for spacecraft will be stored. A "servicing platform", based at this depot, will be used to service a number of client spacecraft and then return to the depot for resupply. Minimizing the total amount of energy or time or a given combination of both during these operations is crucial for the long-term technical and commercial success of OOS.

Modeling OOS operations as a space-variant of the Traveling Salesman Problem (TSP) and the Vehicle Routing Problem (VRP), and developing practical solution methods can enable the mission designers to optimize the use of their on-orbit resources, and in doing so increase the probability of success for such missions.

Beyond their practical significance for OOS missions, space-based TSPs and VRPs also contribute to the effort of modeling and solving time-dependent and moving-target TSPs and VRPs. By its very nature, the OOS problem belongs to this class of problems, since the position of a spacecraft around a planet's orbit changes constantly over time. Compared to the static cost/distance matrices of most Earth-based problems, the dynamic nature of the problem presents a significant challenge. Within the OOS context, we address this challenge by embedding orbital mechanics into optimization algorithms, and creating dynamic cost/distance matrices. For small problem instances (with a handful of client spacecraft), exhaustive search can be quite effective and efficient. A basic tabu search implementation is proposed for larger problem sizes. Note that in the foreseeable future, missions of this kind will only involve a handful of satellites and therefore, exhaustive search should be sufficient.

The main motivation of this chapter is to illustrate the usefulness of Operations Research methodology for modeling and solving a space-related problem that is a variant of the time-dependent, moving-target TSP. It is hoped that this real-life example will rekindle the interest in this class of problems.

3.2 Background

The concept of on-orbit servicing itself is not new, and there have been a number of servicing missions in the past involving astronauts such as the Hubble Space Telescope repair missions (Britt, 2002). However, given the inherent risk of human

spaceflight as evidenced by Challenger and Columbia accidents and the high cost of such flights, there is a clear need for a semi-autonomous, robotic servicing system that can perform servicing operations in a cost-effective manner. There have been many conceptual studies in the past investigating the technical and economic aspects of such systems (Waltz, 1993; Dunn, 1987; Saleh, Lamassoure and Hastings, 2002). As of 2006, there are both governmental and private sector initiatives to provide on-orbit services, and recent developments indicate that an operational system will be launched within a decade (Morris, 2004; David, 2003).

The current mission concepts for servicing systems envision an orbital depot, where propellant, spare parts and other consumables for the client spacecraft will be stored. A servicing platform will bring these consumables to the client spacecraft that need servicing, and the onboard robotic systems will be used to perform the necessary servicing operations. Every servicing tour has to start and end at the orbital depot, where the servicing platform is resupplied. During a servicing tour, the servicing platform can visit more than one client satellite depending on a number of factors such as the nature of the servicing mission and the orbital geometry of spacecraft involved in the servicing tour.

There are over 600 active spacecraft orbiting around the Earth (TEAL, 2001), and most of these space-based assets are satellites that are used for diverse purposes, such as telecommunications, remote sensing and navigation. Other types of spacecraft include scientific satellites and space transportation systems, such as the U.S. Space Shuttle fleet. Initially, only a fraction of these spacecraft are expected to be clients for on-orbit services. From an optimization point of view,

this means small problem sizes, and consequently, higher chances of achieving optimality. Moreover, greedy and exhaustive search methods may prove effective.

Due to the use of more reliable subsystems, advanced materials and increased know-how, the average lifetime of satellites has increased considerably during the last three decades. However, despite all these advances, a satellite's lifetime is still critically dependent on onboard resources. One of the most critical resources is onboard propellant that is used for maintaining a satellite's orbit and *attitude*. The attitude of a satellite determines its orientation in space, and therefore is crucial for pointing onboard instruments and equipment to desired targets. When the onboard propellant is depleted, the satellite cannot be controlled, and it can no longer be kept in a stable orbit, due to various forces acting on it such as perturbations in the gravity field around the Earth. In addition to the depletion of propellant, subsystem failures can also cause a satellite to fail (past data indicates that on-orbit failures account for about 40% of total satellite mission failures (Lacombe, 2002)). Furthermore, solar array problems, battery problems, electronic component failures and thruster failures are the most common sources of on-orbit failure (Frost, 2004).

The robotic capability of maintaining and repairing a satellite can surpass these limitations by providing added operational flexibility through services such as refueling, repair operations, and orbital correction maneuvers. Therefore, coupled with a modular design paradigm, on-orbit servicing can help extend the average lifetime of satellites even further.

However, this additional flexibility comes with a cost: the servicing platform and the required cargo have to be maneuvered from the depot to the client spacecraft. These additional orbital maneuvers require considerable energy and time. Therefore, servicing missions have to be planned meticulously in order to use the energy and time that are required to conduct OOS missions as efficiently as possible.

A significant body of knowledge has been developed in Operations Research that can be used for optimizing OOS missions. That includes TSP and VRP algorithms. So far, almost all of TSP and VRP algorithms have been developed for Earth-based applications. The modification of these algorithms for the OOS problem is not trivial, and the specific nature of the space environment has to be taken into account at every step of implementation. To our knowledge, there has been a limited amount of research in this area. OOS has been modeled as a VRP, and a basic solution method was proposed using the Clarke and Wright algorithm (Gurtuna and Trepanier, 2003). However, this method did not take into account the time-dependent nature of the problem, which is discussed in detail in Section 3.2.1.

The following sections outline the characteristics of the OOS problem, as well as the modeling and solution methods developed specifically for it.

3.2.1 Motion of Satellites

The way satellites move in orbit is completely different from the conventional modes of transportation such as land, air and sea. In space, satellites are in a continuous

state of motion and the distances between objects can change drastically over time.

The orbit of a satellite can be completely described with six variables, the classical orbital elements, as described in Table 3.1, and as shown in Figures 3.1 and 3.2. The *true anomaly* gives the position of the satellite in its orbit and is therefore time-dependent. Moreover, due to perturbations in the Earth's gravity field, the *argument of perigee* and *right ascension of ascending node* change with time as well.

Using orbital mechanics equations, the orbits of spacecraft around the Earth can be calculated with a high level of accuracy. Therefore, for shorter time horizons, the orbits of the spacecraft constitute a deterministic system. The accuracy of calculations decreases over longer time horizons (decades).

In order to move a satellite from one particular orbit to another, thrust is applied in a series of "burns". During these burns, onboard propulsion systems are fired with a specific force and duration in order to change the classical orbital elements in the desired manner. Each burn changes one or more of the classical orbital elements, and the applied energy is quantified by the applied velocity, ΔV. The total amount of time for transfer depends on a number of factors including the type of transfer (see Section 3.3.1), the amount of thrust, and the relative position of origin and destination orbits.

The timing of burns is critical during orbital maneuvers in order to change the orbital elements in the most efficient way. Since the spacecraft are in a dynamic en-

Element	Definition of Element
Semi-major axis, a	Describes the size of the ellipse; for a circular orbit this is equal to the radius of the circle.
Inclination, i	It is the angle between the orbital plane and Earth's equatorial plane.
Right ascension of ascending node, Ω	The angle between the vernal equinox and the ascending node. The ascending node is the point where the satellite passes through the equatorial plane moving from South to North.
Eccentricity, e	The shape of the ellipse (the deviation in form from a circle).
Argument of perigee, ω	The angle from the ascending node and the orbit's point of closest approach to the Earth (perigee).
True anomaly, ν	The angle between perigee and the spacecraft, measured in the direction of motion.

Table 3.1: Classical Orbital Elements

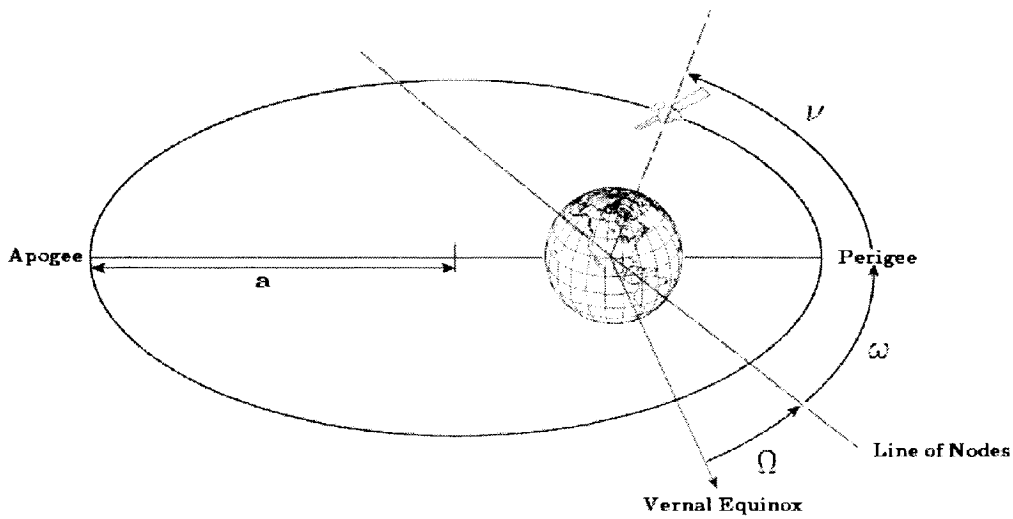

Figure 3.1: Semi-major axis, RAAN, Argument of Perigee and True Anomaly

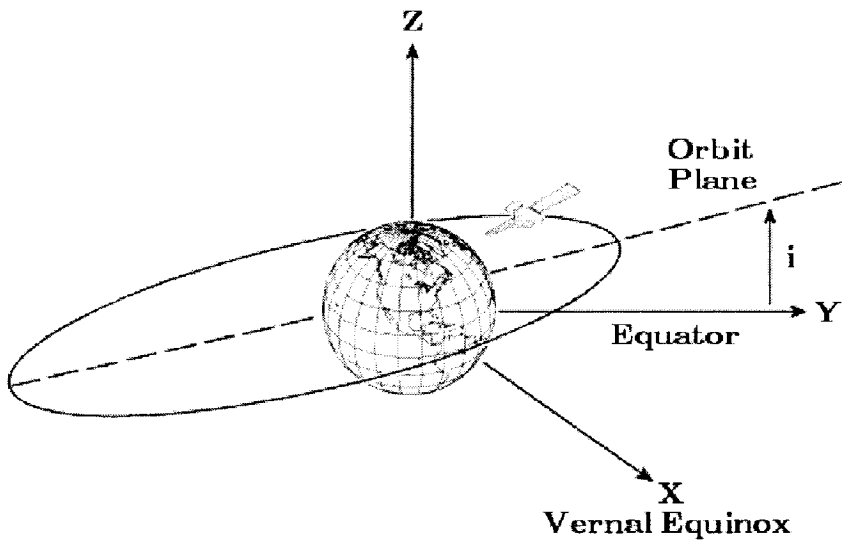

Figure 3.2: Inclination

vironment, the relative distances between them are changing drastically over time. Therefore, in the OOS problem, in order to rendezvous and dock with each client spacecraft and the orbital depot, the servicing platform has to perform a series of burns at precise moments during a servicing operation. This time-dependent nature of the problem has to be incorporated into the modeling of space-based TSPs and VRPs in order to provide practical solution methods.

3.2.2 Time-Dependent and Moving-Target Problems

Time-dependency has been studied in the past for various transportation problems. Ichoua et al. (2003) classify such problems into three main groups: time-dependent shortest path problems, time-dependent path choice problems and time-dependent TSPs and VRPs. Various modeling techniques have been used in the

past to capture the essence of time-dependency, including simple approximations, and discrete, continuous and stochastic travel time and cost functions.

An earlier work, by Malandraki and Daskin (1992) has focused on time-dependent VRPs. Malandraki and Daskin (1992) have developed heuristic algorithms for solving such problems assuming that the travel time between two nodes during a servicing tour is dependent on the distance between the nodes and the time of day. Using a discrete travel time and cost function technique, they generated piecewise linear functions for approximating the travel times between the nodes.

A more recent approach for tackling time-dependent VRPs has focused on implementing an ant colony system as the solution method (Donati et al., 2003) . The researchers concluded that in most experimental cases, the best solutions for constant speed (i.e., time-independent) problem instances might become infeasible or suboptimal if the problem is transformed into a time-dependent one by adding variability in traffic conditions. In this context, infeasibility arises when certain time constraints are violated, such as the time windows of individual clients or the depot. Ichoua et al. (2003) report similar results. They found that in most cases, the best results obtained for constant speed problems deteriorate (in terms of feasibility or optimality) when the problem is transformed into a time-dependent one.

Another line of research that is related to the OOS problem is moving-target (or kinetic) TSP. Helvig et al.(1998) modeled such problems where not only the salesman, but also the clients to be visited are in motion. Problems such as

in-motion refueling fall into this category (e.g., when a refueling aircraft has to rendezvous and refuel a number of other aircraft). Helvig et al.(1998) also provide heuristics and performance bounds for basic moving-target problems in which clients move on a linear path without accelerating.

The motivation for research on Earth-based time-dependent TSPs and VRPs is generally based on the desire for modeling traffic in an urban environment, and as such, these efforts can be seen as an incremental improvement over time-independent algorithms. In comparison, time-dependency is at the core of space applications. OOS definitely shares the characteristics of both time-dependent and moving-target TSPs. In fact, since both the servicing platform and client spacecraft are in a continuous state of motion, and since the duration and cost (ΔV) of travel depend critically on time, OOS can be described as a "time-dependent, moving-target" problem. What makes space-based TSPs and VRPs truly unique, however, is the physical dimension of these problems. Such problems have to be modeled in a three-dimensional space. In the case of OOS, this is achieved by using the classical orbital elements for describing the motion of satellites over time, and calculating the duration and cost of travel based on these elements.

The deterministic nature of satellite orbits facilitates the modeling process to a certain extent. Existing physical models are effective in determining the position of a satellite with a high level of accuracy. Although this deterministic nature does not hold for very long time scales, the rolling horizon for problems such as the OOS are generally less than five years, a time scale that allows for precise position calculations.

There is also an interesting tradeoff between time and energy cost of travel in space. Generally speaking, minimum energy transfers require longer time than energy-intensive transfers; the latter can change the orbital parameters faster. Depending on the nature of the mission, time can be a critical factor. For instance, military missions generally require faster response times, and military operators may want to spend more energy to decrease the response time. On the other hand, civilian operators may prefer lower-energy transfers for activities such as preventive maintenance and scheduled refueling which are not time-critical. Therefore, finding the efficient frontier that captures the tradeoff between energy requirements and servicing time can be very interesting for decision-makers.

3.3 Methodology

The main motivation for this research is to establish a solid framework for modeling and solving space-based TSPs and VRPs. To this end, a dedicated software application was developed in C++, *Orbital Optimizer*, which enables the user to optimize OOS operations, based on different criteria, such as energy, time or a combination of both.

A dataset including the classical orbital elements of 20 satellites was created using the Satellite Database of Satellite Tool Kit® software. This database includes the classical orbital elements of satellites around Earth. Using the Satellite Tool Kit® software, the orbits of the selected satellites were propagated and the orbit of each satellite was modeled for a five-year period (from 1 June 2002 to 1

June 2007). At the end, a data file for each satellite was obtained, which contained information about the satellite's position at any given time during this five-year period. These data files provided the time-dependent input that was needed for the optimization algorithms.

An alternative approach for obtaining datasets is generating random classical orbital elements and obtaining the five-year data for these hypothetical satellites using Satellite Tool Kit®. Such an approach can be useful for studying very large problem instances and improving the performance of optimization algorithms.

In order to achieve a fair representation, satellites in different orbital belts around Earth were included in the dataset. These belts include Low Earth Orbit - LEO (semi-major axis, $a \backsim 7,000km$), Middle Earth Orbit - MEO ($a \backsim 26,000km$) and Geosynchronous Earth Orbit - GEO ($a \backsim 42,000km$). The full list of the selected satellites is presented in Table 3.5. Note that they are in arbitrary order.

3.3.1 Orbital Maneuvers

In order for two spacecraft to rendezvous in space, the classical orbital elements of both spacecraft have to be equalized. Depending on the differences between the origin and destination orbits, one or more classical orbital elements may need to be changed. Table 3.2 lists the types of maneuvers that are most commonly used for rendezvous operations.

Hohmann transfer is considered as the most fuel-efficient transfer between two circular, coplanar orbits. It is composed of two burns: the first burn places the spacecraft on an elliptical transfer orbit. When the spacecraft reaches the

Parameter(s) to change	Symbol	Type of Maneuver
Semi-major axis	a	Hohmann transfer
Inclination	i	Simple plane change
Both the semi-major axis and the inclination	a, i	Hohmann transfer with a combined plane change
Angular phase	ω	Phasing orbits

Table 3.2: Common types of orbital maneuvers

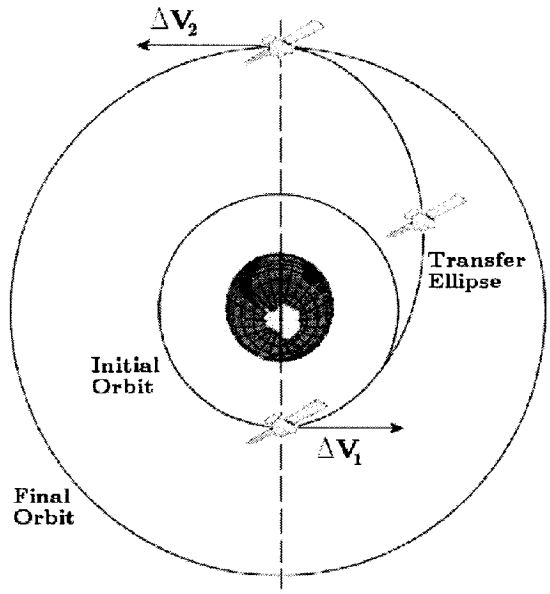

Figure 3.3: Hohmann Transfer

destination orbit, the second burn is applied in order to make the orbit circular again, and to finalize the maneuver (Figure 3.3).

Simple plane change is used when only the orientation of a satellite's orbital plane needs to be changed. With this type of maneuver, the size of the orbit (semi-major axis) remains constant.

Hohmann transfer is the most commonly used type of transfer in space operations, and for the OOS problem, it represents the primary method of transfer when

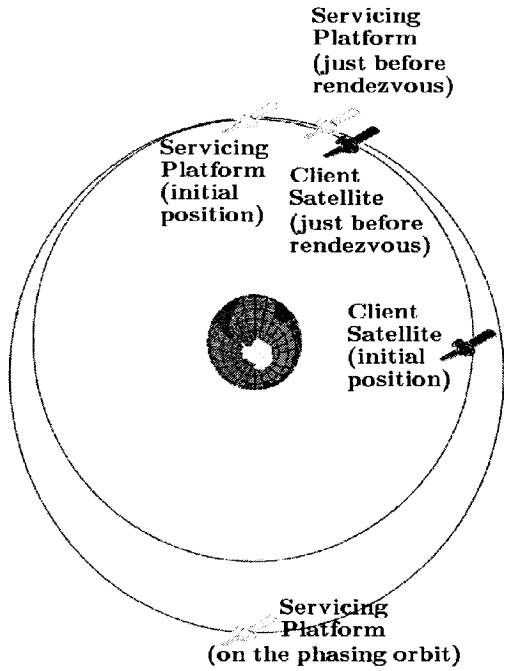

Figure 3.4: Phasing Orbits

the orbits of the servicing platform and the client satellite have different semi-major axes. When the inclinations are different as well, Hohmann transfer with a combined plane change can be used so that both the semi-major axis and the inclination are adjusted during the same set of maneuvers.

Even though the semi-major axes and the planes of two spacecraft are the same, depending on their relative positions (angular phase) these spacecraft can still be far away from each other. In this case, in order to equalize the angular phase, phasing orbit maneuvers are used. In this type of maneuver, thrust is applied in the velocity direction in order to change the semi-major axis of the servicing platform slightly, causing the *orbital periods* between the orbits to change as well (orbital period refers to the amount of time it takes for a satellite to complete a full

revolution around its orbit). Then, when the relative motion between the servicing platform and the client satellite reaches the desired geometry, thrust is applied opposite to the velocity direction so as to return the servicing platform to the original semi-major axis, and complete the rendezvous (Figure 3.4). In general, a larger thrust or larger ΔV leads to shorter transfer times. This is due to the fact that with a small ΔV maneuver, several phasing orbits are required for the rendezvous, while with more energy-intensive maneuvers as few as one phasing orbit may be sufficient.

Figure 3.5 demonstrates the effect of the number of phasing orbits on ΔV and time during OOS operations. As indicated in this graph, ΔV decreases asymptotically with the number of phasing orbits while the duration increases linearly. Therefore, when the number of phasing orbits is less than five, there seems to be a good potential for saving ΔV without increasing the duration of the mission significantly. However, after this point, increasing the number of phasing orbits does not result in sizeable ΔV savings while it increases the duration of the mission. In fact, as the number of phasing orbits increases, ΔV will level off and the transfer time will tend to infinity.

Phasing orbits are used for transfers within the same orbital belt (e.g., GEO to GEO), whereas Hohmann transfer is used for transfers between different types of orbital belts (e.g. LEO to GEO).

The final phase of a rendezvous in space includes proximity operations where smaller burns are used to adjust the orbit of the servicing platform and to enable a smooth capture of the client satellite. In the model, proximity operations are

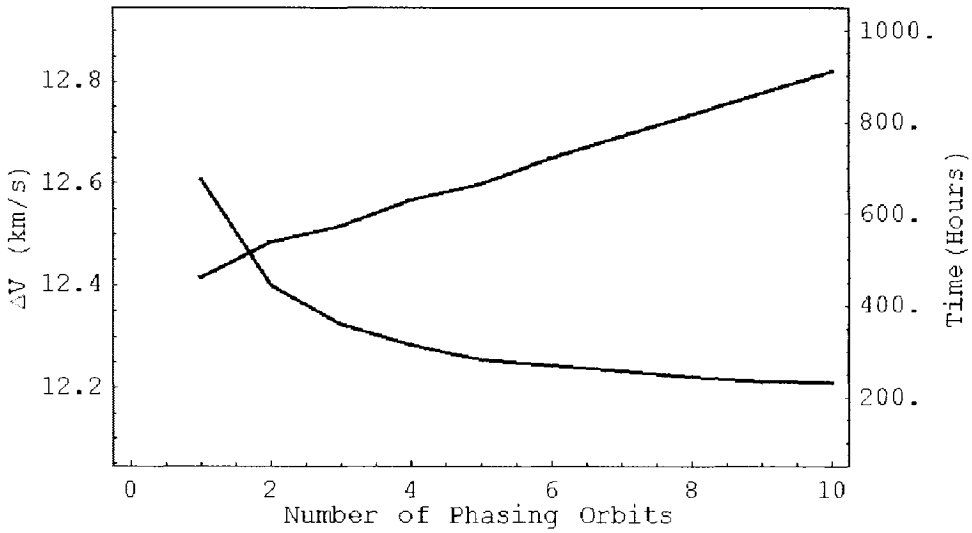

Figure 3.5: The impact of the number of phasing orbits on ΔV and time

approximated by constant ΔV and time values, and these values can be changed by the user according to the specifications of the servicing platform and the client satellites involved.

If a satellite is on an elliptical orbit, its velocity changes significantly during its revolution. It reaches its maximum velocity at the perigee, the point at which it is closest to the Earth. Conversely, it reaches its minimum velocity at the furthest point from Earth (apogee). Despite this change in velocity, computing the orbital period of a satellite is straightforward once we know the classical orbital elements (see Appendix A.1). However, since orbital maneuvers may involve change of orbit type and multiple transfer orbits, and since we can only initiate certain transfers once a satellite reaches a specific point around its orbit, travel times in space cannot be described by a simple, closed-form function. Actually, they depend on the initial and destination orbits, the current location of a satellite on its orbit and the

amount of energy used during a maneuver. Once a satellite is on an orbit, the time it takes to complete one revolution is well defined, i.e., its velocity cannot be changed at will.

3.3.2 OOS Algorithms

In order to solve the OOS problem, two methods are presented: an exhaustive algorithm and a tabu search algorithm. Note that the small size of the instances makes possible the use of exhaustive search and therefore, the number of client satellites, n, is the main factor for selecting the method most suitable for a given problem. Depending on the preferences of the user, the objective function is set to minimize either the total ΔV, or the total amount of time that is required to complete a servicing tour. Alternatively, a bi-criteria objective function can be used in order to analyze the tradeoff between ΔV and time.

For both exhaustive and tabu search algorithms, the choice of the objective function affects the number of phasing orbits that are used during the orbital maneuvers. Based on the insight from Figure 3.5, when ΔV is minimized, the number of phasing orbits is maximized (in order to decrease the total amount of energy). Conversely, when time is minimized, the number of phasing orbits is minimized as well. In the bi-criteria case, each time a particular sequence is evaluated, a greedy algorithm is employed to check the impact of a given number of phasing orbits, and for each transfer, the best number of phasing orbits is chosen.

Modifying the number of phasing orbits is one of the most computationally intensive operations in both exhaustive and tabu search algorithms. Therefore

only three possible values were considered for the number of phasing orbits: 1, 5, and 7 (these values were selected based on Figure 3.5). When the number of phasing orbits exceeds 7, the amount of ΔV gains levels off. For this reason, the area of interest lies within the 1 to 7 phasing orbits range.

A subroutine is nested in both the tabu search and exhaustive search algorithms. For a given sequence, this subroutine calculates the total amount of ΔV, time and fuel that are needed to perform the orbital transfers. In order to make these calculations efficiently, certain assumptions were made. First, all orbits that were investigated had to be circular and all orbits with eccentricity higher than 0.1 were discarded. This assumption holds for most of the satellites that are currently orbiting around the Earth.

Secondly, small differences between the classical orbital elements of initial and destination orbits were ignored. Specifically, if the semi-major axis of the initial orbit is within 10% of the semi-major axis of the destination orbit, then these two orbits are treated as having the same semi-major axis. Recall that phasing orbit maneuvers are used instead of Hohmann transfer for orbital transfers when the semi-major axes of the initial and destination orbits are similar. This simplification has a negligible impact on the overall results.

Finally, in phasing orbit maneuvers, if the given number of phasing orbits (1, 5 or 7) is not sufficient to perform the maneuver, the required (theoretical minimum) number of phasing orbits is enforced by the algorithm. This minimum number of required phasing orbits is calculated using orbital mechanics equations.

Table 3.3 below outlines the orbital transfer subroutine.

Step 1	Initialization
	Get $s_{current}$ from the search algorithm, $s[0] = depot$, $s[n] = last\_client$ Define $origin\_orbit = s[0]$, $destination\_orbit = s[1]$.
Step 2	Hohmann Transfers
	Read $origin\_orbit$ and $destination\_orbit$ check for circularity, discard non-circular orbits. If semi-major axes are similar, calculate ΔV, time and fuel using a transfer with phasing orbits. If semi-major axes are different, calculate ΔV, time and fuel using a Hohmann transfer (which also works for orbits with different inclinations).
Step 3	Calculate total ΔV, time and fuel for the whole sequence
	Update $origin\_orbit$ and $destination\_orbit$ for next transfer. Repeat Step 2 until all servicing transfers are complete.

Table 3.3: Flow of the subroutine for orbital transfers

Exhaustive Search Algorithm

The exhaustive search algorithm was developed in order to solve smaller problem instances ($n \leq 10$) to optimality. It was also used to benchmark the performance of the tabu search algorithm for smaller instances, and to fine-tune it for larger problem instances ($n > 10$). As with any NP-Complete problem, exhaustive search tends to become impractical beyond a certain instance size threshold, due to the enormous amount of computational time required. For this reason, tabu search algorithm emerges as the viable option for larger problem instances.

The exhaustive search algorithm generates all possible permutations for a given servicing mission, such that each mission starts and ends at the orbital depot. Then, the values of ΔV, time and fuel that are required for each permutation are calculated using the subroutine for orbital transfers. During the search, we keep a record of the best solution found since the beginning, including the best values for ΔV, time and fuel consumption.

Tabu Search Algorithm

A solution s is essentially a sequence of client satellites. However, given the nature of orbital transfers and time-dependency, this definition has to be broadened. In addition to this sequence, a solution is also composed of specific information regarding each transfer from one satellite orbit to the next. This information contains the number of transit orbits (if applicable), and also the waiting time. Indeed, due to the time-dependent nature of the problem and the dynamic problem environment, in some cases lower energy transfers can be achieved if the servicing platform waits in a "parked" orbit before initiating the next transfer. Thus, the solution s has to capture not only the sequence, but also the number of transit orbits used for each transfer, and also the amount of time the servicing platform waits at a given orbit (including the wait at the depot at the beginning of the servicing tour).

A basic tabu search algorithm was implemented to tackle larger problem instances. At each iteration, the algorithm operates on the current solution s. The client satellites are coded from 0 to n, 0 denoting the orbital depot. Therefore, a typical solution of the OOS problem is a sequence that starts and ends with 0, with the rest of the satellite codes in between, e.g.$\{0, 3, 2, 4, 1, 0\}$. This particular sequence corresponds to a total of five orbital transfers: from 0 to 3, from 3 to 2, etc. Waiting time and transit orbit information are encoded for every such transfer. For instance we could have {(2,7), (1,0), (0,1), (0,5), (0, 7)}, meaning that for the first transfer (from 0 to 3), after waiting for 2 time units, the transfer is performed using 7 transit orbits, etc.

The neighborhood of s, $N(s)$, is defined as all other solutions that can be reached from s by an elementary move, i.e., by swapping two satellites of the sequence or by a modification of the waiting time or the number of transit orbits. For instance, sequence $\{0, 3, 2, 4, 1, 0\}$ has the following 4 sequences in its neighborhood:

$\{0, 1, 2, 4, 3, 0\}$ by swapping 3 with 1

$\{0, 2, 3, 4, 1, 0\}$ by swapping 2 with 3

$\{0, 3, 4, 2, 1, 0\}$ by swapping 4 with 2

$\{0, 3, 2, 1, 4, 0\}$ by swapping 1 with 4

Unlike conventional applications of tabu search, implementing an elementary move involves intensive calculations since everything that follows the point where the modification is made must be reevaluated. For instance, if an elementary move involves the first two satellites of a solution s, keeping all the rest the same, we still need to recalculate ΔV and the time for the whole of s.

A tabu list is used in order to avoid cycling between successive iterations with a fixed tabu tenure of size 15. This value was found by performing sensitivity analysis through several tests. An aspiration criterion is used to override the tabu status of a move: if a move yields the best solution found so far, it is implemented even if the move still has tabu status. In *Orbital Optimizer*, the user can specify the number of tabu search iterations. When all the iterations are completed, the algorithm stops and the best results are displayed.

Table 3.4 outlines the tabu search algorithm:

Step 1	Initialization
	Construct the initial solution, $s_{initial}$, set $s_{current} = s_{initial}$ Assign current best known solution, $s_{best} = s_{current}$. Define $best\_cost = c(s_{best})$, $best\_cost$ can be ΔV, time or a combination of both.
Step 2	Evaluation
(A)	Generate moves to be evaluated within $N(s)$, pick $s_{next} \in N(s)$, $s_{current} = s_{next}$.
(B)	Evaluate $s_{current}$ based on the sub-algorithm for orbital transfers and the tabu status of the move, if it is the best solution in $N(s)$, then $s^* = s_{current}$ Pick next move to be evaluated, $s_{next} \in N(s)$ and repeat Step 2 (B).
Step 3	Update and Termination
	If a new best solution is found, $s_{best} = s^*$ and $best\_cost = c(s_{best})$. Repeat Step 2 until all tabu search iterations are completed.

Table 3.4: Flow of the tabu search algorithm

Spacecraft and Mission Characteristics

In order to make *Orbital Optimizer* more flexible, additional features were added to the software application. These features enable the user to define various aspects of the servicing missions such as the mass of the servicing platform, the amount of time needed for an average servicing mission after the rendezvous and the amounts of ΔV and time that are required for proximity operations prior to the capture. Based on these characteristics, *Orbital Optimizer* can be used for many systems engineering analyses, such as assessing the total fuel consumption and the impact of different propulsion systems on the feasibility of OOS operations. Fuel itself was not included in the objective function as one of the variables to optimize, since it can be calculated as a function of ΔV and other model variables.

3.4 Results

Orbital Optimizer can be used to conduct various types of analyses for planning OOS missions. In this section, three sets of analysis results will be presented to demonstrate some of these capabilities. These are the results for minimum ΔV and minimum duration servicing tours; the impact of different propulsion systems on the fuel requirement for OOS operations; and the efficient frontier between ΔV and time.

3.4.1 Constructing Servicing Tours

Table 3.6 summarizes the best results that were found for up to $n = 20$ using both exhaustive and tabu search algorithms. The same parameters of servicing missions (e.g., mass of servicing platform, docking maneuver duration, etc.) were used for all problem instances considered.

Each row in Tables 3.5 and 3.6 represents a servicing tour where the row in question and those above it constitute the clients (except for the Depot row). For instance, Code 3 row refers to the servicing tour in which the servicing platform leaves the depot at LEO, and services three satellites (GPS IIR-02, Anik F1 and Thor 2A) before returning to the depot. The depot location is assumed to be the location of an existing satellite (Globalstar-FM64).

The best values for both ΔV and time are given in Table 3.6 as well as the running times for the algorithm. The running times are based on a 1.73 GHz Pentium M processor[2]. The number of tabu search iterations is increased based on the

[2] As expected, as n increases, the running times increase as well. Note that the rate of increase is not

Code	Satellite Name	Location
0	Depot (Globalstar-FM64)	LEO
1	GPS IIR-02	MEO
2	Anik F1	GEO
3	Thor 2A	GEO
4	Turksat 2A	GEO
5	G-star M039	LEO
6	GPS IIR-04	MEO
7	Astra 2A	GEO
8	Iridium 097	LEO
9	GPS IIR-05	MEO
10	Inmarsat 3F5	GEO
11	Iridium 021	LEO
12	GPS IIR-06	MEO
13	HotBird 4	GEO
14	G-star M059	LEO
15	SatMex05	GEO
16	Sirius 3	GEO
17	GPS IIR-07	MEO
18	Iridium 098	LEO
19	Eutelsat 3F2	GEO
20	Intelsat 806	GEO

Table 3.5: Selected Satellites and Their Locations (listed in arbitrary order)

instance sizes. Up to Code 10, 2000 iterations were used. Between Codes 11 and 15, 3000 iterations were used. For Codes 16 to 20, the number of iterations was increased to 3500. The initial solution for the tabu search algorithm is based on the increasing order of the satellite codes (e.g., for Code 4 instance, the initial solution is {0,1,2,3,4,0}). Note that, contrary to intuition, a greedy algorithm that would visit all the satellites in a given orbital belt before moving to another belt does not necessarily gives good solutions since other orbital parameters such as the inclination can have a more significant impact on the values of ΔV and time. Optimal results obtained by tabu search are indicated in bold typeface (see Table 3.6). The corresponding sequences of client satellites can be found in Appendix A.2.

constant (for both exhaustive and tabu search algorithms). One possible explanation for this is the number of calculations needed for a given problem instance. For example, for sequences with no inclination change, certain calculation steps can be skipped and therefore the corresponding running times are lower.

Code	Exhaustive Search				Tabu Search			
	ΔV (km/s)	Running Time (sec.)	Servicing Time (hours)	Running Time (sec.)	ΔV (km/s)	Running Time (sec.)	Servicing Time (hours)	Running Time (sec.)
0	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
1	6.82	<1	123.63	<1	**6.82**	8	**123.63**	9
2	11.50	<1	190.68	<1	**11.50**	16	**190.68**	16
3	11.76	<1	285.68	<1	**11.76**	31	**285.68**	29
4	11.98	<1	369.83	<1	**11.98**	55	**369.83**	51
5	12.23	1	425.29	1	**12.23**	85	**425.29**	76
6	13.43	6	481.85	6	13.47	121	509.00	110
7	13.65	61	563.99	53	13.65	185	583.94	151
8	19.10	570	613.99	497	**19.10**	247	**613.99**	199
9	20.34	5912	674.46	5020	**20.34**	368	688.61	262
10	20.54	71604	765.6	60803	20.59	816	770.14	510
11	-	-	-	-	25.46	1003	836.65	627
12	-	-	-	-	31.57	1234	901.63	781
13	-	-	-	-	32.40	1455	978.38	991
14	-	-	-	-	37.23	1602	1042.77	1151
15	-	-	-	-	33.87	2212	1115.99	1557
16	-	-	-	-	47.12	2577	1200.66	1845
17	-	-	-	-	48.93	2863	1269.35	2142
18	-	-	-	-	54.02	3372	1286.45	2441
19	-	-	-	-	53.94	4205	1382.49	2814
20	-	-	-	-	58.84	4632	1472.59	3222

Table 3.6: Best servicing tours

Comparing the results obtained by the exhaustive vs. tabu search algorithms can give an indication about the performance of the tabu search algorithm. For Codes 9 and 10, results obtained by tabu search for the minimum ΔV problem are 20.34 and 20.59 km/s, respectively. For Code 9, the same result as the exhaustive algorithm is obtained, while for Code 10, the tabu search result is within 0.3% of the result obtained using the exhaustive method.

In order to put in context the values of ΔV that are shown in Table 3.6, consider the following example. Sending a spacecraft all the way to Mars from a geostationary transfer orbit "only" requires 5 km/s of ΔV, while servicing satellites 1 and 2 from 0 and back requires 11.5 km/s of ΔV. Clearly, the amount of ΔV required for OOS operations can be very demanding depending on the orbital parameters of the satellites.

Servicing more than a handful of satellites in a single servicing tour is beyond the reach of current spacecraft systems when the satellites are located in very different orbits. However, when we consider servicing satellites that are in similar orbits, technical feasibility seems to be more promising. For instance, with a depot located in GEO (with the same orbital parameters as the satellite Astra 2A), servicing the first three GEO satellites in Table 3.6 (Anik F1, Thor 2A and Turksat 2A) requires only 0.98 km/s of ΔV. Thus, the technical feasibility of future OOS missions depends on many factors, including the number, capacity and location of the orbital depots, the number and location of client satellites, and the onboard propulsion systems for performing the maneuvers.

Code	Satellite Name	Location	ΔV reduction (km/s)	Waiting Time (days)
0	Depot	LEO	n/a	n/a
1	GPS IIR-02	MEO	0	{0,0}
2	Anik F1	GEO	-0.002	{1,1,0}
3	Thor 2A	GEO	-0.001	{1,0,0,0}
4	Turksat 2A	GEO	-0.004	{1,2,2,0,2}

Table 3.7: Impact of waiting time on the solutions

Table 3.7 demonstrates how the solutions presented in Table 3.6 change when waiting times are taken into account. In Table 3.6, it is assumed that the servicing platform proceeds from one client satellite to the next without delay, whereas for Table 3.7 we allowed the servicing platform to wait for 0, 1 or 2 days before moving to the next client satellite. The last column of Table 3.7 shows the amount of time the servicing platform waits before starting the next transfer. Although theoretically there can be a gain due to waiting, the results in Table 3.7 show that the reductions in energy requirements may sometimes be negligible.

3.4.2 Impact of propulsion system

One of the systems engineering analyses that can be conducted by using *Orbital Optimizer* is the impact of various propulsion systems on mission parameters. Figure 3.6 demonstrates the impact of different propulsion systems on the total amount of fuel that is needed to perform the orbital transfers in a servicing tour ($n = 5$). Specific impulse, I_{sp}, (X-axis) measures the performance of the propulsion system. The higher the specific impulse, the more efficient the propulsion system. I_{sp} for conventional chemical propulsion systems is around the range of

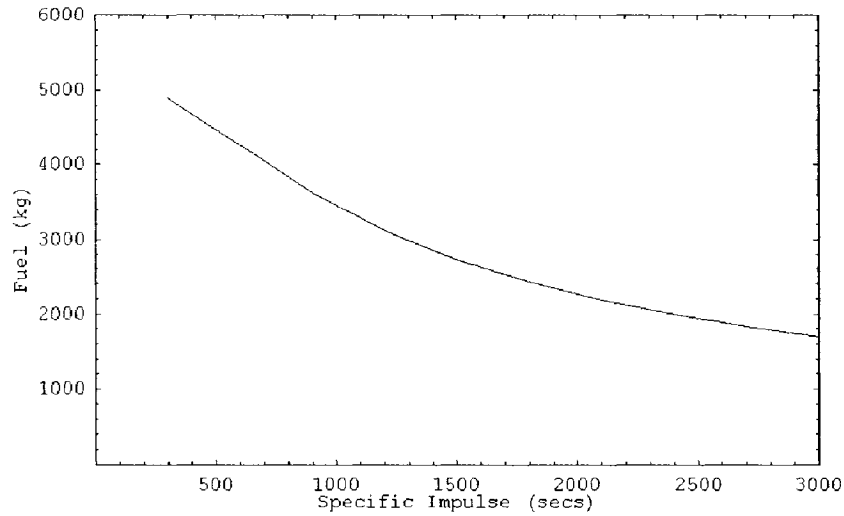

Figure 3.6: Impact of advanced propulsion technologies (n = 5)

200-350 seconds. Advanced technologies, such as ionic propulsion, provide much higher levels of performance, with I_{sp} levels up to 6000 seconds.

As can be seen on the graph, advanced propulsion technologies can decrease the total mass of the servicing platform considerably, since less fuel is needed for orbital transfers. However, since most engines with high specific impulse cannot provide high levels of thrust, the number of phasing orbits will increase during transfers. As a consequence the total mission duration will increase as well.

3.4.3 ΔV / Time Efficient Frontier

Figure 3.7 plots ΔV against time for each servicing sequence when $n = 5$. It is interesting to note a "clustering effect": there seems to be four main groups of solutions, and within each group, dominating solutions can be found that are closest to the efficient frontier. For instance, for a time-conscious decision-maker,

of all the points that are on the top cluster, the one closest to the efficient frontier would be the most attractive, since for a given level of energy, this point represents the minimum time solution. One plausible explanation for this phenomenon may be the impact of phasing orbits. In the sequences where the servicing platform can visit all satellites with the same semi-major axis and inclination successively (e.g., the sequence 0,5,1,2,3,4,0}, phasing orbits can be used for two maneuvers out of a total of six for the whole sequence (i.e., 2 to 3 and 3 to 4). This reduces the required ΔV significantly, but increases the servicing time. At the other extreme, when no phasing maneuvers can be used (e.g., the sequence 0,2,1,3,5,4,0), the required ΔV is much higher (in the 25 km/s range) but servicing times are lower.

Analyzing such solutions in detail can help the decision-makers to quickly converge on efficient solutions that represent a good compromise between ΔV and time. Larger problem sizes also exhibit the clustering effect, although the number of clusters tend to increase as well. The efficient frontier can also be useful to quickly evaluate mission feasibility. A decision-maker with a strict ΔV budget can use the efficient frontier to find the shortest amount of time possible to service all satellites.

3.5 Conclusion

On-orbit servicing is only one of the many space-based problems for which Operations Research can play an important role. Although the physical environment of these problems is fundamentally different from that of Earth-based problems,

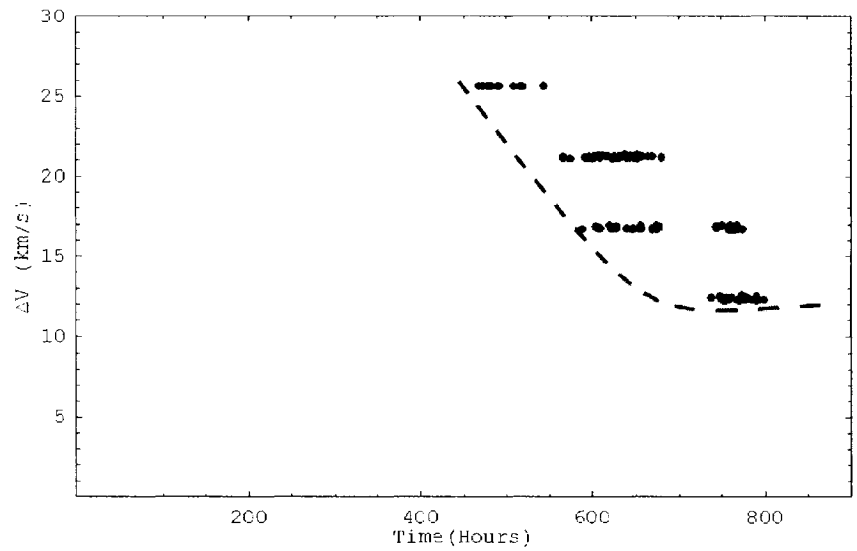

Figure 3.7: ΔV / Time Efficient Frontier

models and methods that have been developed for existing transportation problems provide a strong foundation for space-based problems. In addition to providing practical benefits to space mission designers, these "new generation" problems can also stimulate research in time-dependent and moving-target TSPs and VRPs.

Given the current level of technology and market readiness, it will be a very long time before servicing tours involving a large number of client satellites can be realized. Therefore, from an Operations Research perspective, exhaustive methods will continue to be useful for space applications for quite some time. However, as discussed in the previous sections, determining the optimal servicing tours is not only about the sequence of satellites to be visited, but also about the exact times of initiating transfers and the optimal number of transit orbits. For this reason, as space missions get more complex, there can be an increased role for metaheuristics such as Tabu Search.

One of the obvious extensions of the work presented here is studying the VRP variant of the OOS problem. Analyzing the case of more than one servicing platform can enable mission designers to evaluate the marginal value of deploying multiple servicers. Other questions that remain to be investigated are optimal depot locations and moving-depot TSPs and VRPs (since it is technically possible to relocate an orbital depot). Moving beyond the OOS problem, there are a wide variety of space-based problems that can be studied using Operations Research methods, such as interplanetary trajectory optimization, spacecraft design tradeoffs, and communication down-/uplink scheduling of orbiters and landers.

Chapter 4
Option Valuation Using the Tabu Search Monte Carlo Method

4.1 Introduction

Numerical methods play an increasingly important role in the analysis and valuation of options. This chapter presents a new numerical method for the valuation of options with early exercise features: Tabu Search Monte Carlo (TSMC) method. By integrating numerical simulation and optimization procedures, TSMC method approximates the early exercise boundary (or the early exercise region) of such options, and calculates the option value.

For options on a single underlying asset, two parameters, α and β, define the early exercise boundary. For options on multiple underlying assets, a third parameter, θ, is needed to capture the dynamics of the interactions between the underlying assets.

The problem of finding the set of these parameters that maximizes the option value is modeled as a combinatorial optimization problem. In contrast to the problem presented in the preceding chapter, this problem is not deterministic. The inherent uncertainty of the underlying assets' price dynamics renders it a stochastic one. For this reason, a simulation-based solution approach is developed as part of the Tabu Search Monte Carlo method.

Although, as discussed in the previous chapters, tabu search cannot guarantee optimality, the TSMC method does provide a high level of precision. As the results presented in Section 4.6 indicate, in almost all cases, computed option values are within 0.5% of the best results reported in the literature.

One of the strengths of the TSMC method is the ability to calculate the value of options that are based on multiple underlying assets. Traditional methods, such as binomial trees and finite differences, are not viable when the number of underlying assets is more than three, since the time complexity of these solution methods increases exponentially as more underlying assets are added. Metaheuristics, such as tabu search and genetic algorithms, are suitable tools for dealing with this "curse of dimensionality", since these heuristics can find close to optimal solutions in a reasonable amount of computing time, even if the problem exhibits more than polynomial complexity. As will be discussed in the next chapter, the ability to analyze multidimensional problems is especially relevant for real options analysis, since many real-life projects have multiple sources of uncertainty.

This chapter is structured as follows: following the discussion of the properties of options with early exercise features and the review of relevant literature, the tabu search formulation of the problem is presented. Then, the flow of Monte Carlo simulation and tabu search iterations are demonstrated. The next section includes the results of the valuation of American options with a single underlying asset, the valuation of max-call options with up to ten underlying assets and the valuation of path-dependent options. Relevant performance metrics of the TSMC algorithm, such as accuracy, bounds and computing times are provided. The chapter con-

cludes with the discussion of potential application areas for TSMC, including real options.

4.2 Properties of Options with Early Exercise Features

As discussed in Section 2.5.1, options are financial instruments that give owners *the right but not the obligation* to exercise their rights to buy or sell financial assets in the future. The exercise price and the lifespan (or maturity) of financial options are known at the time of contract creation.

A *call option* gives the option-holder the right to buy the underlying asset, whereas a *put option* gives the right to sell it. A *European option* can only be exercised at maturity. An *American option* can be exercised at any time before (or at) maturity. There are many different underlying assets for options contracts, such as stocks of companies, stock market indices, and currencies.

Combining different features of options has resulted in many "exotic options", such as compound options (options on options), rainbow (or basket) options (options on two or more risky assets) and Bermudan options (options that can be exercised at only certain points before maturity).

4.2.1 Options Written on a Single Underlying Asset

Since an American option can be exercised at any time before maturity, and since the underlying asset's critical values which trigger early exercise are not known *a*

priori, the valuation of American options is considered a free boundary problem (McKean, 1965; Merton, 1973; Geske and Johnson, 1984)[3].

Given an option with early exercise features, at each time step t (when the option can be exercised) the option holder has to determine whether he or she should exercise the option, or continue holding it. The early exercise boundary is the set of critical values (over the lifetime of the option) that the option holder compares with the realization of the underlying asset's value, V_t. For a put (call) option, if V_t is under (above) the boundary, then the option should be exercised. Therefore, the early exercise boundary partitions the (V, t) space into a continuation region and an adjacent exercise region (this boundary is a time-dependent function of the underlying asset's values that partitions the exercise and hold regions in such a way that the value of the option is maximized (Geske and Johnson, 1984)).

The early exercise boundary is also a core concept for the valuation of real options, and it has been studied extensively in the real options literature (Paddock, Siegel and Smith, 1988; Dixit and Pindyck, 1994).

The mathematical properties of American options are as follows[4]:

The payoff function of an American put option is:

$$\Omega(V, t) = \max(0, D - V_t) \tag{4.1}$$

---

[3] Other options with early exercise features, such as the Bermudan options, also have similar characteristics. In this chapter, valuation formulas and mathematical properties of American options are given. With certain modifications, these formulas and properties apply to most other types of options with early exercise characteristics.

[4] Although the properties of the American put option are given in this section, without loss of generality, similar properties hold for call options.

where V_t is the value of the underlying asset at time t, and D is the exercise price of the option.

If the value of the option $F(V,t)$ is not priced accurately, one can make risk-free (arbitrage) profits. For instance, if $F(V,t) < \max(0, D - V_t)$, one can buy the asset in the market for V_t, buy the option for F and by immediately exercising the option, sell the asset for D. In this case, one can make a risk-free profit of $D - V_t - F$.

This arbitrage argument imposes the following condition that an American put option has to satisfy at every instant:

$$F(V,t) \geqslant \max(0, D - V) \text{ for all } t \in [0,T] \tag{4.2}$$

where T is the date of expiration of the option[5].

Based on the arbitrage argument above, the American option valuation problem can be formulated as an optimal stopping problem. The fair value of an American option is obtained when we maximize the expected exercise value taken over all stopping times. In this case, the fair value of an American put option at time $t = 0$ is[6]:

$$F_0 = \max_{\tau \in [0,T]} E[e^{-r\tau} \max(0, D - V_\tau)] \tag{4.3}$$

where τ is the stopping time for any given realization of V and r is the risk-free discount rate.

Assuming that the underlying asset, V, follows geometric Brownian motion,

---

[5] This condition does not apply to European put options and their values can be below the payoff function at some point before maturity.

[6] In comparison, since a European put option can only be exercised at time T (maturity), its value can be calculated as $F_0 = e^{-rT} E[\max(0, D - V_T)]$.

$$dV = \mu V dt + \sigma V dz \tag{4.4}$$

where μ is the drift parameter, σ is the volatility parameter (the standard deviation of the rate of return on the asset) and dz is the increment of the Wiener process defined as $dz = \epsilon\sqrt{dt}$ with $\epsilon \sim N(0,1)$.

Although the valuation of options with early exercise features is different from the valuation of European options, the valuation equations are closely related. As demonstrated by Black and Scholes (1973) and Merton (1973), a European option $F(V,t)$ must satisfy the following partial differential equation:

$$\frac{\partial F}{\partial t} + (r-\delta)V\frac{\partial F}{\partial V} + \frac{1}{2}\sigma^2 V^2 \frac{\partial^2 F}{\partial V^2} = rF \tag{4.5}$$

where r is the risk-free interest rate and δ is the rate of return shortfall (also known as the constant dividend yield). In the case of real options on natural resources, δ can be used to represent the convenience yield (the benefit or premium obtained by holding a physical-asset).

However, in the case of American options, this equality is transformed into an inequality since the option can be exercised before maturity, and the arbitrage argument used in the European option case no longer yields a unique solution (Wilmott, Howison, and Dewynne, 1995).

$$\frac{\partial F}{\partial t} + (r-\delta)V\frac{\partial F}{\partial V} + \frac{1}{2}\sigma^2 V^2 \frac{\partial^2 F}{\partial V^2} \leq rF \tag{4.6}$$

When it is optimal to wait, the equality in equation 4.6 holds (i.e., the Black-Scholes equation in 4.5) and constraint 4.2 must be satisfied. If this is not the case, then it is optimal to exercise the option, and only the inequality in equation 4.6 holds and the equality in constraint 4.2 is satisfied (thus giving the payoff of the option) (Wilmott, Howison, and Dewynne, 1995).

The boundary conditions that complete the characterization of an American put option are:

$$F(V,t) \rightarrow 0 \text{ as } V \rightarrow \infty \tag{4.7}$$

$$F(V^*,t) = D - V^* \tag{4.8}$$

$$\frac{\partial F}{\partial V}(V^*,t) = -1 \tag{4.9}$$

$$F(V,T) = \max(0, D - V_T) \tag{4.10}$$

Condition (4.7) states that the put option has no value when V is infinite. Condition (4.8) is the value-matching condition that gives the payoff of the option when V_t reaches V_t^* (the critical value which triggers early exercise). Condition (4.9) is the smooth-pasting condition ensuring that $F(V)$ is continuous and smooth. It states that just at the point when V_t reaches V_t^*, the put value $F(V)$ would decrease one dollar for a one dollar increase in the underlying asset's value, V_t (Geske and Johnson, 1984).

Condition (4.10) states that at expiration, the put option will be exercised if $D > V_T$. The early exercise boundary approaches D towards the end of the option's maturity since the implicit option value decreases with time.

During the lifetime of the option, when V_t hits the early exercise boundary $V_t^*(t \in [0,T])$ for the first time, early exercise is triggered. If for all $t \in [0,T]$, $V_t > V_t^*$, then the put option expires worthless (at maturity, $V_T^* = D$) .

Similarly, the boundary conditions for solving an American call option are:

$$F(0,t) = 0 \tag{4.11}$$

$$F(V^*,t) = V^* - D \tag{4.12}$$

$$\frac{\partial F}{\partial V}(V^*,t) = 1 \tag{4.13}$$

$$F(V,T) = \max(0, V_T - D) \tag{4.14}$$

4.2.2 Multidimensional Options

For multidimensional options, the early exercise boundary is replaced by a set of early exercise regions which partition the state space over time (Broadie and Detemple (1997) and Kay, Davison and Rasmussen (2003), see Figure 4.1 and Table 4.3 for graphical representations of the early exercise regions). The early exercise region for a particular underlying asset is a function of both time and the values of the other underlying assets.

As in the standard (or unidimensional) case, multidimensional options can have European features or they can be exercised at some point before maturity. American multidimensional options can be exercised at any point before (and at) maturity, while Bermudan multidimensional options have a limited number of early exercise opportunities.

The payoff structure of multidimensional options can also differ considerably. One of the most prominent multidimensional options is the max-option. Broadie and Detemple (1997) refer to such options as the "prototypical contract", since most other multidimensional options can be derived from the max-option.

A max-option is an option written on the maximum of multiple underlying assets. An American max-option can be exercised at any point before (or at) maturity. The payoff of this option at exercise time t is:

$$F_t = \max(\max(V_t^1, ...V_t^n) - D, 0) \tag{4.15}$$

A dual strike option is an option written on multiple underlying assets with different exercise prices. For the American dual strike option, the payoff at exercise time t is:

$$F_t = \max(\max(V_t^1 - D^1, ...V_t^n - D^n), 0) \tag{4.16}$$

A spread option is an option written on the difference (or spread) between two or more underlying assets. For an American spread option written on two assets, the payoff at time t is:

$$F_t = \max(V_t^2 - V_t^1 - D, 0) \tag{4.17}$$

In the special case of $D = 0$, a spread option becomes an exchange option, giving the option holder the right to exchange one asset for the other. Exchange options were first analyzed by Margrabe (1978).

Yet another type of multidimensional option is an option written on the powers of the product of two or more assets. For example, if V^1 and V^2 are the two underlying assets, then the payoff of this option (in the case of an American option) is:

$$F_t = \max((V_t^1 V_t^2)^a - D, 0) \tag{4.18}$$

For values of $a > 0$, the option becomes a "power-product" option. The two special cases of this option are product options ($a = 1$) and options on the geometric average of assets ($a = 0.5$).

In addition to these options, certain other types of options can be modeled as multidimensional options, although they are technically written on a single variable. One class of such options is path-dependent options. The underlying asset of a path-dependent option is the average value of an asset over a specified time period. One way of solving this valuation problem is to structure the problem such that the exercise decision is based on two state variables for each time t: the current value of the asset and its average value for a given time window. Thus, this type of option can be modeled as an multidimensional option.

4.2.3 On the Generality of the TSMC Method

The TSMC framework presented in this chapter is designed principally for max-options. Therefore certain modifications are needed before it can be implemented to other types of multidimensional options. In certain cases, this modification is very simple: for instance, for the valuation of Amerasian (path-dependent) options,

the average of the past values of the asset and the current value of the asset are modeled as two distinct, but related, variables (just like a max-call option written on two assets). For other cases, this modification can be more complex. However the basic premise of the TSMC framework is still valid: deriving policies for early exercise in the presence of inter-asset interactions by approximating the early exercise regions and then valuing the option via Monte Carlo simulation.

4.3 Numerical Methods in Option Valuation

Exact analytical solutions can be obtained for the American option valuation problem only for some special cases, such as the case of a perpetual American option (in which $\frac{\partial F}{\partial t}$ in equation 4.6 disappears). The non-linear nature of the problem and the existence of the free boundary rule out closed-form, exact analytical solutions. Therefore various numerical methods were developed for this problem.

Following the introduction of binomial trees by Cox, Ross and Rubinstein (1979), and finite differences by Brennan and Schwartz (1977), lattice-based methods have evolved into proven tools for American option valuation. These numerical methods are based on the recursive procedures used in dynamic programming, and thus they can be seen as exhaustive methods which trace the value of an option backwards. Starting from maturity, the optimal decision at every step in time is determined until time zero is reached, and the option value is calculated. Lattice-based methods have also been extended to tackle two and three state variables (Boyle 1988). However, for problems of higher dimension, the re-

cursive and exhaustive nature of these methods render them impractical since the memory requirement increases exponentially as the option valuation problem dimension increases. Also, these methods work well when the processes governing the behavior of the underlying assets over time are fairly standard, such as the geometric Brownian motion. Their extension to other types of processes, such as jump-diffusion models, is not trivial.

In addition to lattice-based methods, analytical approximations are also used in option valuation. There are a very wide variety of analytical approaches. Geske and Johnson (1984) present an analytical method that considers an American option as an infinite sequence of options on options (or compound options). They derive an exact formula (in the limit) for the valuation of an American options subject to its early exercise boundary. Although the continuous time solution of this formula would necessitate an infinite series of integrals, Geske and Johnson (1984) model each exercise decision as a discrete event and they extrapolate the actual option price from a small number of exercise points. Another approach by Barone-Adesi and Whaley (1987) is based on using a quadratic approximation for computing the value of the early exercise premium. Other analytical methods, such as those of Kim (1990) and Carr, Jarrow and Myneni (1992) are based on decomposing the American option into two parts: a European option and the early exercise premium. They then compute the value of the premium using various integration methods. Ju (1998) approximates the early exercise boundary as a multipiece exponential function and provides closed form formulae for evaluating this function.

Although both the lattice-based stream and the analytical approximation stream provide a rich set of alternatives for analyzing American options of lower dimensions, these methods are not efficient for multidimensional problems and for processes which are more "exotic" than the geometric Brownian motion. Numerical simulation, on the other hand, provides a very flexible framework for analyzing different types of options problems.

Introduced by the pioneering work of Boyle (1977) for option valuation, Monte Carlo simulation has become a very popular method. Initially, Monte Carlo simulation was used for the valuation of European options only, however recent advances in theory and practice have opened the doors for many other types of applications, including the valuation of options with early exercise features such as American and Bermudan options.

Generally, simulation-based methods approximate the value of an American option by dividing the lifespan of an option into a finite number of exercise points (e.g., an American option is modeled as a Bermudan option). This approximation enables computing the value function or the early exercise boundary of the option using Monte Carlo simulation. Methods based on computing the value function include those described in Longstaff and Schwartz (2001) and Barraquand and Martineau (1995). Those that are based on approximating the early exercise boundary include Ibanez and Zapatero (2001), Andersen (2000), Garcia (2003) and the evolutionary real options method of Dias (2001). Modeling American options with a small number of early exercise opportunities (i.e., Bermudan options) and Monte Carlo simulation are common features in these efforts.

Ibanez and Zapatero (2001), compute some of the fixed points of the boundary and then use regression and interpolation to approximate the boundary (by "fixed points", they refer to the critical values V^* at which the exercise value of the option is equal to its continuation value). Their method is applicable to both options written on a single underlying asset and multidimensional options. They provide results for various max-call options with up to 5 underlying assets.

Basso, Nardon and Pianca (2002) use a two-step procedure. In the first step, the early exercise boundary is estimated using an improved binomial approximation, and in the second step, early exercise features are studied using Monte Carlo simulation. Although their procedure provides accurate results, it is still subject to the curse of dimensionality, since they are using the binomial method to approximate the early exercise boundary.

Garcia (2003) presents a parametric method that is effective for both standard American options and max-options. The main idea behind his algorithm is to represent the early exercise rules by a finite number of parameters, and then using Monte Carlo simulation to approximate the value of options for a given set of parameters. Therefore, finding the set of parameters that maximizes the option value is the main objective of Garcia's algorithm. In order to represent the boundary of a standard American option, he uses two parameters: the boundary values at points 0 and $0.7T$. Since at time T the boundary value is equal to D, this gives him three points that lie on the boundary. He then uses a cubic Hermite polynomial to interpolate between these points and approximates the whole boundary. In

essence, this approximation enables him to measure how deep in the money the option should be for early exercise to be optimal.

In the case of max-options, Garcia (2003) adds another set of parameters to approximate the early exercise regions. Given a vector of underlying asset prices $\mathbf{v}_t = \langle V_t^1, V_t^2, ..., V_t^n \rangle$, this additional set of parameters measures how far the value of the highest element of $\mathbf{v}_t$ should be from the second order statistic of $\mathbf{v}_t$ (second lowest element) for early exercise to be optimal. For a two-asset max-option, he uses the following representations for the optimal stopping times: $\tau = \inf\{t : \max(\max(V_t^1, V_t^2) - D, 0) > \theta_t^1; |V_t^1 - V_t^2| > \theta_t^2\}$, where θ_t^1 measures how deep in the money the option should be for early exercise, and θ_t^2 measures the minimum acceptable spread between the two assets for early exercise[7]. He observes that the smaller the spread, the less likely it will be to exercise the option. That is, even though the highest element in $\mathbf{v}_t$ may be very deep in the money, if the highest element and the second order statistic are within θ_t^2 of each other, the option holder should wait and not exercise immediately. This rule can be interpreted as follows: if the highest element of $\mathbf{v}_t$ and the second order statistic are within θ_t^2 of each other, there is a chance that the second order statistic can exceed the highest element (at time t) at some point between $T - t$. Even with the impact of discounting, this can yield a higher option value than exercising at time t.

Dias (2001) proposes an alternative way to approximate the early exercise boundary. He uses a logarithm function and two free points to represent the bound-

[7] In the TSMC method, a different set of parameters are used to measure how deep in the money an option should be for early exercise. However, the θ_t parameter of the TSMC method has a very similar function to the θ_t^2 parameter of Garcia (2003).

ary parametrically. Free points are used to approximate the boundary when the option is close to its maturity, since the logarithm function diverges from the actual value of the boundary. The "fitness" of the boundary (i.e., the quality of the approximation) is evaluated by genetic algorithms and Monte Carlo simulation. By using a forward Monte Carlo procedure and an optimization heuristic, Dias (2001) provides a new way to approximate the early exercise boundary. Dias (2001) demonstrates his method in the context of a real options problem where the decision-maker has to select one of the investment alternatives in presence of both market and technical uncertainties. Although Dias (2001) does not provide an extension of his method for a max-call option, he nevertheless tackles a problem based on multiple sources of uncertainties (i.e., market and technical).

One significant strength of Ibanez and Zapatero's (2001) and Garcia's (2003) approach is the ability to provide upper and lower bounds for option values, and not just point estimates. This feature is also implemented in the TSMC framework.

It is important to note that, for options written on a single underlying asset, the critical values on the early exercise boundary are independent of the current (or initial) value of the underlying asset (Geske and Johnson, 1984; Garcia, 2003). This property is particularly useful for the TSMC method, since once the early exercise boundary is found for a particular option, the same boundary can be used to calculate the values of options with similar characteristics, but with different initial values. However, this property does not hold in the case of multiple underlying assets, since the relative values of the assets have a direct impact on the exercise decision.

4.4 Computational Complexity and the Curse of Dimensionality

Lattice-based methods are based on the discrete approximation of the dynamic programming principle of Bellmann (1957). This principle, in turn, is based on modeling the underlying uncertainty as a Markov process, where each source of uncertainty constitutes one dimension of the problem at hand. Therefore, within the dynamic programming framework, modeling multiple sources of uncertainty necessitates introducing additional dimensions to the problem. Since each dimension evolves independently over time, the complexity of this system increases exponentially as new dimensions are added. The curse of dimensionality refers to this exponential increase in problem complexity (Bellmann, 1957). Due to this characteristic, dynamic programming based methods are not practical for problems which have more than three sources of uncertainty.

To date, there have been various attempts to develop methods that can bypass this limitation. Barraquand and Martineau (1995) provide what they call "the first capable method" for computing American option prices and exercise strategies in higher dimensional cases. They partition the space of underlying assets into a tractable number of cells, and they approximate the early exercise strategy for these cells. They use a payoff function as the partition, and call this method "stratified state aggregation along the payoff" (SSAP). For the three underlying asset test case, they provide results that are very similar to those of the lattice based methods (while the running time of SSAP is only a fraction of the time needed for lattice based methods, since it increases linearly in n, the number of underlying

assets). Barraquand and Martineau (1995) provide results for up to ten underlying assets with additional experimental data indicating that their method works for up to 400 underlying assets. Although some convergence issues were reported about the SSAP method (Coyle and Yang, 1999), it remains as one of the first methods capable of calculating the value of multidimensional options with early exercise characteristics. Recently, other methods have also emerged as discussed in section 4.3.

4.5 Tabu Search Monte Carlo Algorithm

The TSMC method is based on the approximation of the early exercise boundary in the case of a single underlying asset, and on the approximation of the early exercise region in the case of multiple underlying assets. The two core features of the TSMC method, the logarithm function representing the early exercise boundary, and the θ parameter which captures the interaction among multiple assets were originally proposed by Dias (2001) and Garcia (2003), respectively.

Although the TSMC method follows the pathway of Dias (2001) and Garcia (2003), it has the following unique features. Instead of using free points, the portion of the boundary that is close to maturity is calculated by linear interpolation, and thus the number of parameters to be optimized are reduced. Also, a flexible optimization method, tabu search, is used with varying numbers of Monte Carlo paths to create good solutions faster with smaller number of paths, and then searching

Parameter	Function
	For a standard option with early exercise features (single underlying asset)
α	Moves the boundary up and down along the y-axis. The higher the volatility and the number of assets, the lower (higher) α for a put (call) option.
β	Scales the logarithm function. As β approaches 0, the early exercise boundary flattens.
Boundary curve equation	$V_t^* = D(\alpha + \beta \ln(T-t))$
	For a max-call option with early exercise features (multiple underlying assets)
α	Same α value for each underlying asset if the asset characteristics are the same, if not, a unique α value for each asset
β	Same β value for each underlying asset if the asset characteristics are the same, if not, a unique β value for each asset
θ_t	For max-call options, θ_t delays option exercise based on the relative values of the highest asset price and the rest of the asset prices (it represents the opportunity cost of early exercise)
Boundary region equation	$(\mathbf{v}_t^{\max})^* = (\mathbf{v}_t^{\max})\prime + \max[0, (V_t^i - (1-\theta_t)(\mathbf{v}_t^{\max})\prime)]$ where $(\mathbf{v}_t^{\max})\prime = D(\alpha + \beta \ln(T-t))$

Table 4.1: TSMC Parameters

for better, more accurate solutions around the set of previously found solutions, thus increasing the speed of the optimization procedure.

The Tabu Search Monte Carlo (TSMC) algorithm calculates option values by approximating the early exercise boundary (or the region) in successive rounds of optimization in order to find the set of parameters that yields the maximum option value. These parameters are given in Table 4.1.

For an option with a single underlying asset, the early exercise boundary is approximated by the following function (based on Dias (2001)):

$$V_t^* = D(\alpha + \beta \ln(T-t)) \tag{4.19}$$

where $\alpha \in [0,1]$ and $\beta \in [-0.6, 0]$, for a put option and $\alpha \in [1,2]$ and $\beta \in [0, 0.6]$ for a call option[8]. The values of α and β are not time-dependent, since

[8] These ranges are selected based on the boundary conditions of the options discussed above, as well as numerical experiments. Note that, for a call option the minimum value that V_t^* can achieve is D. This

the evolution of the early exercise boundary over time is approximated using the logarithm function[9].

As stated in Conditions 4.10 and 4.14, at maturity, $V_T^* = D$ (thus the approximation of equation 4.19 does not hold at maturity). However, by using this feature of the early exercise boundary, we can approximate the values of V_{T-1}^* and V_{T-2}^* by interpolating between D and V_{T-3}^* , the last point computed using the logarithm function. Thus, by using only two parameters, α and β the whole boundary can be approximated.

4.5.1 Early Exercise Rules

Options on a Single Underlying Asset

For an option written on a single underlying asset, the optimal values of α and β determine the stopping times which maximize equation 4.3. In turn, these stopping times form an early exercise policy, which can be used to compute the price of the put option on a single underlying asset. The first time the underlying asset's value V_t hits the boundary value V_t^* (as determined by equation 4.19), early exercise is triggered. In order to compute the option value, both the exercise points, and the parameters of the logarithm function are transformed into discrete points in the TSMC method (effectively transforming a continuous problem into a discrete one).

---

minimum value corresponds to the lowest possible values for α and β for a call option (1 and 0, respectively). Conversely, for a put option, these values correspond to the maximum possible value of V_t^* (which is again D).

[9] In comparison, the values of θ_t are time-dependent, as it will be explained in the next section.

Options on Multiple Underlying Assets

In the case of American options written on multiple assets, instead of a single boundary curve, there are multiple *early exercise regions*. Broadie and Detemple (1997) give a formal definition of the early exercise regions, as well as their mathematical properties. The shape of the regions are documented by Kay, Davison and Rasmussen (2003), Ibanez and Zapatero (2001) and Broadie and Detemple (1997).

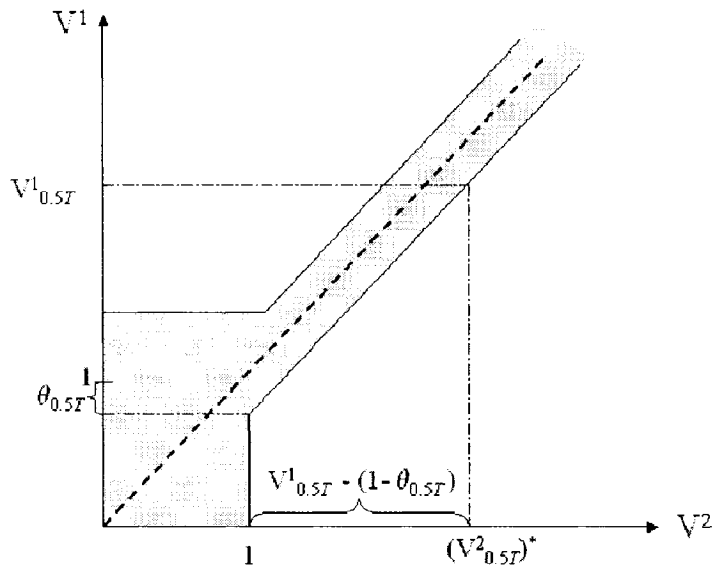

Figure 4.1: V^1V^2space at $t = 0.5T$ (max-option on two assets)

One key observation of Broadie and Detemple (1997) is that, for a max-option on two assets V^1 and V^2, $max(V_t^1, V_t^2)$ is not a sufficient statistic for early exercise decisions (the payoff of this option is given in equation 4.15). In other words, since the option is written on multiple assets, even if $V_t^1 > V_t^2$ there is always a chance

that the value of V^2 may surpass that of V^1 at any point in time in $(t, T]$. Therefore Broadie and Detemple (1997) state that when the two underlying assets prices are equal, it is never optimal to exercise the option (even if the option is deep in the money and dividend rates are substantial). Early exercise rules for American options on multiple assets need to take this "opportunity cost" of early exercise into consideration (otherwise the option will be undervalued).

For a max-option on two (or more) underlying assets, the *relative* values of the assets are very important factors that affect the early exercise decision. Therefore, early exercise is triggered only when an underlying asset's value surpasses its critical value which includes an opportunity cost premium based on the value of the other asset. In other words, the early exercise region of one asset is a function of both time and the value of the other underlying asset, i.e., $(V_t^1)^* = f(V_t^2, t)$.

In the TSMC implementation, this relationship is represented as follows:

$$(V_t^1)^* = (V_t^1)\prime + \max\left[0, (V_t^2 - (1-\theta_t)(V_t^1)\prime)\right] \tag{4.20}$$

where $(V_t^1)\prime = D_1(\alpha_1 + \beta_1 \ln(T-t))$.

Similarly, for V^2:

$$(V_t^2)^* = (V_t^2)\prime + \max\left[0, (V_t^1 - (1-\theta_t)(V_t^2)\prime)\right] \tag{4.21}$$

where $(V_t^2)\prime = D_2(\alpha_2 + \beta_2 \ln(T-t))$.

For a finite time option valuation problem, these linear approximations provide a reasonably good representation of the early exercise regions in the V^1V^2 plane.

In the case of $D_1 = D_2$, this option becomes an American max-call option. If $D_1 \neq D_2$, it becomes an American dual strike option. If the underlying assets' characteristics are the same (i.e., if they both have the same σ and δ values), then $\alpha_1 = \alpha_2 = \alpha$ and $\beta_1 = \beta_2 = \beta$.

Thus, the option should only be exercised if $V_t^1 \geq (V_t^1)^* > V_t^2$ or $V_t^2 \geq (V_t^2)^* > V_t^1$. Graphically, in Figure 4.1, if the (V_t^1, V_t^2) pair is located in the upper left exercise region or in the lower right exercise region, the option should be exercised. If the pair is located inside the shaded region, then the owner of the option should wait.

For max-options with $n > 2$, equation 4.20 can be generalized as follows:

$$(\mathbf{v}_t^{\max})^* = (\mathbf{v}_t^{\max})\prime + \max\left[0, (V_t^i - (1-\theta_t)(\mathbf{v}_t^{\max})\prime)\right] \tag{4.22}$$

where $\mathbf{v}_t$ is a vector containing all underlying asset prices $\mathbf{v}_t = \langle V_t^1, V_t^2, ..., V_t^n \rangle$, when $\mathbf{v}_t$ is sorted from the smallest element to the highest, we get $\mathbf{v}_t = \langle V_t^{(1)}, V_t^{(2)}, ..., V_t^{(n)} \rangle$. In this sorted (or ordered) vector, $\mathbf{v}_t^{\max} = V_t^{(n)}$ is the highest element, and $V_t^{(n-1)}$ is the second highest. $(\mathbf{v}_t^{\max})\prime$ is the early exercise boundary value *if this was a standard option written on a single underlying asset*, thus $(\mathbf{v}_t^{\max})\prime = D(\alpha + \beta \ln(T-t))$.

In equation 4.22, $(\mathbf{v}_t^{\max})^*$ approximates the early exercise region for the highest element in $\mathbf{v}_t$: for a max-call option, if the value of $\mathbf{v}_t^{\max}$ is higher than $(\mathbf{v}_t^{\max})^*$, then it is optimal to exercise the option. If the underlying assets' characteristics are the same (i.e., if they all have the same σ, δ, and D), then only three parameters are needed for the TSMC method (α, β, and θ_t) to evaluate the option regardless

of the value of n. If the asset characteristics are different, then each asset needs its own α and β values.

In equation 4.22, the selection of V_t^i deserves a detailed discussion. Since θ_t measures the minimum acceptable spread between $\mathbf{v}_t^{\max}$ and V_t^i, the selection of V_t^i has a direct impact on the calculated option value. There are a number of candidates for V_t^i : the second highest element in the ordered vector $\mathbf{v}_t$; the lowest element; the second lowest element; or a composite value (such as the average of lowest and second lowest elements). It is important to note that, whatever the choice is, this simplification is likely to cause a slight bias and undervalue the option, since not all the interactions between the assets can be captured with θ_t. In the TSMC method, after some experimentation, the second highest element $V_t^{(n-1)}$ was selected as V_t^i.

At time T, the option is exercised if and only if $\mathbf{v}_t^{\max} > D$. At any time period before T, the option is exercised only if the value of the leading asset is higher than its critical value (as defined in equation 4.22).

The role of θ in the case of a max-option on two assets is illustrated in Figure 4.1. The critical value of V^2 without the presence of V^1 is 1 (i.e., $(V_t^2)\prime = 1$). For the indicated value of $V_{0.5T}^1$, notice that $(V_{0.5T}^2)^*$ includes an "opportunity cost premium" of $V_{0.5T}^1 - (1 - \theta_{0.5T})$ on top of $(V_t^2)\prime$. Thus the relevant critical value for V^2 is $V_{0.5T}^1 + \theta_{0.5T}$. The shaded region represents the area of continuation (holding the option).

This property can also be used to reduce the dimensionality of the option valuation problem when $n > 2$, since the valuation problem boils down to finding a suitable representation of the premium based on θ and a selected metric that

represents the assets trailing the maximum asset price (e.g., in TSMC, this metric is chosen as the second highest price of the vector that holds the asset prices at each time t. Other metrics are also possible, such as the average of assets, minimum asset, etc).

As maturity approaches, the probability of significant changes in $\mathbf{v}_t$ decreases (e.g., for the geometric Brownian process, the variance of the change grows linearly with the time horizon). Therefore, depending on the remaining maturity of the option, $\theta_{initial}$ needs to be decreased so that the exercise rule is relaxed as T approaches. This is achieved by a linear reduction in $\theta_{initial}$ using the scaling factor $(1 - \frac{t}{T})$. In other words, $\theta_t = \theta_{initial}(1 - \frac{t}{T})$. At time T, this scaling factor goes to zero, effectively removing the θ from the decision rule (since no more future price movements are possible).

The following example demonstrates how the early exercise decision is made for a max-call option. Assume that at time $t = 0.5T$, there are four underlying assets $\mathbf{v}_{0.5T} = \langle V^1_{0.5T}, V^2_{0.5T}, V^3_{0.5T}, V^4_{0.5T} \rangle$ and the stochastic processes result in the following vector of prices $\mathbf{v}_{0.5T} = \langle 120, 110, 66, 115 \rangle$. Assuming that $(\mathbf{v}^{\max}_{0.5T})\prime = 100$, and $\theta_t = 0.2$, exercise is not optimal. Indeed, by applying equation 4.22, we can compute the opportunity cost premium for early exercise: $110 - (0.8)(100) = 30$. Thus the relevant critical value becomes $100 + 30 = 130$, and given that $120 < 130$, early exercise is not optimal. In other words, there is still a reasonably high probability that one of the other assets (V^2, V^3, V^4) can exceed the value of $V^1_{0.5T}$ and yield a higher option value.

4.5.2 Steps of the Tabu Search Monte Carlo Method

TSMC method uses an embedded sorting algorithm to determine $\mathbf{v}_t^{\max}$ and the ranking of other assets in the asset vector. During all m Monte Carlo runs, the option values from individual sample paths,F_0, are calculated using Equation 4.3. At the end of all Monte Carlo iterations, the mean option value from m simulation runs is calculated along with other descriptive statistics (such as the standard deviation of the mean option value).

The steps of the Tabu Search Monte Carlo algorithm are given in Table 4.2. The algorithm is designed to find the near-optimal[10] values of α, β and θ_t which define the stopping times and, thus, can be used to establish early exercise rules.

For an option problem with a single underlying asset, the tabu search algorithm is run with an initial solution[11] $s_{initial} = \{\alpha_{initial}, \beta_{initial}\}$ which defines a unique early exercise boundary (for an option problem with multiple underlying assets, $s_{initial} = \{\alpha_{initial}, \beta_{initial}, \theta_{initial}\}$)[12]. Since tabu search is a heuristic based on improving the initial solution in successive iterations, the better the initial solution, the closer the output of the algorithm will be to the optimal value. Therefore starting with high-quality initial solutions increases the chances of obtaining a better solution, and at the same time, reduces the amount of computational time needed to achieve close-to-optimal solutions.

---

[10] As discussed in Chapter 2, the tabu search method does not guarantee optimality. Therefore the objective is finding near-optimal solutions in a reasonable amount of computing time.

[11] Here, the term "solution" refers to the set of TSMC parameters that are changed from one iteration to the next. For a more detailed discussion of concepts related to Tabu Search, see Section 2.4.

[12] Note that the values of $\alpha_{initial}$ and $\beta_{initial}$ are not time dependent, and they apply to all $t \in [0, T]$. Whereas, $\theta_{initial}$ refers to the value of θ_t at time $t = 0$. The value of θ_t values for $t \in (t, T]$ is calculated based on the linear reduction formula $\theta_t = \theta_{initial}(1 - \frac{t}{T})$.

For the option valuation problem, certain characteristics of the problem, such as σ, have a direct impact on the initial solution. For instance, for a call option with $\sigma = 0.2$ and $T = 1$, an initial α value of 1.2 seems to be reasonable, while increasing σ to 0.4 makes 1.4 a better choice for the initial α value. Therefore, changes in other factors which affect the option values, such as changes in r, δ and T will also have an impact on the values of α and β. As a general guideline, any changes that may increase the value of an option will result in higher (lower) values for α for a call (put) option (assuming that the value of β remains the same).

Given $s_{initial}$, the tabu search algorithm creates the initial early exercise boundary and then computes the corresponding initial option value with a Monte Carlo routine. At each tabu search iteration, parameters α and β are incremented based on moves of ± 0.01, which gives a total of four moves to be evaluated per iteration $\{\alpha+0.01; \alpha-0.01; \beta+0.01; \beta-0.01\}$. For max-options, there are a total of six moves $\{\alpha + 0.01; \alpha - 0.01; \beta + 0.01; \beta - 0.01; \theta_t + 0.01; \theta_t - 0.01\}$. These moves also define the neighborhood of a solution. For instance, the solution $s_1 = \{1.2, 0.1, 0.2\}$, has the following neighborhood $N(s_1)$:

$$s_{11} = \{1.19, 0.1, 0.2\}$$

$$s_{12} = \{1.21, 0.1, 0.2\}$$

$$s_{13} = \{1.2, 0.09, 0.2\}$$

$$s_{14} = \{1.2, 0.11, 0.2\}$$

$$s_{15} = \{1.2, 0.1, 0.19\}$$

$$s_{16} = \{1.2, 0.1, 0.21\}$$

The simple operation of changing a parameter by ± 0.01 connects s_1 to all of these solutions in its neighborhood[13].

For options on a single underlying asset, each move corresponds to a unique early exercise boundary (a sufficient early exercise decision rule). For multidimensional options, moves also include the values for θ_t, the parameter used for calculating the opportunity cost "premium" of early exercise when multiple assets are present.

In order to evaluate each move, Monte Carlo paths are generated first. Using the generated paths and the given early exercise boundary, each move is evaluated. Based on this evaluation and the tabu status of the move, the best move is selected for each iteration. Then, the move is implemented and the current solution $s_{current}$ is updated. If $s_{current}$ is the best solution found so far during the tabu search iterations, s_{best} is updated as well. One of the core concepts of tabu search is that $s_{current}$ is the best solution *in its neighborhood*, and not necessarily the best solution *overall*. This ensures that the algorithm does not advance in a greedy fashion, and preserves the option of reaching the global optimum by not locking the search process into local optima.

[13] Note that other operations can also be used to modify the solutions. The value of 0.01 is chosen based on the intervals of α, β and θ (see section 4.5.5). These intervals determine the size of the search space and the corresponding range of tabu search iterations needed to perform an effective search. Although not implemented in TSMC, one possible search strategy can be based on variable size moves: starting with higher increments, the size of the changes in α, β and θ can be reduced as the search progresses.

Step 1	**Initialization**
	Define the initial solution, $s_{initial} = \{\alpha_{initial}, \beta_{initial}, \theta_{initial}\}$ $s_{current} = s_{initial}$ Assign current best known solution, $s_{best} = s_{current}$. Define $option\_value = \text{eval}(s_{best})$, where $eval$ is the function that computes the option value using equation 4.29 using the stopping times defined by $s_{initial}$.
Step 2	**Evaluation**
(A)	Generate the moves to be evaluated within $N(s_{current})$, pick $s_{next} \in N(s_{current})$, $s_{current} = s_{next}$.
(B)	Generate m sample paths using Monte Carlo simulation Evaluate $s_{current}$ using the $eval$ function (equation 4.29) and the stopping times given by $s_{current}$. Also check the tabu status of the move. If a new best solution is found, $s_{best} = s_{current}$ and $option\_value = \text{eval}(s_{best})$. Pick next move to be evaluated, $s_{next} \in N(s_{current})$, assign $s_{current} = s_{next}$ and repeat Step 2 (B).
Step 3	**Update and Termination**
	Repeat Step 2 until a specified number of tabu search iterations are completed.

Table 4.2: TSMC algorithm flow

4.5.3 Monte Carlo Simulation in TSMC

The TSMC method uses an embedded Monte Carlo algorithm that generates random paths for the underlying assets. Equation 4.4 can be used to describe the behavior of the asset over time, assuming that it follows geometric Brownian motion.

Substituting $\epsilon\sqrt{dt}$ for dz, equation 4.4 now becomes:

$$dV = \mu V dt + \sigma V \epsilon \sqrt{dt} \tag{4.23}$$

Transforming the variables by $G = \ln V$ and using Ito's lemma:

$$dG = \left(\frac{\partial G}{\partial V}\mu V + \frac{\partial G}{\partial t} + \frac{1}{2}\frac{\partial^2 G}{\partial V^2}\sigma^2 V^2\right) dt + \frac{\partial G}{\partial V}\sigma V \epsilon \sqrt{dt} \tag{4.24}$$

now substituting

$$\frac{\partial G}{\partial V} = \frac{1}{V}, \frac{\partial^2 G}{\partial V^2} = -\frac{1}{V^2}, \frac{\partial G}{\partial t} = 0 \tag{4.25}$$

into equation 4.17:

$$dG = \left(\mu - \frac{\sigma^2}{2}\right) dt + \sigma\epsilon\sqrt{dt} \tag{4.26}$$

By reversing the log transformation, we can estimate the value of the underlying asset over time using the following formula:

$$V_{t+dt} = V_t \exp(\left(\mu - \frac{\sigma^2}{2}\right) dt + \sigma\epsilon\sqrt{dt}) \tag{4.27}$$

For an underlying asset that yields continuous dividend, δ, equation 4.27 becomes:

$$V_{t+dt} = V_t \exp(\left(r - \delta - \frac{\sigma^2}{2}\right) dt + \sigma\epsilon\sqrt{dt}) \tag{4.28}$$

Therefore, given V_0, r, δ, σ and dt increment, the value of the underlying asset can be calculated using equation 4.28 at any point in time. By making random draws from $N(0,1)$, we can simulate the future trajectory of the underlying asset price.

Geometric Brownian motion is certainly not the only process that can be used to model the behavior of the underlying asset in Monte Carlo simulation. Other dynamics are possible for the movement of V over time, such as the diffusion model of Bass (1969), jump-diffusion and mean reversion.

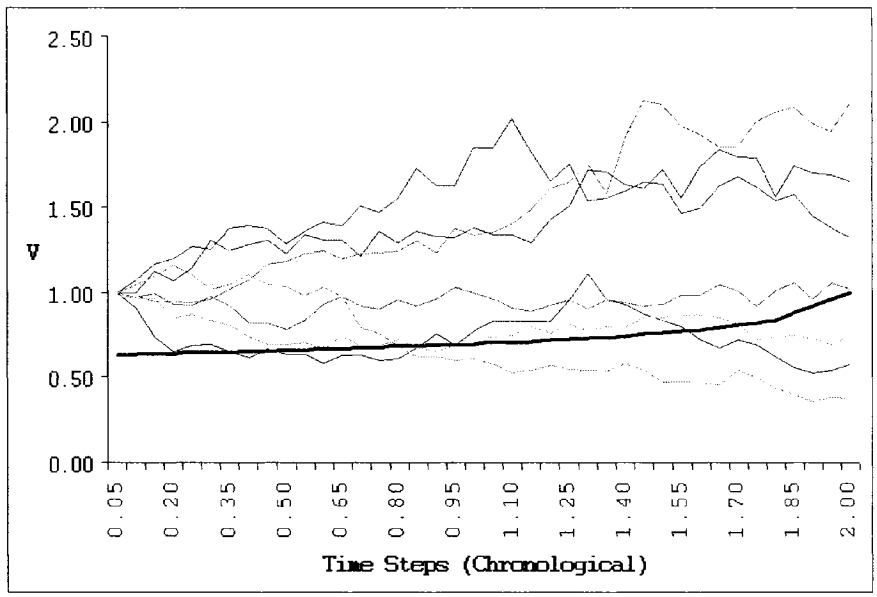

Figure 4.2: Monte Carlo simulation paths and the early exercise boundary (single asset)

4.5.4 Graphical Representation of the TSMC Algorithm

The following graphs demonstrate how the tabu search algorithm and Monte Carlo simulation work together. In Figure 4.2, sample paths of the underlying asset and the initial early exercise boundary are shown for a put option written on a single asset with a maturity date (T) of 2 years.

Tabu search iterations modify this early exercise boundary through successive iterations. At the end of the iterations, the best solution found during the search corresponds to the final option value. Figure 4.3 shows the evolution of the early exercise boundary during the search process. The bold curve corresponds to the final boundary that was found at the end of the tabu search iterations, which yields the maximum option value.

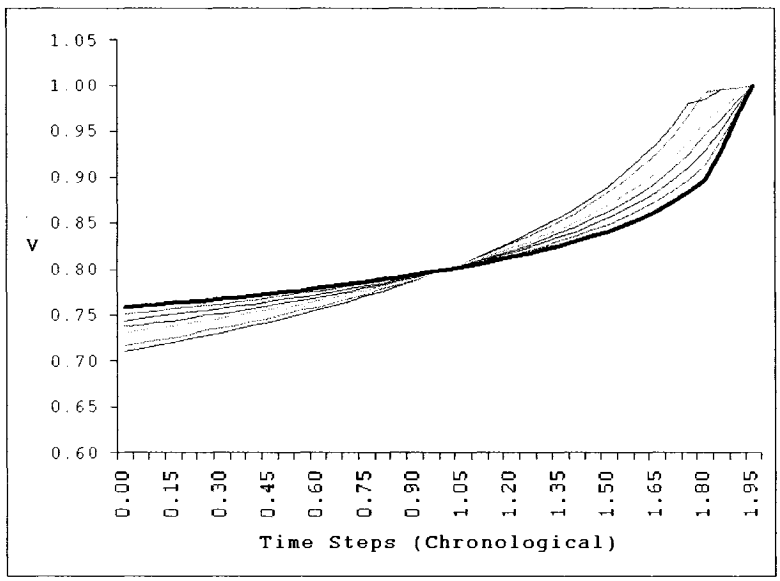

Figure 4.3: Evolution of the early exercise boundary (single asset)

For the case of a max-option on two assets, the panel in Table 4.3 demonstrates the evolution of the early exercise regions as the option approaches maturity. There are three main regions in each graph, the waiting region (closest to the 45^o line partitioning the graph), early exercise region for Asset 1 (below the curve under the 45^o line), and the early exercise region for Asset 2 (above the curve over the 45^o line).

It is interesting to note that the waiting region shrinks as the maturity approaches, and (for $t < T$), it is not optimal to exercise the option when the two asset prices are the same (even if both prices are very high). At time T, the waiting region is confined to the square formed by the exercise prices of the two assets and anywhere outside the square, the option will be exercised. Furthermore, as the option's maturity is increased, the area of the waiting region also increases.

During the option's lifetime, if Asset 1 or Asset 2 hit the zero boundary, then due to the nature of the stochastic process of equation 4.4, that asset's value stays at zero for the rest of the options' lifetime. In this case, the option is transformed into an option written on a single underlying asset.

4.5.5 Constructing Probabilistic Bounds

Since simulation-based methods can be prone to significant biases in estimated values, determining the quality of a given solution is not trivial. The ability to create lower and upper bounds for a given point estimate is desirable in order to assess the overall accuracy of an estimate, and also to evaluate the results based on a well-defined confidence interval.

Garcia (2003) provides a technique for constructing lower and upper bounds that is directly applicable to the TSMC method. Let $F_M(s)$ be the mean of the discounted payoffs of an American option based on M simulations for a given early exercise policy defined by $s = \{\alpha, \beta, \theta\} \in \Theta$. The set Θ is composed of subsets for each parameter. For α, Θ_α is $[0, 1]$ for a put option and $[1, 2]$ for a call option. For β, Θ_β is $[-0.6, 0]$ for a put option and $[0, 0.6]$ for a call option. For θ, Θ_θ is $[-3, 3]$.

$$F_M(s) = \frac{1}{M} \sum_{i=1}^{M} \exp(-r\tau(s))\Omega(X^i_{\tau(s)}, \tau(s)) \tag{4.29}$$

In this equation, the payoff of the option, Ω, is calculated based on a given realization of the stochastic process X^i and the corresponding stopping time τ. The stopping time τ is determined by the exercise policy dictated by the early exercise policy, s.

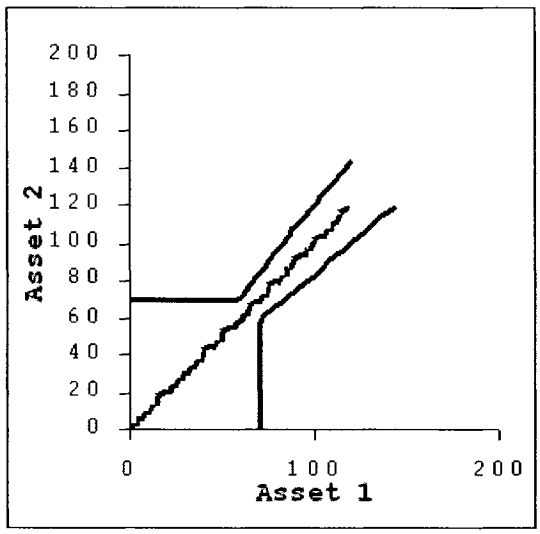

$t = 0$

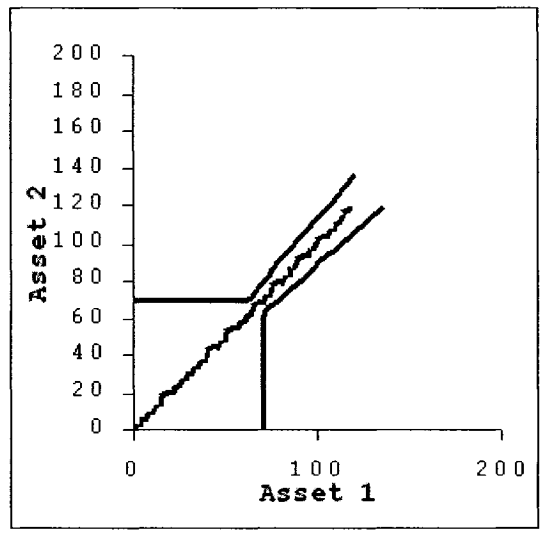

$t = 0.33T$

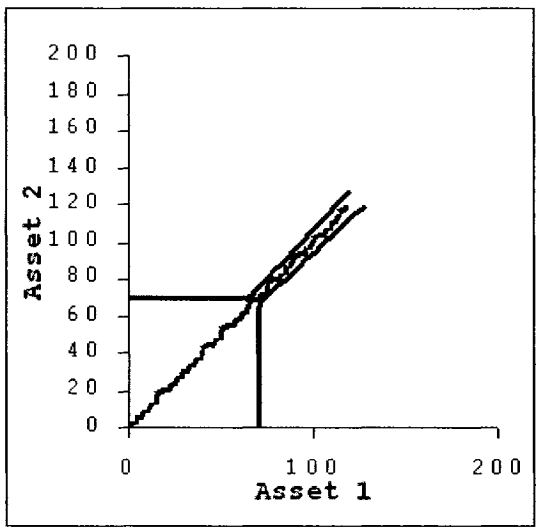

$t = 0.66T$

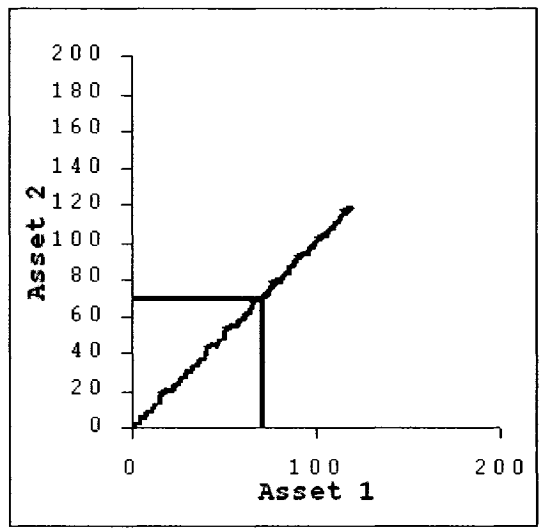

$t = T$

Table 4.3: Evolution of early exercise regions over time

The objective of the optimization algorithm is to find the optimal early exercise policy such that F_M is maximized.

$$\hat{F}_M \equiv F_M(\hat{s}_M) \equiv \max_{s \in \Theta} F_M(s) \tag{4.30}$$

Since the estimate $\hat{F}_M$ is obtained using a simulation procedure embedded into an optimization routine, for finite M, $\hat{F}_M$ is biased high (Garcia, 2003). Tabu search will pick $\hat{s}$ such that the estimate of the option value $F_M(\hat{s}_M)$ is the maximum over all parameter values that were evaluated, even the true optimal exercise policy.

To counter the impact of this high biased estimator, we need a second estimator which can be used to construct a lower bound for the true option price. For this low biased estimator, Garcia (2003) suggests generating a new set of Monte Carlo paths, and evaluating the option using the early exercise policy found by the high biased estimator, $\hat{s}_M$. Specifically,

$$\hat{f}_M = \frac{1}{m} \sum_{j=1}^{m} \exp(-r\tau(\hat{s}_M)) \Omega(X^j_{\tau(\hat{s}_M)}, \tau(\hat{s}_M)) \tag{4.31}$$

Simulated X^j is independent from the initial set of Monte Carlo paths used for finding $\hat{s}_M$. For finite M, $\hat{f}_M$ is biased low. Garcia (2003) provides the proofs for both the low and high biased estimators and he also states that both of these estimators are asymptotically unbiased.

Therefore, once $\hat{F}_M$ and $\hat{f}_M$ are estimated, probabilistic confidence intervals can be constructed using the following bin:

$$\left[\hat{f}_M - z_{\frac{\psi}{2}} \sigma(\hat{f}_M), \hat{F}_M + z_{\frac{\psi}{2}} \sigma(\hat{F}_M)\right] \tag{4.32}$$

where $z_{\frac{\psi}{2}}$ is the $1-\psi$ quantile of the standard normal distribution, and $\sigma(\hat{f}_M)$ and $\sigma(\hat{F}_M)$ are the standard deviations of the low- and high-biased estimators, respectively. All TSMC results in section 4.6 are reported using these confidence bins. The TSMC point estimates in the result tables refer to the arithmetic average of $\hat{F}_M$ and $\hat{f}_M$.

4.6 Results

In this section, results obtained by implementing TSMC for both unidimensional and multidimensional option problems are presented. In order to provide a clear indication of the performance of the TSMC algorithm, well-documented option problems from the literature are used in the numerical experiments.

4.6.1 American Put Options (on a single asset)

Table 4.4 presents the results obtained by the TSMC method for an American put option, with $D = \$40$ and $r = 0.06$. In order to evaluate the performance of TSMC, the results of finite difference and Least Squares Monte Carlo (LSM) methods are also presented. These results are taken from Longstaff and Schwartz (2001). The finite difference method is based on 40,000 time steps per year and 1,000 steps for the stock price, and it is taken as a benchmark for accuracy. The LSM method is based on 100,000 simulation paths using 50 exercise points per year.

As it can be seen in Table 4.4, the differences between the TSMC results and the finite difference results are very small. Of the 20 differences shown, 13 are less than or equal to one cent in absolute value. Also based on the presence of both

V	σ	T	Finite Difference	LSM	TSMC (point estimate)	$\alpha_{best}, \beta_{best}$	TSMC 95% C.I.	LSM accuracy	TSMC accuracy
36	.20	1	4.478	4.472	4.488	{0.83, -0.05}	[4.454, 4.523]	.006	-.010
36	.20	2	4.840	4.821	4.824	{0.84, -0.04}	[4.770, 4.877]	.019	.016
36	.40	1	7.101	7.091	7.073	{0.61, -0.1}	[6.984, 7.163]	.010	.028
36	.40	2	8.508	8.488	8.499	{0.64, -0.1}	[8.406, 8.595]	.020	.009
38	.20	1	3.250	3.244	3.244	{0.83, -0.05}	[3.217, 3.277]	.006	.006
38	.20	2	3.745	3.735	3.749	{0.84, -0.04}	[3.707, 3.791]	.010	-.004
38	.40	1	6.148	6.139	6.148	{0.61, -0.1}	[6.074, 6.222]	.009	.000
38	.40	2	7.670	7.669	7.654	{0.64, -0.1}	[7.577, 7.730]	.001	.016
40	.20	1	2.314	2.313	2.316	{0.83, -0.05}	[2.278, 2.356]	.001	-.002
40	.20	2	2.885	2.879	2.886	{0.84, -0.04}	[2.855, 2.918]	.006	-.001
40	.40	1	5.312	5.308	5.297	{0.61, -0.1}	[5.236, 5.359]	.004	.015
40	.40	2	6.920	6.921	6.899	{0.64, -0.1}	[6.828, 6.972]	-.001	-.021
42	.20	1	1.617	1.617	1.611	{0.83, -0.05}	[1.577, 1.645]	.000	.006
42	.20	2	2.212	2.206	2.208	{0.84, -0.04}	[2.179, 2.236]	.006	.004
42	.40	1	4.582	4.588	4.574	{0.61, -0.1}	[4.521, 4.626]	-.006	.008
42	.40	2	6.248	6.243	6.226	{0.64, -0.1}	[6.164, 6.289]	.005	.022
44	.20	1	1.110	1.118	1.102	{0.83, -0.05}	[1.080, 1.125]	-.008	.008
44	.20	2	1.690	1.675	1.685	{0.84, -0.04}	[1.662, 1.709]	.015	.005
44	.40	1	3.948	3.957	3.953	{0.61, -0.1}	[3.910, 3.995]	-.009	-.005
44	.40	2	5.647	5.622	5.608	{0.64, -0.1}	[5.525, 5.692]	.025	.039

Table 4.4: Comparative results for unidimensional options

positive and negative differences, just like LSM, TSMC seems to be free of any significant bias in estimation. Overall, the performance of TSMC is comparable to that of LSM, and therefore it can be used to calculate the value of financial options.

In order to obtain the results presented in Table 4.4, the following initial solutions were used. For options with $\sigma = 0.2$, $s_{initial} = \{0.8, -0.1\}$; for options with $\sigma = 0.4$, $s_{initial} = \{0.6, -0.1\}$. In order to determine the best α and β parameters, TSMC was run with 10,000 Monte Carlo paths for 15 tabu search iterations, with a running time of 81 seconds (using a 1.2 GHz Celeron processor). Then, the best parameters were used to run another 5 tabu search iterations, this time with 50,000 Monte Carlo paths to increase the accuracy and to calculate the upper and lower bounds. This second set of iterations requires 135 seconds. All iterations were run with 72 time steps (and 72 exercise points).

4.6.2 Path-Dependent Options

Path-dependent options are a special class of options, where the payoff of the option depends on the average value of an asset over a specified time period. This feature of path-dependent options makes them desirable as risk management tools, since the payoff from path-dependent options can mirror the risk exposure of firms (Grant, Vora and Weeks, 1997).

One of the most prominent path-dependent options is the Asian option. An "Eurasian" option can only be exercised at maturity, while "Amerasian" options can be exercised at any time before maturity. Other variations are also possible, such as imposing an initial lockout period, limiting the number of early exercise

opportunities (i.e., adding Bermudan option characteristics), and making the payoff depend on the geometric mean.

Longstaff and Schwartz (2001) report that the valuation of path-dependent options can be particularly complex, since on top of the American exercise features, the average value of the asset needs to be taken into account for each exercise opportunity. One way of solving this valuation problem is to structure the problem such that the exercise decision is based on two state variables for each time t: the current value of the asset and its average value for a given time window.

When simulation-based methods are used, adding this second variable is straightforward, and it does not increase the computational time considerably. However, for other methods which are prone to the curse of dimensionality (such as finite differences), path-dependent option valuation can become quite complex.

Table 4.5 presents the results for an Amerasian call option written on the arithmetic average of an asset, V_t. The second state variable (the arithmetic average) is denoted by $\bar{V}_t$.Therefore the payoff of this option is given by

$$\Omega(\bar{V}, t) = \max(0, \bar{V}_t - D) \tag{4.33}$$

There is a critical average price $\bar{V}_t^*$ above which early exercise is optimal. As in the case of max-call American options, this critical price not only depends on the value of $\bar{V}$ for a given t, but also on the value of V at the same time period. In the TSMC framework, this relationship between the average of an asset's value and the asset value itself can be expressed as

D	GVW	TSMC Point Estimate	TSMC 95% C.I.	$\{\alpha, \beta, \theta\}$	% difference
$\sigma = 20\%$ p.a.					
90	13.22	13.21	[13.092, 13.347]	{1.02, 0, 0.23}	0.075
95	9.14	9.16	[9.079, 9.232]	{1, 0, 0.21}	-0.219
100	5.80	5.81	[5.756, 5.879]	{1, 0, 0.27}	-0.172
105	3.35	3.36	[3.319, 3.412]	{1, 0, 0.2}	-0.298
110	1.75	1.76	[1.719, 1.806]	{1, 0, 0.2}	-0.571
$\sigma = 30\%$ p.a					
90	14.51	14.58	[14.414, 14.785]	{1.09, 0.02, 0.3}	-0.482
95	10.92	10.95	[10.805, 11.133]	{1.12, 0.02, 0.25}	-0.274
100	7.92	7.95	[7.839, 8.098]	{1.09, 0.02, 0.25}	-0.378
105	5.53	5.54	[5.459, 5.648]	{1.1, 0.02, 0.27}	-0.180
110	3.72	3.75	[3.689, 3.838]	{1.05, 0.01, 0.26}	-0.806

Table 4.5: Comparative results for Amerasian options

$$(\bar{V}_t)^* = (\bar{V}_t)\prime + \max\left[0, (V_t - (1-\theta_t)(\bar{V}_t)\prime)\right] \tag{4.34}$$

where $(\bar{V}_t)\prime = D(\alpha + \beta \ln(T-t))$.

Grant, Vora and Weeks (1997) provide results for an Amerasian call option with the following characteristics: $V_0 = \$100$, $D = \$100$, $T = 120$ days, $r = 9\%$ p.a., $\sigma = 30\%$ p.a. Furthermore, the average asset price is calculated from daily prices starting from day 91 and ending at day 120. There is also a lockout period until the end of day 105. Therefore, including the expiration day, there are a total of 16 exercise opportunities. The results of Grant, Vora and Weeks (1997) are presented in the second column of Table 4.5. Using TSMC, each row of the table was calculated in 412 seconds (with a 1.2 GHz Celeron processor).

4.6.3 American Max-Call Options

TSMC method can also be used for analyzing multidimensional options. One of the classical multidimensional problems in the options literature is an American call option written on multiple assets, the American max-call option. This problem serves as a benchmark for comparing the performance of various methods developed for the valuation of multidimensional options. In recent years, significant progress has been made in this area, and many alternative methods have been developed, including Barraquand and Martineau (1995), Broadie and Glasserman (2004), Andersen and Broadie (2004), Haugh and Kogan (2004), Kay, Davison and Rasmussen (2003), Ibanez and Zapatero (2001), Longstaff and Schwartz (2001) and Garcia (2003).

The American max-call option problem can be described as follows. In a risk-neutral framework, n assets follow correlated geometric Brownian motion processes, such that:

$$dV_t^i = (r - \delta_i)V_t^i dt + \sigma_i V_t^i \epsilon_t^i \sqrt{dt} \tag{4.35}$$

The instantaneous correlation of $\epsilon^i$and $\epsilon^j$is given by ρ_{ij}. The problem can be simplified by assigning $\delta_i = \delta$ and $\rho_{ij} = \rho$ for all $i, j = 1, ..., n$ where $i \neq j$. Furthermore, the interest rate, r, is assumed to be constant throughout the maturity of an option. At time t, the values of the underlying assets form vector $\mathbf{v}_t$.

Similar to the unidimensional case, early exercise features are added to the problem by specifying a number of early exercise opportunities. Thus, American option values are approximated by calculating the corresponding Bermudan op-

V_0	Binomial	95% C.I. (A&B)	TSMC Point Estimate	TSMC 95% C.I.	$\{\alpha, \beta, \theta\}$		% difference
$n = 2$ assets							
90	8.075	[8.053, 8.082]	8.056	[7.95, 8.199]	{1.2, 0.25}	0.08,	-0.23
100	13.902	[13.892, 13.934]	13.846	[13.717, 14.001]	{1.2, 0.25}	0.08,	-0.39
110	21.345	[21.316, 21.359]	21.318	[21.186, 21.481]	{1.2, 0.25}	0.08,	-0.13
$n = 3$ assets							
90	11.29	[11.265, 11.308]	11.263	[11.188, 11.424]]	{1.19, 0.41}	0.08,	-0.23
100	18.69	[18.661, 18.728]	18.603	[18.439, 18.806]	{1.19, 0.41}	0.08,	-0.46
110	27.58	[27.512, 27.663]	27.457	[27.24, 27.711]	{1.19, 0.41}	0.08,	-0.44
$n = 5$ assets							
90	n/a	[16.602, 16.655]	-	[16.36, 16.799]	{1.19, 0.6}	0.08,	-
100	n/a	[26.109, 26.292]	-	[25.756, 26.192]	{1.19, 0.6}	0.08,	-
110	n/a	[36.704, 36.832]	-	[36.339, 36.802]	{1.19, 0.6}	0.08,	-

Table 4.6: Comparative results for max-call American options

tion values. Increasing the number of early exercise opportunities increases the accuracy of this approximation.

The payoff of this option at exercise time t is:

$$F_t = \max(\max(V_t^1, ...V_t^n) - D, 0) \tag{4.36}$$

The most accurate results for this particular problem were obtained by using multidimensional binomial methods up to $n = 3$. In order to determine the accuracy of the results for $n > 3$, lower and upper bounds for the option value can also be calculated, see for example Andersen and Broadie (2004), Haugh and Kogan (2004), Ibanez and Zapatero (2001), and Garcia (2003).

TSMC results for $n = 2$, $n = 3$,and $n = 5$ are presented in Table 4.6. The first column of Table 4.6 indicates the initial value of all underlying assets. As a benchmark, second and third columns indicate the binomial results and 95% confidence intervals reported in Andersen and Broadie (2004). TSMC results for n=10 are presented in Table 4.7. Again, the first column of Table 4.7 indicates the initial value of all 10 underlying assets. As a benchmark, the second column shows the multidimensional European option values, and the third column indicates the results reported by Haugh and Kogan (2004).

The parameters of the options in Table 4.6 are: $D = \$100, r = 5\%, \delta = 10\%, \rho = 0, T = 3$, and $\sigma = 20\%$ (all underlying assets have the same parameters). For the max-call option, TSMC was run with the following initial parameters: $\alpha = 1.2$, $\beta = 0.1$, and $\theta = 0.2$. With these initial parameters, mean running times of the TSMC algorithm were 237 seconds for $n = 2$; 352 seconds for $n = 3$;and 591 seconds for $n = 5$ (using a 1.2 GHz Celeron processor). Similar to the unidimensional case, TSMC was run with 10,000 Monte Carlo paths for 15 tabu search iterations to determine the best α, β, and θ parameters. Then, these best parameters were used to run another 5 tabu search iterations with 50,000 Monte Carlo paths to increase accuracy[14].

The parameters of the options in Table 4.7 are: $D = \$100, r = 5\%, \delta = 10\%, \rho = 0, T = 1$,and $\sigma = 20\%$ (all underlying assets have the same parameters). 25 time steps were used in the calculations as well as 25 exercise opportuni-

[14] This benchmark max-call problem has a correlation coefficient (ρ) of zero. TSMC was also used to calculate the option values for non-zero values of ρ. As a general guideline, the smaller the ρ, the higher the option value. Therefore higher α values are more suitable for options with lower ρ values (all other factors remaining constant).

V_0	European Price	Lower and Upper Bounds (H&K)	TSMC Point Estimate	TSMC 95% C.I.	$\{\alpha, \beta, \theta\}$		% difference
$n = 10$ assets		25 exercise periods					
90	14.747	[15.158, 15.172]	15.109	[15.03, 15.22]	{1.21, 1.14}	0.05,	-0.37
100	26.403	[26.960, 26.970]	26.807	[26.65, 26.99]	{1.2, 1.15}	0.05,	-0.59
110	38.522	[39.182, 39.193]	39.039	[38.92, 39.19]	{1.22, 1.15}	0.11,	-0.38

Table 4.7: Comparative results for ten assets

ties. The reported percentage differences are based on the difference between the TSMC point estimate and the midpoint of the bounds reported by Haugh and Kogan (2004).

4.7 Discussion

As the results indicate, TSMC is a very versatile method for calculating the value of options with early exercise features. Although only the results for standard American options, max-call options, and path-dependent options are provided in this chapter, TSMC has the flexibility needed for studying other types of multidimensional options (for instance, in Chapter 5, dual strike American options are analyzed within a real options framework).

TSMC algorithm brings together the power and robustness of a proven combinatorial optimization heuristic with the flexibility of Monte Carlo simulation. This fusion results in a intuitive and transparent way for calculating the value of options. In TSMC, the way option values are calculated is straightforward and easy to explain.

Only three parameters are required to guide the optimization process, and there is no need to use any basis functions (unlike the LSM method of Longstaff and Schwartz (2001)) or a computationally intensive stochastic mesh method (Broadie and Glasserman, 2004). Being able to compute the early exercise boundary is an additional strength, since different options can be analyzed using graphs that plot the evolution of the boundary over time.

Since TSMC is a forward looking method based on numerical simulation, it is not limited by the curse of dimensionality and it can be used for calculating the value of multidimensional options. Furthermore, it is not based on any assumptions regarding the process of the underlying asset. Although the classical geometric Brownian implementation is demonstrated in this chapter, the method is compatible with other types of asset dynamics (as shown in Chapter 5). The core properties of the early exercise boundary remain the same for different stochastic processes and Monte Carlo simulation has the flexibility that is needed to model these processes.

TSMC algorithm is capable of providing a "custom" level of accuracy depending on the needs of the user. It provides a rough approximation of the true value very quickly (by picking a smaller number of Monte Carlo paths). Using 10,000 Monte Carlo paths, obtaining results within 2% of the true value takes less than 3 minutes (using a 1.2 GHz Celeron processor) even for multidimensional problems. The desired level of accuracy can be adjusted by changing the number of Monte Carlo paths and tabu search iterations.

TSMC is suitable for parallel computation since the Monte Carlo paths for each asset n can be computed on a different computer and then the option value

can be determined by merging the results of these paths. In addition to parallel processing, advanced simulation methods such as variance reduction can also increase the performance of the algorithm and help decrease the computation times even further.

As any other method, TSMC also has certain disadvantages:

It requires developing a dedicated software application, hence it has a high level of complexity for implementation. The object-oriented nature of the C++ code developed as part of this thesis work somewhat reduces this complexity and facilitates its implementation to different classes of problems.

Since it is based on Monte Carlo simulation, TSMC is computationally intensive. However, this shortcoming is partly overcome by the polynomial complexity of the method (with respect to the number of assets) and suitability for parallel processing.

As a versatile tool, TSMC method can be used to analyze both financial and real options. The accuracy of TSMC, as shown in the results section, is sufficient for most real options problems. Typically, these problems do not have very precise data, and the quality of what may seem to be an optimal solution is ultimately limited by the quality of the data and model assumptions. Therefore, by providing the ability to find solutions to multidimensional problems, TSMC can solidify the applicability of real options analysis to complex, real-life decision making problems.

4.8 Conclusion

The Tabu Search Monte Carlo method provides an alternative way for analyzing financial and real options. It represents a synergy between the fields of Operations Research and quantitative finance. TSMC is a hybrid method: it represents the fusion of a proven combinatorial optimization heuristic, with a very versatile numerical simulation technique.

Further work areas that remain are testing the TSMC method for stochastic exercise prices (an important element of real options problems) and max-options with different characteristics (including different values of ρ, σ and δ across assets). Another area of research can be customizing TSMC to tackle other types of multi-dimensional options, such as options on the product of underlying assets or spread options.

In the TSMC implementation presented in this chapter, θ_t operated on the values of the highest element and the second highest element of the vector containing asset prices. Alternative representations using θ_t are possible and require more research. One particular research possibility is to test other factors which may affect early exercise decisions, such as the second moments of assets.

Finally, the total time for computation can be further decreased by introducing more advanced simulation techniques (such as variance reduction).

Chapter 5
Car of the Future: A Real Options Analysis

5.1 Introduction

During the last two centuries, the phenomenal increase in energy production and consumption has been one of the major drivers of economic growth. Moving from solids, such as wood and coal, to liquids, such as petroleum, the efficiency and capability of our energy systems increased dramatically. During the past century, the supply of energy has grown at an average of 2 percent per year, stimulating a growing world economy at roughly 3 percent per year, on the average (Hamel, 2000).

Today, our energy systems are largely based on hydrocarbons, and although many initiatives exist to create future energy sources and technologies, we are still largely dependent on finite natural resources to satisfy our ever increasing appetite for energy. The recent surge in crude oil prices is a very strong sign indicating that "cheap oil" is a phenomenon of the past and that we need new technologies and energy sources to satisfy the growing demand for energy.

One of the biggest sources of demand for hydrocarbons is the transportation sector. The largest oil consumer of the world, the United States, accounts for about 25% of worldwide oil demand (Dunn, 2001). Of this amount, around 50% is used for fueling gasoline-powered automobiles (Oman, 2003). In other words, about

12% of worldwide demand for oil comes from the owners of passenger vehicles in the U.S., alone. It is clear that the passenger vehicle market has a pivotal role in the future dynamics of demand for energy in the U.S. and other countries.

This chapter examines the prospects for two emerging technologies in the automotive market: the hybrid electric vehicle (HEV), and the fuel cell vehicle (FCV). This "alternative vehicles problem" is modeled within a real options framework as an American dual strike option. The Tabu Search Monte Carlo (TSMC) method (introduced in Chapter 4) is used to calculate the value of this option and to determine the conditions which render switching from an internal combustion engine vehicle (ICEV) to an HEC or an FCV feasible.

Today, an overwhelming majority of the automobiles in use are internal combustion engine vehicles running on fossil fuels. However, arrival of new alternatives in the automotive market, more efficient renewable energy technologies, and increased volatility of oil prices are all factors that can change the balance and reshape the automotive market in the long run.

There are strong indications that the current energy regime, based mainly on fossil fuels, is not sustainable. Many experts believe that the economic rationale prevalent in the energy markets today does not internalize the effects of environmental damage such fuels are causing. Air pollution in major cosmopolitan centers, signs of climate change due to greenhouse gas emissions and occasional oil spills are exacting a toll on the environment.

Another dimension of the problem is the political one. The world's dependence on the Middle East for a steady supply of oil is a major concern. As OPEC

gains more strength in the world oil markets in the coming decades, and as the competition intensifies between the developed and developing economies for accessing remaining oil supplies, energy insecurity concerns and disruptions in the energy supply chain are likely to cause additional strains on the energy market.

Environmental and political dimensions add additional layers of complexity to the alternative vehicles problem. This chapter focuses primarily on the lifecycle ownership cost of an automobile from an individual consumer's perspective. As such, environmental and political dimensions of the problem are not addressed in the present framework. While this creates a limitation on the depth of this research, it also keeps the focus on the modeling of market and technical uncertainties using the TSMC method.

This chapter is structured as follows. Following a discussion of the current energy regime and emerging technologies such as hybrid electric vehicles and fuel cells, real options rationale of the alternative vehicles problem will be discussed. Then, the option to switch model will be presented along with the stochastic processes used in the model. Presentation of the analysis results and a discussion of future research areas conclude the chapter.

5.2 Current Energy Regime and Future Alternatives in the Passenger Vehicle Market

The oil crises in the 1970s encouraged Western economies to launch major oil and gas exploration efforts in the last two decades. However, most of these projects have resulted in few new discoveries of oilfields, and some experts believe that

about 90 percent of crude oil reserves in the world has already been discovered (Campbell and Laherrere, 1998). This makes the eventual decline in worldwide oil production inevitable, after a period of steady increase in production and hitting the peak. The balance between the oil demand and supply is a classical "stock and flow" problem, and the stock of existing oil supplies is finite (for all practical purposes). The flow (or demand induced production rates) of oil, on the other hand, is on a steady rise.

On the demand side, International Energy Agency, an OECD organization, estimates that by 2030 the worldwide demand for energy will be 60% higher than its 2004 levels (IEA, 2004). The growing economies of developing countries put additional strain on the balance between demand and supply. As India and China are demanding more and more energy to fuel their booming economies, the discrepancy between supply and demand in the crude oil market becomes even more pronounced.

When it comes to future of oil, the question is not whether or not we will run out of oil anytime soon, but it is when the price of oil will render it less desirable in the presence of a cost-competitive substitute, such as hydrogen. When the oil production peaked in the U.S. in 1970, the U.S. oil imports rose dramatically to satisfy the increasing demand. However, when the same phenomenon occurs on a worldwide scale, there will not be any other conventional oil reserves to satisfy the excess demand, and it is likely that there will be significant upside pressure on oil prices.

The estimates for the timing of worldwide oil production peak are causing an intense debate. Expert estimates range from 2010 (Campbell and Laherrere, 1998; Deffeyes, 2001) to 2030 and beyond (IEA, 2004). Even if the production will not peak before 2030, given the environmental and political dimensions of oil dependence, the need for alternative sources of energy is clear.

Experts think that using "unconventional" sources of oil, such as the heavy oil reserves in Venezuela and the tar sands deposits in Canada, can only postpone the inevitable. Furthermore, heavy oil has a much higher environmental cost, since its extraction requires more aggressive techniques than conventional oil (Roberts, 2004). Therefore, unconventional sources of oil cannot be seen as a long-term solution, and the quest for alternative sources of energy continues.

5.2.1 The Promise of Renewable Energy

Research and development (R&D) efforts on renewable energy sources and technologies have been gaining momentum in recent years, after a period of inertia largely due to relatively low oil prices. A wide range of renewable energy sources and technologies has been identified, including solar, wind, biomass and hydrogen based systems.

A fundamental difference between renewable energy and fossil fuels is the issue of long-term sustainability. Since all fossil fuels have been created over millennia of biological and geological activity, they are not regenerative (especially at the current levels of demand) and therefore, they are finite resources. Renewable energy generation, on the other hand, is based on developing advanced technolo-

gies for harnessing the power of natural processes, such as wind or solar radiation. Therefore, as long as these natural processes continue to act on Earth, they can be converted to usable energy. Moreover, harvesting these sustainable energy sources with more efficient technologies can only result in decreasing long-term production costs. When coupled with a supply source that is not finite, the efficiency gains through technological advances can be phenomenal. A technology driven, predominantly renewable energy regime would have more characteristics in common with the recent information technology revolution than the current paradigm of a highly centralized system based on "mining" finite resources.

Currently, production, storage and distribution costs of renewable energy sources are much higher than those of fossil fuels. However, given the trend of increasing crude oil prices, and the potential for stable (if not decreasing) prices for renewable energy, sooner or later these two price curves are going to intersect. In the long-run, this changing balance can trigger fundamental changes in our energy mix.

In the following sections, the interaction between oil and renewable energy sources will be analyzed within a real options framework for the case of passenger vehicles.

5.2.2 Hybrid Electric Vehicles

A very interesting transition is taking place in the automotive market today. Spearheaded by the models of Toyota and Honda, hybrid gas-electric vehicles (HEVs) are enjoying a very strong interest from consumers.

First introduced by Toyota in 1997, HEVs combine the benefits of gasoline and electric cars. At lower speeds, when the internal combustion engine is more polluting and has a low level of efficiency, the electric motor provides the propulsion. At higher speeds, when the vehicle needs more power, the small gas engine kicks-in (Fairley, 2004). Through a new technology called "regenerative braking", when the brakes are used, kinetic energy is transformed into electric energy to recharge the batteries. Therefore, there is no need to plug-in the vehicle to an electric outlet to recharge its batteries (a major advantage over earlier electric vehicles).

For consumers, HEVs provide a number of advantages over conventional ICEVs: fuel economy levels roughly twice as those of ICEVs, lower emissions, and use of advanced electronics. The combination of these features seems to be very attractive for the early adopters of this technology. From the day it was introduced into the market, Toyota's HEV model, Prius, sold a cumulative of 100,000 units in 2002, and 200,000 units in 2004 (Economist, 2004; Fairley, 2004; Lipman and Delucchi, 2003). While these are relatively small numbers compared to Toyota's total annual sales of about 7 million units in 2003, they also indicate a strong growth potential for HEVs.

Although HEVs provide solid advantages over ICEVs, compared to the potential of FCVs, they have some limitations. Since they are still dependent on gas engines, there is not as much room to increase their fuel economy. Moreover, although they have about 50% less emissions than ICEVs, they are nowhere close to providing near zero emissions. Finally, they still require a steady supply of oil, and

for this reason, energy security is still an issue for an automotive market composed largely of HEVs.

5.2.3 Hydrogen Energy and Fuel Cell Vehicles

Some of the most promising alternative energy technologies are based on hydrogen. Hydrogen is not a source of primary energy such as crude oil, coal or natural gas, and it does not exist freely in nature. It is a secondary form of energy, and it can be considered as an energy carrier (much like electricity) that needs to be generated using primary energy sources (Hoffmann, 2001). Hydrogen energy has numerous potential application areas, including transportation, residential and industrial power generation, and development of portable power sources for consumer electronics (Rifkin, 2002; Hoffmann, 2001).

In 1839, Sir William Grove demonstrated the basic principles of fuel cells by combining hydrogen and oxygen electrochemically to create water and electricity. Unlike batteries, fuel cells do not store chemical energy, and unlike conventional engines, they are not based on the thermal and mechanical processes based on combustion. Fuel cells are electrochemical devices that directly convert hydrogen (or hydrogen-rich fuels) into electricity (Sodal, 2003; Oueslati, 1999).

For more than a century, fuel cells were confined to laboratories. In 1960s, NASA decided to use fuel cells in the U.S. human spaceflight program, and fuel cells were used to power the Gemini and Apollo series spacecraft (Dunn, 2001). This heritage is also transferred to the U.S. Space Shuttle Program.

In the automotive sector, research on fuel cells was dormant until 1990s. Starting from early 1990s, automotive companies launched ambitious R&D projects to investigate the applications of fuel cells in passenger vehicles and alliances were built between automotive and fuel cell companies. This created a competitive R&D environment and resulted in the steady improvement of fuel cell efficiency (NAVC, 2003).

The world's first commercially available fuel cell vehicles were rolled out in December 2002: Toyota's FCHV and Honda's FCX (Dawson, 2002). Although these vehicles are much more expensive than conventional ICEVs today, learning curve effects and mass production can decrease the unit prices of FCVs in the long-run. When the first commercially available HEVs were launched (again by the same two companies) in late 1990s, their prospects were uncertain as well. However, the current demand for HEVs by the consumers is a very good indication that the automotive market is ripe for change, and consumers are willing to switch to alternative vehicles if the right conditions are met. Nevertheless, automotive experts do not expect the retail prices of FCVs to fall below $US 100,000 before the end of this decade (Dawson, 2002).

One of the major advantages of FCVs over ICEVs and HEVs is their potential for very low or zero emissions. However, for a consumer, a direct (and perhaps more relevant) benefit is the FCVs' superior efficiency and increased fuel economy. Today's ICEVs convert only about 15% of the heat content of petroleum into energy, whereas fuel-cells currently provide at least twice of this level of efficiency with a

	ICEV	HEV	FCV
Glider	Mid-sized sedan	Mid-sized sedan	Mid-sized sedan
Power Unit	Engine, Cooling System	Engine, Cooling System	Fuel Cell Module, Fuel Processor, Cooling System
Fuel Economy	25-45 mpg	50-70 mpg	80-140+ mpg
Range	350 miles	350 miles	350 miles
Fuel Type	Gasoline	Gasoline	Hydrogen
Retail Price (2004)	$US 15,000 - $US 17,000	$US 20,000 - $US 28,000	$US 200,000

Table 5.1: Comparison of Alternative Passenger Vehicles

significant upside potential for further efficiency gains through R&D (Economist, 2001).

5.2.4 A Framework for Comparing Future Vehicles

Previous research on future passenger vehicles has focused on different aspects related to vehicle ownership. Various studies examined the impact of different vehicles on air pollution, greenhouse gas emissions and lifecycle ownership costs. Ogden, Williams and Larson (2004), Weiss et al. (2003), Padro and Putsche (1999), MacLean and Lave (2003), and Hackney and de Neufville (2001) all provided comparative analyses of ICEVs, HEVs and FCVs. Most studies assumed an incremental increase in the performance of all three classes of vehicles within the next two decades. The rate of increase is expected to be highest for the FCVs (assuming that the current level of R&D investments continue). Table 5.1 summarizes the characteristics of the vehicles analyzed in this chapter (Table sources: Lipman and Delucchi, 2003; Ogden, Williams and Larson, 2004; Weiss et al., 2003).

Fuel economy is generally measured in "miles per gallon -mpg", the distance that the vehicle can be driven using a U.S. gallon of fuel. The mpg values in this chapter (either as mpg or as mpgge - miles per gallon of gasoline equivalent for fuel cell vehicles) represent a combined driving schedule of 55% urban and 45% highway driving.

5.3 Rationale for Real Options Analysis

As discussed in Chapter 4, the Tabu Search Monte Carlo (TSMC) method is a suitable tool for analyzing options that have early exercise features and multiple underlying assets. This capability renders TSMC a suitable choice for studying the alternative vehicles problem.

In order to evaluate future scenarios in the alternative vehicles problem, the future trajectory of oil and hydrogen prices and the associated R&D efforts need to be modeled. Therefore, both the natural resources and R&D streams of real options literature are relevant for this particular problem. A framework based on TSMC allows us to use different price dynamics, such as mean reversion and jump diffusion processes, and to determine the value of real options for both HEVs and FCVs from the consumer's perspective. In this way, both technical and market-based uncertainties can be assessed, and multiple sources of uncertainty can be incorporated into the model.

5.3.1 Financial Options vs. Real Options

As discussed in Chapter 2, the theory of real options is based on the conceptual framework of financial options. However, the analogy between a real asset (such as a new HEV) and a financial asset is imperfect and real options have certain characteristics that make their valuation different from financial options.

There are typically many sources of uncertainty that can affect the value of a real asset. For example, in a construction project, not only the future revenue stream, but also the cost of construction is uncertain at the beginning of a project. Therefore, one needs to model the uncertainty of both the revenue and cost streams of such a project.

Furthermore, the data used in real options calculations are generally not as precise as those for financial option calculations. The estimation of future revenue and cost streams are generally not very accurate, and in certain problem types, such as research and development (R&D) projects, these estimates can be very subjective. A general rule of thumb in mathematical modeling states that "the results from a model are as good as the data". Therefore, instead of striving for an arbitrary level of accuracy in real options analysis, the analysis should focus on testing many different assumptions in a reasonable amount of time. The TSMC method enables this; since the search process is relatively quick and transparent, many different alternatives that provide close-to-optimal solutions can be presented to the decision-maker in a timely manner.

Finally, the owner of a real option can drastically influence the value of the underlying asset, whereas the value of an underlying asset of a financial option is

determined by market interactions. In many cases, the decisions of a real option holder directly affect the value of a real option over time, and these decisions need to be modeled explicitly.

Since the TSMC method uses simulation, embedding these decision points into the models and assessing their potential impact on the underlying asset(s) are possible. Events like real option interactions, policy decisions, and external shocks can all be built into the simulation paths.

5.3.2 Real Options in Past HEV and FCV Research

As outlined in Section 5.2.4, there has been extensive research on alternative propulsion technologies and fuels. However, an overwhelming majority of the studies were based on methods other than real options. This is quite surprising, since a real options approach seems to be well suited to study the impact of major market and technical uncertainties in the alternative vehicles problem. Two studies that have implemented a real options analysis to study fuel cells are Oueslati (1999), in the automotive context, and Sodal (2003), in the maritime applications context.

Using a real options perspective, Oueslati (1999) has studied the decision of Ford to build an alliance with Ballard Power Systems to establish itself in the future FCV market. Oueslati (1999) also considers the prospects of investing in fuel cells for other application areas, such as power generation. For the passenger vehicle problem, he focuses mostly on the market penetration of FCVs and he does not take into account the fuel cost of competing systems. He uses a compound option framework coupled with parallel options (the use of parallel options to study the

interaction among projects was initially suggested by Childs, Ott and Triantis (1998) and Trigeorgis (1996)).

Another real options analysis on fuel cells was conducted by Sodal (2003), focusing on the potential use of fuel cells in the shipping industry for propulsion.

5.3.3 Sources of Uncertainty in the Alternative Vehicles Problem

The problem outlined in the previous sections consists of many different types of uncertainty, including market, technical, environmental and political uncertainties. The first two groups of uncertainties will be studied in this chapter.

Market uncertainties stem mainly from the following sources:

(a) the balance of demand and supply in the crude oil market and the future trajectory of oil prices

(b) the retail prices of alternative vehicles and their adoption rates in the market

(c) the price dynamics of hydrogen as fuel

Technical uncertainties include:

(d) rate of progress in the fuel efficiency of FCVs through R&D

These uncertainties can be captured using various stochastic processes and different model assumptions in a numerical simulation environment as explained in the next section.

5.4 The Car of the Future: An Option to Switch Model

Consider a prospective buyer in the automotive market who currently owns an automobile with an internal combustion engine running on petroleum (ICEV). The buyer's objective is to pick a new automobile such that the lifecycle cost of owning and operating the vehicle is minimized. The other factors that affect the purchasing decision are assumed to be the same for all vehicle alternatives (such as safety, quality of after-sales service, insurance costs, etc.) and therefore lifecycle ownership cost is assumed to be the primary decision criterion.

At any point in time during the option's lifetime, the buyer can make one of the following three decisions:

(a) Keep/renew the internal combustion engine vehicle (ICEV)

(b) Switch to a hybrid electric vehicle (HEV)

(c) Switch to a hydrogen fuel cell vehicle (FCV)

These decisions can be modeled as American options. If only a single switching alternative is available, then the corresponding option is an American option with a single underlying asset (i.e., switching from an ICEV to an HEV, or switching from an ICEV to an FCV). If either one of the switching alternatives is available to the prospective buyer, then the option is a multidimensional American option.

The real option model presented in this section is based on the same principles as the American option problems studied in Chapter 4. However, for the multidimensional option, this time there can be different exercise prices for each underlying asset (since HEVs and FCVs have different retail prices, options based

on these assets will have different exercise prices). Therefore, in this chapter, the multidimensional real option problem will be modeled as an American dual strike option. Broadie and Detemple (1997) give the formal definition and mathematical properties of dual strike options, these are discussed in detail in Section 5.4.5.

The switching model presented in this chapter is similar in spirit to Dixit and Pindyck's (1994) model on operating costs and temporary suspension. In their model, a decision maker has to decide operating a project or suspending it at any time period, based on the expected profits and costs in future time periods. Similarly, in the alternative vehicles problem, the decision maker has to decide between two "states": buying a new vehicle versus keeping the existing vehicle. One major difference between the approach of Dixit and Pindyck (1994) and this chapter is the treatment of time. In their model, it is assumed that the option is perpetual with an infinite time horizon. This assumption makes it possible to derive analytical solutions. However, if the finite time version of the same problem needs to be solved, analytical solutions are not possible. This can be seen as one advantage of the simulation-based TSMC method, since a wide range of finite time horizons can be studied. Furthermore, Dixit and Pindyck (1994) assume that no technological advancements are possible in future time periods. In the TSMC formulation, technological advancements can be incorporated into the option valuation problem as described in Section 5.6.

5.4.1 The Decision-Making Process of the Buyer

It is assumed that the customer making the purchasing decision is not an early adopter of the new HEV or FCV technology. He or she is part of the "early majority", who are typically risk-averse consumers and very price sensitive (see Moore (2002) for a discussion of technology adoption life cycles). Therefore the buyer in this problem has a primarily economic rationale for the purchasing decision and he or she follows the trends in the market, rather than leading the wave of technology adoption. His or her decision rule is primarily economic and is based on the expected savings that can be achieved by switching to a particular vehicle. In other words, the amount of expected savings is the underlying asset of the switching option. By exercising the option and purchasing a new vehicle (either HEV or FCV) at any time t, the decision maker is entitled to the expected savings that will accrue during the lifetime of the new vehicle. The exercise price of this switching option is the incremental price difference between a new ICEV and a new HEV or FCV (depending on the option that is studied). Therefore the payoff of the option to switch is the difference between the expected savings and the incremental price difference of vehicles.

If the expected savings is less than the price difference, then the buyer can wait until the next time period and re-evaluate. Once the decision is made to switch, it is assumed that the option has been exercised. In other words, this is not a perpetual option to switch.

It is further assumed that the aggregate amount of savings at time t is the discounted value of expected future savings for the next 10 years of vehicle own-

ership (in the case that the option's maturity is less than 10 years, only the savings within the option's maturity date are calculated). This assumption can be easily relaxed and other periods of vehicle ownership can be studied.

The equations for calculating the savings from different options are as follows.

Savings from option to switch from ICEV to HEV:

$$S_t^H(P^{oil}, mpg_I, mpg_H) = q\sum_{i=t}^{t+9}\exp(-r(i-t))P_i^{oil}(\frac{1}{mpg_I} - \frac{1}{mpg_H}) \qquad (5.1)$$

Savings from option to switch from ICEV to FCV:

$$S_t^{Fcell}(P^{oil}, P^{h2}, mpg_I, mpg_{Fcell}) = q\sum_{i=t}^{t+9}\exp(-r(i-t))(\frac{P_i^{oil}}{mpg_I} - \frac{P_i^{h2}}{mpg_{Fcell}}) \qquad (5.2)$$

where q is the average number of miles driven (per year), P_t^{oil} and P_t^{h2} are the retail cost per gallon of oil and hydrogen, respectively. Savings are cumulative and they are discounted to time t. When the value of option to switch from ICEV to HEV or FCV is evaluated, both S_t^H and S_t^{Fcell} are calculated and compared.

In equations 5.1 and 5.2, P_t^{oil} and P_t^{h2} are given by the realizations of the stochastic processes during the Monte Carlo runs. Note that a risk-free discount rate is used in the savings calculations. This risk-neutral framework is a common feature in many real options models, and it assumes that the drift rates of the stochastic processes of oil and hydrogen are adjusted for risk instead of using a risk-adjusted discount rate in the savings equations. For more information about risk-neutral valuation in real options, see Dixit and Pindyck (1994) and Kulatilaka (1993).

As discussed in detail in Chapter 4, the following equation (equation 4.29 in Chapter 4) gives a high-biased estimator for the option price.

$$F_M(s) = \frac{1}{M}\sum_{i=1}^{M} \exp(-r\tau(s))\Omega(X^i_{\tau(s)}, \tau(s)) \tag{5.3}$$

In this equation, the payoff of the option, Ω, is calculated based on a given realization of the stochastic process X^i and the corresponding stopping time τ. The stopping time τ is determined by the exercise policy dictated by the early exercise policy, s. The stochastic process for each Monte Carlo path contains all the information needed to calculate the savings (S_t^H or S_t^{Fcell}). The details of the stochastic processes used to model oil and hydrogen prices are given in Section 5.5. As explained in Chapter 4, a low-biased estimator can also be calculated by running a new set of Monte Carlo simulations given the best early exercise policy found using tabu search.

5.4.2 Isoquants for FCV Savings

As can be seen in equation 5.2, the savings from switching to an FCV depend on four separate variables. We can construct isoquants in order to determine the combined impact of these variables. This analysis is particularly useful for understanding the relative weights of the price of hydrogen and the FCV fuel economy technology on the overall fuel savings an FCV can offer over an ICEV.

The points (P_t^{h2}, mpg_{Fcell}) located on a particular isoquant correspond to a set of inputs (hydrogen prices and FCV fuel efficiency) which gives an average cost metric for driving an FCV. Depending on the realizations of P^{oil} and mpg_I for a

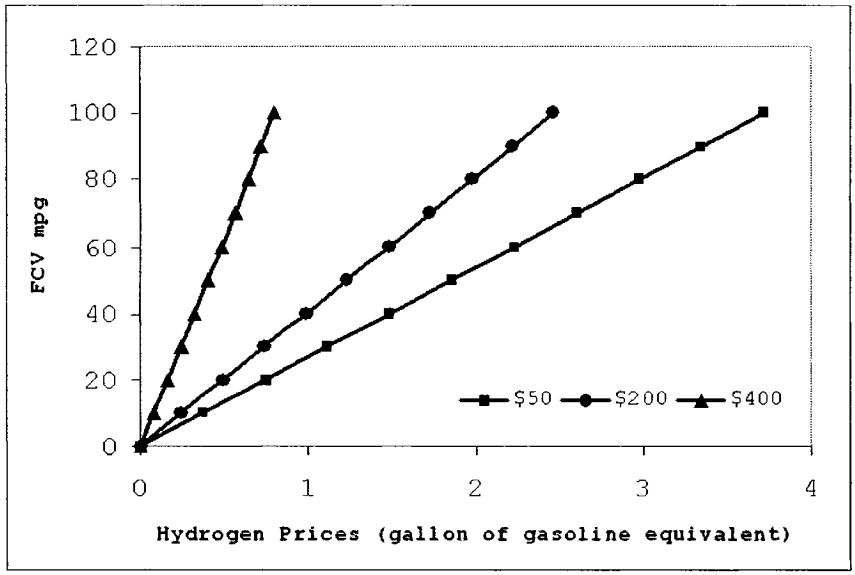

Figure 5.1: Isoquants for FCV savings

given time period, this isoquant determines the savings that can be achieved by switching to an FCV.

In Figure 5.1, each line represents a realized level of annual savings by switching to an FCV. For instance, the line with the steepest slope (closest to the y-axis) represents annual savings of $400. As it can be seen on the graph, in order to realize this level of savings, hydrogen prices need to stay below the $1 per gallon of gasoline equivalent (gge) mark (assuming that oil prices and the fuel efficiencies of the ICEVs stay unchanged for the period in question).

5.4.3 Option to Switch to a Single Alternative

During the lifetime of the option, at any point in time t, the payoff of the option to switch from an internal combustion engine vehicle to a hybrid electric vehicle is given by the following equation:

$$F_t^H = \max(S_t^H - D_H, 0) \tag{5.4}$$

Where S_t^H denotes the expected savings in the lifecycle fuel cost of the vehicle if the buyer switches from an ICEV to an HEV (the underlying asset of the option). D_H represents the price difference between buying a new ICEV and buying a new hybrid electric vehicle (the exercise price of the option).

The boundary conditions for solving this American call option are:

$$F^H(0,t) = 0 \tag{5.5}$$

$$F^H(S^{H*},t) = S^{H*} - D_H \tag{5.6}$$

$$\frac{\partial F^H}{\partial S^H}(S^{H*},t) = 1 \tag{5.7}$$

$$F^H(S^H,T) = \max(S_T^H - D_H,\ 0) \tag{5.8}$$

These boundary conditions are explained in detail in Chapter 4. Similarly, the payoff of the option to switch from an ICEV to a fuel cell vehicle is given by:

$$F_t^{Fcell} = \max(S_t^{Fcell} - D_{Fcell}, 0) \tag{5.9}$$

The boundary conditions for F^{Fcell} are identical to the conditions for F^H.

5.4.4 Option to Switch to Either One of the Alternatives

The buyer can also consider switching to either the HEV or FCV alternatives simultaneously. In this case, the payoff at time t is:

$$F_t^{I\rightarrow(H,Fcell)} = \max(\max(S_t^H - D_H, S_t^{Fcell} - D_{Fcell}), 0) \quad (5.10)$$

$S_t^{I\rightarrow(H,Fcell)}$ depends on the same variables as in Equations 5.1 and 5.2: the price of oil, the price of hydrogen and the fuel economies of the three types of vehicles. As it will be discussed in the subsequent sections, all of these variables can change over time.

5.4.5 Early Exercise Region for Multiple Assets

A standard American option written on a single asset has a continuous early exercise boundary that is monotonically decreasing with time (see Chapter 4). In the case of American options written on multiple assets, instead of a single boundary curve, there are multiple *early exercise regions*. Broadie and Detemple (1997) give a formal definition of the early exercise regions, as well as their mathematical properties. The shape of the regions are documented by Kay, Davison and Rasmussen (2003), Ibanez and Zapatero (2001) and Broadie and Detemple (1997).

One key observation of Broadie and Detemple (1997) is that, for an option on two assets V^1 and V^2, $max(V_t^1, V_t^2)$ is not a sufficient statistic for early exercise decisions. In other words, since the option is written on multiple assets, even if $V_t^1 > V_t^2$ there is always a chance that the value of V^2 may surpass that of V^1 at any point in time in (t,T]. Early exercise rules for American options on multiple assets need to take this "opportunity cost" of early exercise into consideration, as discussed in Chapter 4 (otherwise the option will be undervalued).

The properties of the early exercise regions for a dual strike option are given in Broadie and Detemple (1997). When applied to the alternative vehicles problem, these properties can be stated as follows:

E^D denotes the set of all early exercise regions for a dual strike option, it is continuous for all t in [0,T]. Subregions can be defined as:

$E_i^D = E^D \cap \{(S_t^H, S_t^{Fcell}, t) : S_t^i - D_i = \max(S_t^H - D_H, S_t^{Fcell} - D_{Fcell})\}$ for $i =_{H,Fcell}$

The following properties hold for the subregions and E^D:

(i) $(S_t^H, S_t^{Fcell}, t) \in E^D$ implies $(S_t^H, S_t^{Fcell}, s) \in E^D$ for all $t \leq s \leq T$.

(ii) $(S_t^H, S_t^{Fcell}, t) \in E_H^D$ implies $(\gamma S_t^H, S_t^{Fcell}, t) \in E_H^D$ for all $\gamma \geq 1$.

(iii) $(S_t^H, S_t^{Fcell}, t) \in E_H^D$ implies $(S_t^H, \gamma S_t^{Fcell}, t) \in E_H^D$ for all $0 \leq \gamma \leq 1$.

(iv) $(S_t^H, 0, t) \in E_H^D$ implies $S_t^H \geq B_t^H$.

(v) $(S_t^H, S_t^{Fcell}, t) \in E_H^D$ and $(\bar{S}_t^H, \bar{S}_t^{Fcell}, t) \in E_H^D$ implies $\gamma(S_t^H, S_t^{Fcell}, t) + (1-\gamma)(\bar{S}_t^H, \bar{S}_t^{Fcell}, t) \in E_H^D$ for all $0 \leq \gamma \leq 1$.

(vi) If $S_t^H - D_H = S_t^{Fcell} - D_{Fcell}$ and $min(S_t^H, S_t^{Fcell}) > 0$ and $t < T$ then $(S_t^H, S_t^{Fcell}, t) \notin E^D$.

Properties analogous to (ii) to (v) hold for subregion E_{Fcell}^D.

Property (i) states that the continuation region shrinks as maturity approaches. Property (ii) states that the exercise subregion is connected in the direction of increasing S^H. Given that early exercise is optimal by assumption at (S_t^H, S_t^{Fcell}, t),

the option value at $(\gamma S_t^H, S_t^{Fcell}, t)$ is bounded above by its early exercise value (this property is also called right connectedness). Property (iii) implies that the early exercise subregion is connected in the direction of decreasing S^{Fcell}(down connectedness). Property (iv) states that if S^{Fcell}hits 0 at any point in time, the option becomes an option on a single asset, and B_t^H (the early exercise boundary for a standard American option) becomes the sufficient statistic for early exercise. Property (v) implies that the subregions of the exercise region are convex. Property (vi) states that it is always suboptimal to exercise an American dual strike option (before maturity) if intrinsic values of both options are the same.

Early exercise rules for multidimensional options developed in Chapter 4 can be used to represent these subregions and the entire early exercise region. One possible TSMC implementation is presented below. The critical values at time t that trigger early exercise (for option to switch to a single alternative) can be modeled as[15]:

$$S_t^{*H} = D_H(\alpha_H + \beta_H \ln(T-t)) \tag{5.11}$$

$$S_t^{*Fcell} = D_{Fcell}(\alpha_{Fcell} + \beta_{Fcell} \ln(T-t)) \tag{5.12}$$

The critical values for option to switch to either one of the alternatives can be modeled as:

$$(S_t^H)^* = (S_t^H)\prime + \max\left[0, (S_t^{Fcell} - (1-\theta_t)(S_t^H)\prime)\right] \tag{5.13}$$

[15] Note that in the special cases of $\{D_H, D_{Fcell}\} = 0$, the boundary values are set to zero. This causes a low bias in the valuation of options with zero exercise price, since at the first occurance of positive savings, the option will be exercised, destroying the possibility of achieving higher payoffs in future time periods.

where $(S_t^H)\prime = D_H(\alpha_H + \beta_H \ln(T-t))$ and

$$(S_t^{Fcell})^* = (S_t^{Fcell})\prime + \max\left[0, (S_t^H - (1-\theta_t)(S_t^{Fcell})\prime)\right] \tag{5.14}$$

where $(S_t^{Fcell})\prime = D_{Fcell}(\alpha_{Fcell} + \beta_{Fcell} \ln(T-t))$

These representations can be used to generate the early exercise boundaries and show their evolutions over time. The graphs in Table 5.2 show this evolution.

The early exercise boundaries contain a lot of useful information for decision-making. For instance, following the example given in Table 5.2, at $t = 0$, the decision-maker will hold the option if the expected savings from switching to an FCV or to an HEV is below $6000. Moreover, for values above this level, as long as the expected savings of FCV and HEV are close to each other, it is still better to hold the option (even for very high saving values around $14,000 range). We observe that as T approaches, the decision-maker relaxes the conditions for holding and at maturity, even a savings level just above $4000 will trigger exercise.

5.4.6 Maturity of the Option to Switch

The maturity of the option to switch is limited due to a number of factors. Since the analysis is based on the decision making process of an individual buyer, the life expectancy of the buyer generates a natural time limit. More importantly, given an analysis horizon of a few decades, model assumptions about the nature of technological change and resource prices can easily become unrealistic. For this reason, the maturity of the option to switch was varied between 5 and 30 years during the analysis.

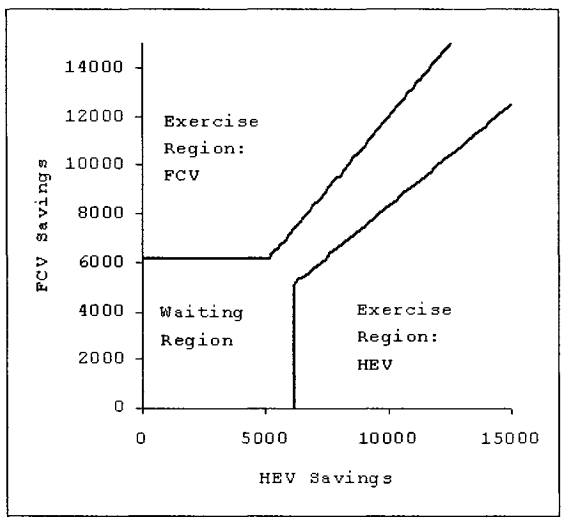

$t = 0$

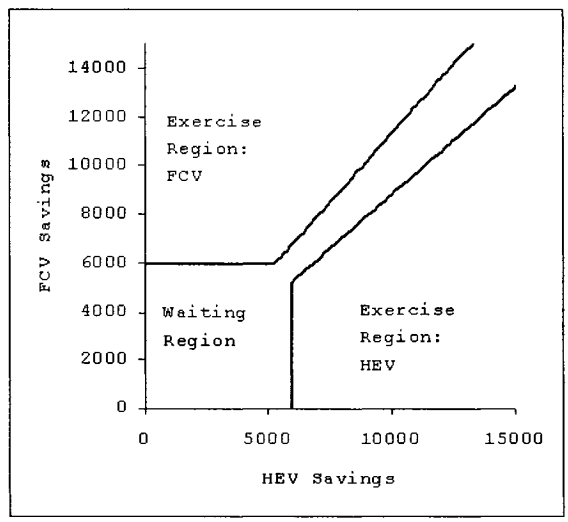

$t = 0.33T$

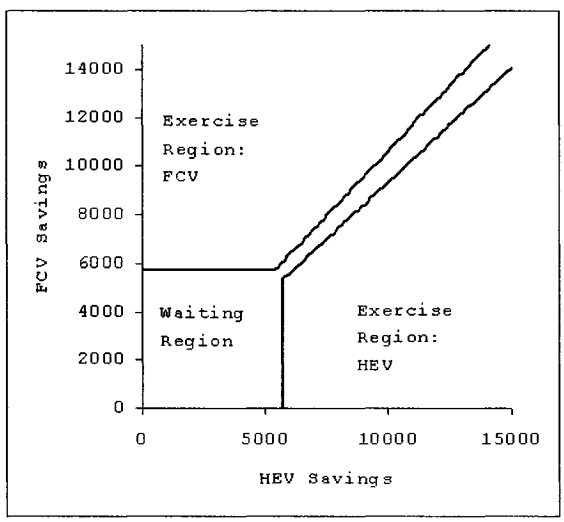

$t = 0.66T$

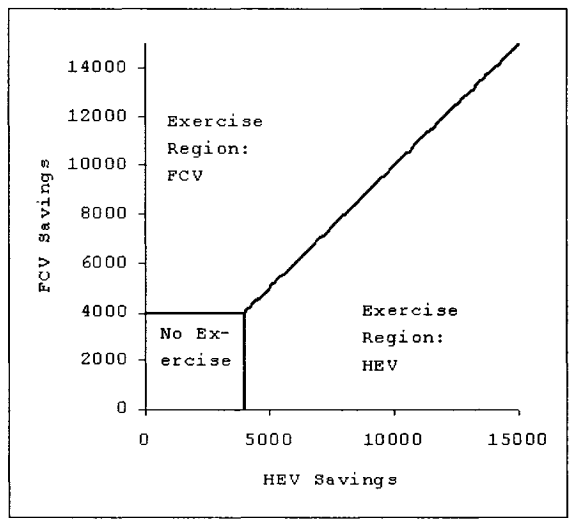

$t = T$

Table 5.2: Early exercise regions

5.5 Modeling the Price Dynamics of Oil and Hydrogen

The prices of the fuel alternatives, oil or hydrogen, are key factors in determining the expected savings in the lifecycle fuel cost of the vehicles studied. In the real options literature, oil price dynamics have been modeled using various stochastic processes. One of the earliest models for real options in oil and gas investments is the model of Paddock, Siegel and Smith (1988). In their model, Paddock, Siegel and Smith (1988) use a stochastic process based on geometric Brownian motion for the value of petroleum production. They state that the payoff from an oil and gas investment is not only based on the profits from production, but also from the capital gain on holding the remaining reserves (the convenience yield)[16]. They also discuss the behavior of the early exercise boundary over the lifetime of the option.

Smith and McCardle (1999) and Dias, Rocha, and Teixeira (2004) demonstrate the use of both GBM and mean-reverting processes for oil and gas investments. Dias and Rocha (1999) also demonstrate the use of a hybrid jump-diffusion mean-reverting stochastic process by allowing large jumps in the oil prices that follow a Poisson process. In this way, Dias and Rocha (1999) incorporate the impact of shocks (such as political crises, OPEC interventions) in the oil and gas market.

Other models for oil price dynamics include the three-factor model of Cortazar and Schwartz (2003), the model of Gibson and Schwartz (1990) that incorporate a stochastic convenience yield and the model of Robel (2001) that demonstrates the use of arithmetic mean reversion and inhomogeneous GBM.

---

[16] Dixit and Pindyck (1994) define convenience yield as "the flow of benefits (less storage costs) that the marginal stored unit provides".

Of these models, geometric Brownian motion and arithmetic mean-reversion were implemented within the TSMC real option valuation model.

5.5.1 Geometric Brownian Motion

One of the simplest models for oil prices is the geometric Brownian motion (GBM). Assuming that oil prices follow a Markovian process without any long-term "attractors" in their future trajectory, the following equation can be used to create simulation paths for oil prices:

$$P_{t+dt}^{crude} = P_t^{crude} \exp\left(r - \delta - \frac{\sigma^2}{2}dt + \sigma\epsilon\sqrt{dt}\right) \tag{5.15}$$

This is the same equation as equation 4.28 (Chapter 4). However, this time P_t is the spot price of one barrel of crude oil, and σ is the volatility parameter. In Equation 4.28, δ represents the continuous dividend that a stock provides. In the case of a storable commodity (such as oil or metals), δ denotes the net marginal convenience yield from storage.

5.5.2 Arithmetic Mean-Reverting Process

An alternative stochastic process for oil prices is the mean-reverting one. Dixit and Pindyck (1994) argue that if the prices of commodities like oil are believed to be drawn back to their long-run marginal costs, then a mean-reverting process would be more suitable than a GBM approach. The properties of this process are as follows[17]:

[17] Also called the Ornstein-Uhlenbeck process

$$dP = \eta(\bar{P} - P)dt + \sigma_p \epsilon \sqrt{dt} \tag{5.16}$$

where η is the reversion rate to the long-run price of oil, $\bar{P}$, and σ_p is the volatility parameter.

The following differential equation governs the value of an option $F(P,t)$, written on an underlying asset P, that has mean-reverting properties (Dias et al., 2004):

$$\frac{1}{2}\sigma_p^2 P^2 \frac{\partial^2 F}{\partial P^2} + \left(r - \rho - \frac{\eta(\bar{P} - P)}{P}\right) P \frac{\partial F}{\partial P} + \frac{\partial F}{\partial t} - rF = 0 \tag{5.17}$$

For the Monte Carlo simulation of mean-reverting paths, Dias (2004) proposes the following equation:

$$\ln(P_{t+\Delta t}) = P_t e^{(-\eta \Delta t)} + \ln(\bar{P})(1 - e^{-\eta \Delta t}) + \sigma_p \left(\epsilon \sqrt{\frac{1 - e^{(-2\eta \Delta t)}}{2\eta}} - \sigma_p \left(\frac{1 - e^{(-2\eta \Delta t)}}{4\eta}\right)\right) \tag{5.18}$$

Given $P_0, \bar{P}, \eta, r, \delta, \sigma$ and Δt increment, mean-reverting simulation paths of the price of oil (or hydrogen) at any point in time can be created.

5.5.3 Calculating Retail Gas Prices

The following simple regression equation is used to convert crude oil prices to "prices at the pump", the average price (in $) that a consumer pays for one U.S. gallon of oil :

$$P_t^{pump} = \left(60.205 + 3.358 P_t^{crude}\right) 0.01 \tag{5.19}$$

The parameters in this regression equation were computed using a dataset of crude oil prices and retail prices of gas in the U.S. spanning the last 10 years (dataset was obtained from U.S. Energy Information Administration, the adjusted R-squared of the regression is 93%). The prices for hydrogen do not require a regression equation, since gallon of gasoline equivalent (gge) prices were used in the stochastic processes.

5.6 Modeling the Technical Uncertainty

The technical uncertainty associated with the future performance of FCVs are modeled using an S-curve framework. Technology S-curves are frequently used in technology forecasting, and historical analysis of past R&D data suggest that there is a regular pattern in the trajectory of incremental R&D projects (Burgelman, 1996; Twiss, 1992). This pattern can be observed when a technical performance metric (such as fuel cell efficiency) or economic performance metric (such as cost-efficiency of hydrogen generation) is plotted against time. Generally there is a slow initial growth followed by exponential growth and eventual decrease in the rate of progress as a physical limit is reached (Twiss, 1992; Boer, 1999).

Davis and Owens (2003) develop a real options model to investigate a different, but related problem: renewable electric R&D projects. They argue that when modeling the impact of R&D, one should use stochastic processes such that deviations from the current values are permanent. In other words, once the technology is created, the benefits it brings are never lost. Compared to the price dynamics

of natural resources, this is a fundamental difference, and the impact of R&D on the alternative vehicles problem should be modeled by taking this observation into consideration.

In the TSMC model, as can be seen in Figure 5.2, the future performance of vehicles were simulated using technology S-curves. The occurrence of incremental increases in the fuel economy were modeled as random events following a Poisson distribution. For each time period t, the fuel efficiency of a given vehicle can increase by a certain "jump" (J_i) with a probability $p(J_i)$. It is assumed that there can be no jumps down, since the technology either stays at the same maturity level at each time period, or it improves.

J_i is expressed in terms of the percentage increase in fuel efficiency. The jump size of each vehicle type i can depend on a number of factors, including the amount of progress already made, the nature of the technological challenge, and the amount of funds allocated for R&D activities.

Since $p(J_i)$ follows a Poisson process, the probability of different number of jumps occurring in each period can be calculated using the Poisson distribution formula:

$$p(J_i) = \frac{\lambda^{\phi}}{\phi!} e^{-\lambda} \tag{5.20}$$

where λ is the mean number of jumps in one time period, and ϕ is the number of jumps that the probability is calculated for. Incremental innovations in the performance of a technology can be modeled with relatively lower J_i values with higher λ

values. Conversely, radical innovations can be modeled with higher J_i values with smaller mean number of jumps.

Based on the Poisson distribution formula, the following process is developed in order to model the evolution of the fuel efficiencies (E) over time:

$$E_{t+1} = E_t dj \tag{5.21}$$

$$dj = \begin{Bmatrix} 1 & p(J_i) \mid \phi = 0 \\ (1+J_i)^1 & p(J_i) \mid \phi = 1 \\ (1+J_i)^2 & p(J_i) \mid \phi = 2 \\ (1+J_i)^3 & p(J_i) \mid \phi = 3 \end{Bmatrix} \tag{5.22}$$

It is possible that more than one jump can happen in a given time period. In this case, the increase in efficiency is assumed to be compounded. The probability of more than three jumps occurring per time period is very low given a λ range of $[0.5, 1.5]$. Therefore, the possibility of more than three jumps per time period has not been included in the model.

Furthermore, due to the S-curve nature of technological evolution, λ decays over time. In other words, the mean number of jumps in one time period decreases as the technology matures. This decay is modeled using an exponential decay formula based on time:

$$\lambda_t = \lambda_0 e^{-ta} \tag{5.23}$$

where $a \in [0, 1]$ is the decay factor and λ_0 is the starting value for mean number of jumps per time period (a variable that is dependent on the nature of the technology in question).

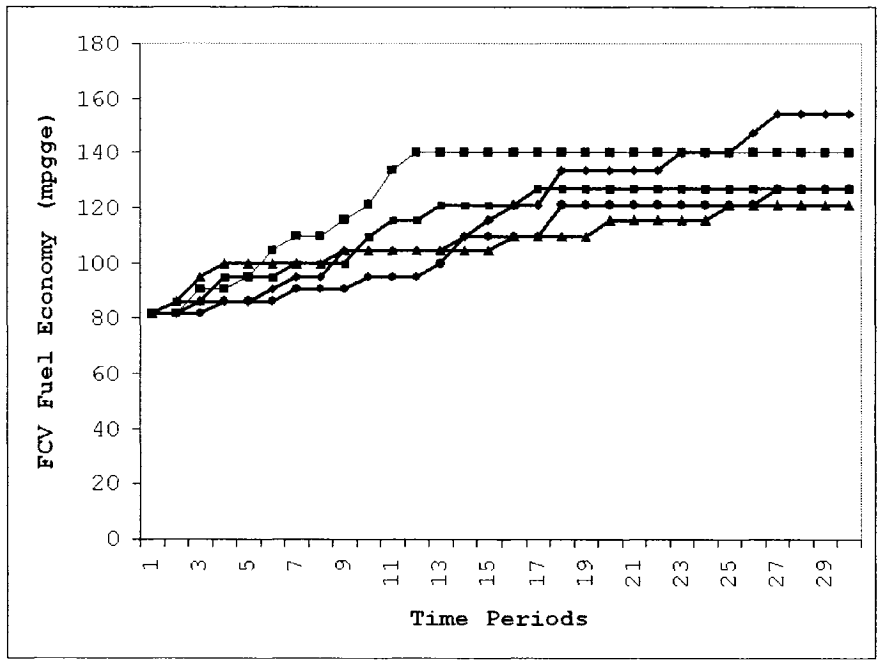

Figure 5.2: Sample Technology S-curves ($\lambda = 1$, $J_i = 0.05$)

Starting from their current fuel economy levels, the fuel efficiency of vehicles are assumed to follow a Poisson jump process. Figure 5.2 shows an example of simulated random paths for an FCV with $J_{FCV} = 0.05$, $\lambda_0 = 1$, $a = 0.1$, and $E_0 =$ 82 mpgge (miles per gallon of gasoline equivalent). As the fuel economy of FCVs increases over time, the expected lifetime savings from switching to FCVs increase as well (all other factors remaining constant).

Technology S-curves have been incorporated into Monte Carlo simulation using the cumulative Poisson probabilities. Below is an example of the cumulative probabilities of increase in fuel efficiency (in a single time period) for $\lambda = 1$ and $J_i = 0.1$. Percentage increase in fuel efficiency over the preceding time period are also given.

ϕ	$p(J_i)$	Cumulative Probability	% increase	Interval
0	0.3678	0.3678	0%	(0, 0.3678]
1	0.3678	0.7356	10%	(0.3678, 0.7356]
2	0.1839	0.9195	21%	(0.7356, 0.9195]
3	0.0613	0.9808	33.1%	(0.9195, 1)

In order to determine if an increase occurs at each time step, a random draw between 0 and 1 is made. Then this random number is compared to the intervals corresponding to the cumulative probability of different number of jumps. For instance, if the random number is 0.43, only 1 jump takes place. If the random number is 0.33, no jump occurs. If the random number is greater than 0.9195, 3 jumps take place. In this fashion, by changing the λ and J_i values, cumulative probabilities and expected increase in fuel efficiencies can be modeled.

It is important to note that the increases in fuel efficiencies are effectively "frozen" when a new FCV is purchased. In other words, the consumer is locked into the level of fuel efficiency of the vehicle at the time of purchase, since future increases in fuel efficiency can only be incorporated into new vehicles in each time period (assuming that no "upgrades" after sales are possible). Therefore, if the decision maker has the expectation that fuel efficiencies will increase considerably in future time periods, he/she may decide to wait longer before switching.

5.7 Results

There are many variables in the alternative vehicles problem, some having a stronger impact on the option values than others. By using the TSMC method together with sensitivity analysis, these leading variables can be determined.

Model Parameter	Value
Risk-free rate of interest r (p.a.)	5%
Annual volatility of oil prices σ_P	20%
Convenience yield of oil δ_P	10%
Long-term average oil price P (US$/barrel)	[$20 - $100]
Reversion rate of oil prices, η_P	0.1
Initial oil price (US$/barrel)	$40
Number of time steps used in calculations	30
Annual volatility of hydrogen prices (σ_H)	10%
Convenience yield of hydrogen (δ_H)	10%
Long-term average hydrogen price H (gge)	[$1 - $5]
Reversion rate of hydrogen prices, η_H	0.1
Initial hydrogen price (gge)	$2.2
Fuel economy of ICEV (mpg) / retail price of ICEV	27/$20K
Fuel economy of HEV (mpg) / retail price of HEV	52/$28K
Fuel economy of FCV (mpgge) / retail price of FCV	82/$100K

Table 5.3: Base case scenario parameters

In this section, the results of this sensitivity analysis will be presented along with the option values. The base case scenario for the model variables is given in Table 5.3. The values in the base case scenario were determined based on the current literature on passenger vehicles and other sources. Specifically, the FCV and HEV data is from Thomas et al. (2000), Weiss et al. (2003), and Lipman and Delucchi (2003). The initial fuel efficiency (mpg) values and hydrogen prices are from Ogden, Williams and Larson (2004). Parameters for mean-reverting oil price dynamics are from Dias (2004). For other parameter values, see Dunn (2001), IEA's World Energy Investment Outlook (2003), and Fairley (2004).

5.7.1 Option to switch from ICEV to HEV

This option's exercise price is the incremental price difference between buying a new ICEV vs. buying an HEV. In the base case scenario, this incremental difference is $8,000 (derived from the retail prices of hybrid vehicles in 2005). It is

Incremental Price Difference	Maturity of Option - years (T)	Option Value ($) 95% C.I.	$\{\alpha, \beta\}$
$8000	5	0	-
$8000	15	[0.16, 0.74]	{1.17, 0.11}
$8000	30	[2.43, 3.73]	{1.16, 0.01}
$6000	5	0	-
$6000	15	[1.31, 3.11]	{1.21, 0.06}
$6000	30	[5.27, 8.13]	{1.2, 0.1}
$4000	5	0	-
$4000	15	[32.86, 35.04]	{1.16, 0}
$4000	30	[32.19, 37.50]	{1.2, 0.08}
$2000	5	[21.19, 31.56]	{1.15, 0}
$2000	15	[960.88, 981.07]	{1.15, 0}
$2000	30	[1107.35, 1137.10]	{1.18, 0}
$0	5	[1753.71, 1759.77]	{1.25, 0.1}
$0	15	[2972.37, 2989.98]	{1.21, 0.1}
$0	30	[3126.59, 3136.41]	{1.22, 0.09}

Table 5.4: Results for Option to Switch from ICEV to HEV

expected that the prices of HEVs will decrease in the future, making the option to switch more valuable (all other variables remaining constant).

Table 5.4 presents the results for option to switch from an ICEV to a HEV. For this option, oil prices are assumed to follow Geometric Brownian Process, with the initial price per barrel being $40. Other option parameters are as follows, $r = 0.05$, $\sigma_P = 0.2$, $\delta_P = 0.1$. 30 time steps were used for the calculations. ICEV and HEV mpg values are 27 and 52, respectively. The incremental price difference is varied between $8000 and $0 to analyze the impact of retail prices of the vehicles on option values. 95% confidence bins for option values are given in the table (the calculation of the confidence bins is described in detail in Chapter 4). Tabu Search was run with 10,000 Monte Carlo paths for 15 iterations, followed by 5 iterations with 50,000 paths to increase the accuracy of the results. Solution times for each row are about 1000 seconds using a 1.2 GHz Celeron processor. Initial values for $\{\alpha_{initial}, \beta_{initial}\} = \{1.2, 0.1\}$.

5.7.2 Option to switch from ICEV to FCV

This option's exercise price is the incremental price difference between buying a new ICEV vs. buying an FCV. In the base case scenario, this incremental difference is $80,000 as of 2005. However, at this price level, the option price has zero value. Moreover, if the exercise price of the switching option is higher than the $8000 range, then the value of the option is negligible. For this reason, in this section, results for incremental price differences between $0 and $8000 will be presented. Although this retail price level is unrealistic today, it is expected that the prices of FCVs will decrease significantly in the future, making the option to switch more valuable (all other variables remaining constant).

Table 5.5 presents the results for option to switch from an ICEV to a FCV. For this option, oil prices are assumed to follow a geometric Brownian process, with the initial price per barrel being $40. Other option parameters are as follows, $r = 0.05$, $\sigma_P = 0.2$, $\delta_P = 0.1$. 30 time steps were used for the calculations. Hydrogen prices are assumed to follow a mean-reverting process, with $\sigma_H = 0.1$, $\delta_H = 0.1$, $\eta_H = 0.1$, and $\bar{H} = \$2$. The initial price of hydrogen (per gge) is $2.2 (this is the current price level of hydrogen generated from natural gas). ICEV and FCV mpg values are 27 and 82, respectively (and no further efficiency gains are assumed for both types of vehicles). The incremental price difference is varied between $8000 and $0 to analyze the impact of retail prices of the vehicles on option values.

Tabu Search was run with 10,000 Monte Carlo paths for 15 iterations, followed by 5 iterations with 50,000 paths to increase the accuracy of the results. Solution

Incremental Price Difference	Maturity of Option - years (T)	Option Value ($) 95% C.I.	$\{\alpha, \beta\}$
$8000	5	0	-
$8000	15	[16.08, 19.87]	{1.17, 0.03}
$8000	30	[27.06, 36.95]	{1.22, 0.09}
$6000	5	0	-
$6000	15	[46.76, 55.66]	{1.16, 0.1}
$6000	30	[89.55, 101.43]	{1.17, 0}
$4000	5	[0.38, 1.68]	{1.2, 0.07}
$4000	15	[328.80, 359.46]	{1.16, 0}
$4000	30	[401.94, 440.96]	{1.17, 0}
$2000	5	[324.25, 339.06]	{1.12, 0}
$2000	15	[1616.62, 1627.41]	{1.12, 0.07}
$2000	30	[1798.36, 1825.49]	{1.14, 0}
$0	5	[2204.77, 2217.58]	{1.17, 0.1}
$0	15	[3601.59, 3636.57]	{1.21, 0.09}
$0	30	[3790.16, 3818.86]	{1.21, 0.06}

Table 5.5: Results for Option to Switch from ICEV to FCV

times for each row are about 1300 seconds using a 1.2 GHz Celeron processor. Initial values for $\{\alpha_{initial}, \beta_{initial}\} = \{1.2, 0.1\}$.

5.7.3 Option to switch from ICEV to HEV or FCV

When the prospective buyer is given the option to switch to an HEV or an FCV from an ICEV, the option value is expected to be higher than the maximum of the single switching options, but less than the sum of these options (due to the nature of options on multiple assets). Moreover, in this option, the dual-strike nature of the option to switch can also be observed: since the retail prices of HEVs and FCVs are likely to be different, the exercise prices of switching are different as well.

Table 5.6 presents the results for this multidimensional option. Again, oil prices are assumed to follow a geometric Brownian process, with the initial price per barrel being $40. Other option parameters are as follows, $r = 0.05$, $\sigma_P = 0.2$, $\delta_P = 0.1$, $T = 15$ years. 30 time steps were used for the calculations. Hydrogen

$T = 15$ years	Incremental Price Difference - HEV					
Incremental Price Difference - FCV	$8000		$4000		$0	
	Option Value ($)	$\{\theta\}$	Option Value ($)	$\{\theta\}$	Option Value ($)	$\{\theta\}$
$8000	[15.36, 21.56]	{0.03}	[32.49, 35.34]	{-0.02}	[2973.95, 2984.32]	{0.03}
$4000	[329.86, 349.03]	{0.02}	[325.59, 347.35]	{-0.04}	[2976.75, 2981.68]	{-0.01}
$0	[3608.06, 3624.60]	{-0.01}	[3610.77, 3628.15]	{0.06}	[3669.32, 3678.96]	{0.03}

Table 5.6: Results for Option to Switch from ICEV to HEV or FCV

prices are assumed to follow a mean-reverting process, with $\sigma_H = 0.1$, $\delta_H = 0.1$, $\eta_H = 0.1$, and $\tilde{H} = \$2$. The initial price of hydrogen is $2.2. ICEV, HEV and FCV mpg values are 27, 52 and 82, respectively (and no further efficiency gains are assumed for all types of vehicles).

Tabu Search was run with 10,000 Monte Carlo paths for 25 iterations, followed by 5 iterations with 50,000 paths to increase the accuracy of the results. Solution times for each row are about 430 seconds using a 1.2 GHz Celeron processor (assuming that both ICEV to HEV, and ICEV to FCV switching options have been evaluated and $\alpha_{best}, \beta_{best}$ are known for each individual option).

By comparing the midpoints of upper and lower bounds, one can get an idea about the percentage differences between the value of this option and the value of option to switch to a single alternative. Note that one would expect this particular option's value to be higher than the corresponding values in Tables 5.4 and 5.5, since this option has a more valuable payoff than the previous two. However, this is not always the case as can be seen in Table 5.6. This disparity is most likely a

result of the simulations used in the TSMC method and the true value of this option should be more than the individual values of the previous two options, but less than the sum of them.

5.7.4 Impact of Technological Advancements

One other important factor in the alternative vehicles problem is the future trajectory of technological advancements in fuel efficiency. As discussed in Section 5.6, through R&D, fuel efficiencies of different types of vehicles can be increased further.

In the case of the FCV, the starting mpgge value of 82 (base case scenario value) can significantly increase over the next few decades. One way of modeling this technical uncertainty is using the technology S-curve model. Parameters of this model can be changed to create future trajectories based on different assumptions. For instance, incremental innovation can be modeled as more frequent, but lower jumps in the fuel efficiency of vehicles. Radical innovation, on the other hand, can be modeled as less frequent, but higher jumps.

Given different values of λ (mean number of jumps for each time period), and different J_i (jump size) values, the Poisson probabilities of technological advancement that can take place in a given year can be calculated. By generating random numbers between 0 and 1, and by comparing these random draws with Poisson probabilities, future fuel efficiency levels can be simulated. Table 5.7 presents the results of option to switch from ICEV to FCV when S-curve dynamics are included in the analysis.

$T = 15$ years	J_i (jump size per period)					
λ	5%		10%		15%	
	Option Value ($) 95% C.I.	$\{\alpha, \beta\}$	Option Value ($) 95% C.I.	$\{\alpha, \beta\}$	Option Value ($) 95% C.I.	$\{\alpha, \beta\}$
0.5	[346.45, 376.35]	{1.15, 0}	[383.84, 402.88]	{1.17, 0}	[420.33, 443.89]	{1.14, 0.02}
1	[371.45, 392.98]	{1.19, 0}	[445.18, 462.36]	{1.16, 0.02}	[517.96, 538.17]	{1.16, 0.03}
1.5	[405.97, 418.03]	{1.15, 0.03}	[501.06, 524.21]	{1.18, 0.02}	[596.46, 618.19]	{1.16, 0.05}

Table 5.7: Results for Option to Switch from ICEV to FCV - Technology S-curve model

Base case parameter values are used for obtaining the results in Table 5.7 (except for the incremental price difference, which is $4000 for this case). J_i (jump size per period) is given in percentages (the percentage increase in fuel efficiency if a jump occurs).

Tabu Search was run with 10,000 Monte Carlo paths for 15 iterations, followed by 5 iterations with 50,000 paths to increase the accuracy of the results. Solution times for each cell are about 740 seconds using a 1.2 GHz Celeron processor.

Table 5.7 can be useful to compare the impact of incremental versus radical innovation. In the incremental innovation scenario, fuel efficiencies are expected to increase in small increments, but with a relatively high probability of increase per time period (e.g., $\lambda = 1.5$ and $J_i = 0.05$). In the radical innovation scenario, fuel efficiencies are expected to increase in larger jumps, but with relatively small probability of increase per time period (e.g., $\lambda = 0.5$ and $J_i = 0.15$). Therefore, the lower left side of Table 5.7 represents the incremental innovation scenario, while the upper right side represents the impact of radical innovation. Not surprisingly,

the highest option value in Table 5.7 (lower right) corresponds to the highest λ and J_i values.

As it can be seen in this table, including S-curve dynamics increases the value of the option: in both incremental and radical innovation scenarios, option values (shown in Table 5.7) are about 25% higher than the value of the option to switch from ICEV to FCV without S-curve dynamics (Table 5.5, row 9) when $T = 15$ years and the incremental price difference is $4000.

It is important to note that there are two main effects of increasing λ:

(a) The higher the mean number of jumps (λ), the higher the expected savings will be. This would make the FCV switching option more valuable (all other factors remaining constant) and increase the likelihood of exercising the switching option.

(b) Higher mean number of jumps (λ) can also translate into an incentive to wait, since more savings can be achieved if the decision maker postpones the switching decision. This can also have an impact on the likelihood of exercising the option.

The current model does not separate the impact of these two effects.

5.7.5 Impact of Hydrogen Generation from Renewable Energy Sources

Generating hydrogen for FCVs in a cost-effective and environmentally friendly way is one of the greatest challenges in the automotive and energy industries. Hydrogen generated from renewable energy sources has a very low environmental

footprint, but currently this type of production is much more costly than generating hydrogen from fossil fuels (which cost around $2.2 per gge today). In order to determine the impact of hydrogen prices on the value of option to switch from ICEV to FCV, a sensitivity analysis was conducted. Different values of long-term average hydrogen price ($\bar{H}$) and the reversion rate of hydrogen prices (η_H) yield drastically different option values.

For obtaining the results in Table 5.8, base case parameter values were used (except for the incremental price difference, which is zero in this case). Furthermore, the cost of hydrogen (gge) at time zero is assumed to be $5, closer to the current cost of hydrogen generation from renewable sources.

Tabu Search was run with 10,000 Monte Carlo paths for 15 iterations, followed by 5 iterations with 50,000 paths to increase the accuracy of the results. Solution times for each cell are about 640 seconds using a 1.2 GHz Celeron processor.

As can be seen in the table, even in the case of zero exercise price, when the hydrogen is obtained from renewable sources the value of the option decreases considerably compared to the case of hydrogen generation from natural gas (Table 5.5, row 15).

5.7.6 Mean-reverting Oil Prices

Another interesting case occurs when the Geometric Brownian Motion assumption for oil prices is relaxed. For this analysis, oil prices are assumed to follow a mean-reverting process with the reversion rate of oil prices, $\eta_P = 0.1$. All other parameters are as stated in the base case scenario, except for the incremental price difference

$T = 15$ years	η_H (reversion rate of hydrogen prices)			
H (long-term average hydrogen price in gallon of gasoline equivalent)	0.3		0.6	
	Option Value ($) 95% C.I.	$\{\alpha, \beta\}$	Option Value ($) 95% C.I.	$\{\alpha, \beta\}$
$5	[578.27, 596.47]	{1.16, 0.12}	[563.60, 580.53]	{1.17, 0.13}
$4	[1047.82, 1077.07]	{1.17, 0.1}	[1166.52, 1193.69]	{1.21, 0.16}
$3	[1763.05, 1797.66]	{1.2, 0.1}	[2073.84, 2104.42]	{1.18, 0}
$2	[2674.60, 2694.76]	{1.19, 0.14}	[3103.25, 3138.39]	{1.2, 0.1}
$1	[3727.95, 3748.27]	{1.12, 0.1}	[4224.78, 4252.31]	{1.18, 0.1}

Table 5.8: Results for Option to Switch from ICEV to FCV - Different Hydrogen Prices

between an FCV and an ICEV, which is assumed to be $4,000. The maturity of the option (T) is 15 years.

Figure 5.3 plots the values of the option to switch from an ICEV to an FCV for different long-term average prices of oil and hydrogen. To obtain each data point on the graph, 15 Tabu Search iterations were run with 10,000 Monte Carlo paths. It took about 300 seconds (using a 1.2 GHz Celeron processor) to calculate each data point.

The values on x- and y-axes are the long-term average prices of oil and hydrogen, respectively. As it can be seen in the graph, oil prices have a very significant impact on the option values, while the impact of hydrogen prices is not as drastic.

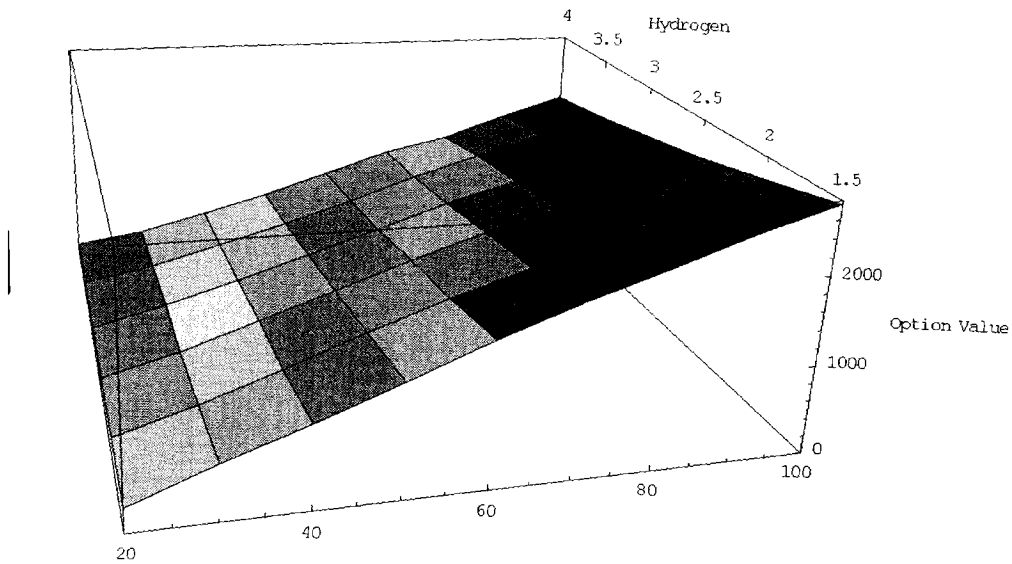

Figure 5.3: Option to sw ›m ICEV to FCV - mean reverting oil and hydrogen prices

5.8 Discussion and Conclusions

All of the option values given in Section 5.7 are calculated by taking into account the decision making process of an individual buyer in the automotive market. As such, these values reflect the dollar amount of benefits the options of buying an HEV or an FCV bring to an individual consumer. It is important to emphasize that these benefits can occur if and only if, these vehicles are commercially available to the average consumer. It is also important to note that, the option values only capture the impact of fuel economy. On top of these values, the "value package" of an automobile includes other components as well (such as safety, convenience, prestige, etc.). These other components are not captured in the results presented in this chapter.

From the perspective of automotive companies, reported option values can be useful for making investment decisions in alternative vehicles. One possibility is to compute aggregate option values from the reported option values of individual consumers. Even a rough calculation (such as multiplying individual option values with the expected number of consumers who may make a switching decision) can give an indication of the aggregate option value of having HEVs and FCVs in the market. Such a figure can be an interesting benchmark to compare the amount of R&D, production, marketing and distribution costs required to bring such vehicles into the market.

As it can be seen in the result tables, even in the most favorable cases for FCVs (e.g., competitive retail pricing for FCVs and lower hydrogen prices), the option value is around US$ 3000. When the retail prices of such vehicles are considered, this may not look like a very significant amount at first (about 15% of the retail price of an average ICEV). However, these are values for a single consumer, and again, the aggregate value from all potential consumers can be significant for the automotive market.

When the environmental and energy security benefits of alternative vehicles are considered, the associated benefits of ownership are likely to be higher. Thus, moving beyond the direct fuel-efficiency benefits, such vehicles can offer higher indirect benefits to the society. One way of realizing these indirect benefits could be by putting into place certain incentive mechanisms to stimulate the demand for alternative vehicles.

It is clear that the retail pricing of HEVs and FCVs is one of the most important factors determining the value of the options to switch. This is hardly surprising, since retail prices determine the exercise prices of the options. Therefore any policy that has a favorable impact on retail prices (such as tax incentives) deserves to be studied in detail.

Another important factor is hydrogen prices. In order to maximize the environmental benefits of FCVs, the option of hydrogen generation from renewable energy sources has to be considered. Although this alternative is currently not cost-competitive compared to hydrogen generation from natural gas, indirect benefits and costs need to be considered as well. Again, policy actions that target making hydrogen from renewable energy sources more competitive can help change the fuel mix in the transportation sector.

Finally, it is critical to realize that the base case scenario assumptions, model parameters and other model components presented in this chapter can be inaccurate. The main contribution of the real options model is not the accuracy of the model inputs and outputs (although these have to be reasonably accurate) but the insight that the analysis process brings. By testing different assumptions, studying alternative future scenarios and comparing the advantages and disadvantages of different types of vehicles, individual consumers, automotive and energy companies and policy makers can make healthier decisions.

Some future research areas are identified as follows: studying the value of flexibility that can be achieved by distributed power generation (e.g., hydrogen generation in households using solar panels); network effects in a distributed energy

system; option value of dual-fuel engines (e.g., ICEV/FCV hybrid vehicles); and assessing the conditions of feasibility for distributed hydrogen production.

Chapter 6
Conclusion

Moving from a deterministic optimization problem to stochastic ones, this thesis has demonstrated the versatility of tabu search as an optimization technique.

One common thread throughout the thesis is the power of a multidisciplinary approach for solving decision-making problems. In Chapter 3, orbital mechanics and operations research were used in conjunction to create a practical tool for the analysis of on-orbit servicing systems. In Chapter 4, operations research met finance to tackle a well-established valuation problem. In Chapter 5, techniques of operations research were used in a real options context. Although conducting multidisciplinary research is a challenging task, the benefits far outweigh the burdens of this challenge.

One other common thread is the use of a very flexible programming environment, C++, for building dedicated software tools. Compared to standard software packages, such as MS Excel, the advantages of building dedicated software applications in C++ are clear: increased accuracy and speed, absolute control on the contents of the application, reusable modules that enable portability of most used routines from one application to another.

The topics that were covered in this thesis can only be the first few steps of a more comprehensive research agenda. There are still many areas that need in-depth study and analysis. The applicability of the Tabu Search Monte Carlo method

in other problem domains which require a hybrid simulation/optimization approach (such as bioinformatics) needs to be investigated.

This thesis was an attempt to develop an innovative way for modeling and solving decision making problems. The main emphasis was on using a versatile metaheuristic within a dedicated software application environment. It is the author's hope that, within a multidisciplinary research setting, this approach can provide a powerful mix that can be used to tackle tough decision making problems.

References

[1] ANDERSEN, L. B. A simple approach to the pricing of bermudan swaptions in the multi-factor libor market model. *Journal of Computational Finance 3* (2000), 5–32.

[2] ANDERSEN, L. B., AND BROADIE, M. A primal-dual simulation algorithm for pricing multi-dimensional american options. *Management Science 50*, 9 (2004), 1222–1234.

[3] ARISTOTLE. *Politics*. Hackett Publishing, Indianapolis, IN, 1998.

[4] BACHELIER, L. Théorie de la spéculation. Tech. Rep. 3, Annales Scientifiques de l'École Normale Supérieure, 1900.

[5] BARONE-ADESI, G., AND WHALEY, R. E. Efficient analytic approximation of american option values. *Journal of Finance 42* (1987), 301–320.

[6] BARRAQUAND, J., AND MARTINEAU, D. Numerical valuation of high dimensional multivariate american securities. *The Journal of Financial and Quantitative Analysis 30*, 3 (1995), 383–405.

[7] BASS, F. A new product growth model for consumer durables. *Management Science 15* (1969), 215–227.

[8] BASSO, A., NARDON, M., AND PIANCA, P. Discrete and continuous time approximations of the optimal exercise boundary of american options, 2002. Working Paper, Dept. of Applied Mathematics University of Venice, Dorsoduro, Italy.

[9] BELLMAN, R. *Dynamic Programming*. Princeton University Press, Princeton, New Jersey, 1957.

[10] BLACK, F., AND SCHOLES, M. The pricing of options and corporate liabilities. *Journal of Political Economy 81* (1973), 637–654.

[11] BLUM, C., AND ROLI, A. Metaheuristics in combinatorial optimization: Overview and conceptual comparison. Tech. Rep. 2001-13, IRIDIA, October 2001.

[12] BOER, F. P. *The Valuation of Technology : Business and Financial Issues in R and D*. Wiley Operations Management Series for Professionals. Wiley, New York, 1999.

[13] BOYLE, P. Options: A monte carlo approach. *Journal of Financial Economics 4*, 4 (1977), 323–338.

[14] BOYLE, P. A lattice framework for option pricing with two state variables. *Journal of Financial and Quantitative Analysis 23*, 1 (1988), 1–12.

[15] BOZKAYA, B. E., ERKUT, G., AND LAPORTE, G. A tabu search heuristic and adaptive memory procedure for political districting. *European Journal of Operational Research 144*, 1 (2003), 12–26.

[16] BRENNAN, M. J., AND SCHWARTZ, E. The valuation of american put options. *Journal of Finance 32*, 2 (1977), 449–462.

[17] BRENNAN, M. J., AND SCHWARTZ, E. S. Evaluating natural resource investments. *The Journal of Business 58*, 2 (1985), 135–157.

[18] BRITT, R. R. How hubble has survived a decade of impacts. *Space News* (2002).

[19] BROADIE, M., AND DETEMPLE, J. B. The valuation of american options on multiple assets. *Mathematical Finance 7*, 3 (1997), 241–286.

[20] BROADIE, M., AND DETEMPLE, J. B. Option pricing:valuation models and applications. *Management Science 50* (2004), 1145–1177.

[21] BROADIE, M., AND GLASSERMAN, P. A stochastic mesh method for pricing high-dimensional american options. *Journal of Computational Finance 7*, 4 (2004), 35–72.

[22] BURGELMAN, R. A., MAIDIQUE, M. A., AND WHEELWRIGHT, S. C. *Strategic Management of Technology and Innovation*, 2nd ed. Irwin, Chicago, 1996.

[23] CAMPBELL, C., AND LAHERRERE, J. The end of cheap oil. *Scientific American 278*, 3 (1998), 76–83.

[24] CARR, P. P., JARROW, R., AND MYNENI, R. Alternative characterizations of american put options. *Mathematical Finance 2*, 4 (1992), 87–106.

[25] CHILDS, P. D., OTT, S., AND TRIANTIS, A. Capital budgeting for interrelated projects: A real options approach. *Journal of Financial and Quantitative Analysis 33*, 3 (1998), 305–334.

[26] CORDEAU, J.-F., GENDREAU, M., AND LAPORTE, G. A tabu search heuristic for periodic and multi-depot vehicle routing problems. *Networks 30* (1997), 105–119.

[27] CORDEAU, J.-F., AND LAPORTE, G. Tabu search heuristics for the vehicle routing problem. Tech. Rep. G-2002-15, GERAD, Montreal, QC, 2002.

[28] CORTAZAR, G., AND SCHWARTZ, E. Implementing a stochastic model for oil futures prices. *Energy Economics 25* (2003), 215–238.

[29] COX, J. C., ROSS, S. A., AND RUBINSTEIN, M. Option pricing: A simplified approach. *Journal of Financial Economics 7* (1979), 229–263.

[30] COYLE, L., AND YANG, J. Analysis of the ssap method for the numerical valuation of high-dimensional multivariate american securities. *Algorithmica 25* (1999), 75–98.

[31] DAVID, L. Satellite savior: A look at an earth-orbit rescue vehicle. *Space News* (2003).

[32] DAVIS, G. A., AND OWENS, B. Optimizing the level of renewable electric r and d expenditures using real options analysis. *Energy Policy 31* (2003), 1589–1608.

[33] DAWSON, C. Fuel cells: Japan's carmakers are flooring it. *Business Week*, 3813 (2002), 50.

[34] DEFFEYES, K. S. *Hubbert's Peak : The Impending World Oil Shortage*. Princeton University Press, Princeton, NJ, 2001.

[35] DIAS, M. A. G. Selection of alternatives of investment in information for oilfield development using evolutionary real options approach, 2001. Working Paper, Real Options in Petroleum website, www.puc-rio.br/marco.ind/.

[36] DIAS, M. A. G. Monte carlo simulation of stochastic processes, 2004. Real Options in Petroleum Website www.puc-rio.br/marco.ind/.

[37] DIAS, M. A. G., AND ROCHA, K. Petroleum concessions with extendible options using mean reversion with jumps to model oil prices, 1999. Working Paper, Betrobras, Rio de Janeiro.

[38] DIAS, M. A. G., ROCHA, K., AND TEIXEIRA, J. P. The optimal investment scale and timing: A real option approach to oilfield development, 2004. Presented in the 8th Annual International Real Options Conference Montréal Canada, June 17-19.

[39] DIXIT, A. K., AND PINDYCK, R. S. *Investment under Uncertainty.* Princeton University Press, Princeton, N.J., 1994.

[40] DONATI, A. V., MONTEMANNI, R., CASAGRANDE, N., RIZZOLI, A. E., AND GAMBARDELLA, L. M. Time dependent vehicle routing problem with a multi ant colony system, 2003. Working Paper, Istituto Dalle Molle di Studi sull'Intelligenza Artificiale (IDSIA), Manno, Switzerland.

[41] DUNN, B. P. High-energy orbit refueling for orbital transfer vehicles. *Journal of Spacecraft and Rockets 24*, 6 (1987), 518–522.

[42] DUNN, S. Hydrogen futures: Toward a sustainable energy system. Tech. Rep. Worldwatch Paper No. 157, WorldWatch Institute, 2001.

[43] ECONOMIST. Survey: The fuel cell's bumpy ride. *The Economist 358*, 8214 (2001).

[44] ECONOMIST. Survey: Clean machine. *The Economist 372*, 8391 (2004), 13.

[45] FAIRLEY, P. Hybrid's rising sun. *Technology Review* (2004).

[46] FIECHTER, C.-N. A parallel tabu search algorithm for large traveling salesman problems. *Discrete Applied Mathematics 51*, 3 (1994), 243–267.

[47] FROST. Satellite bus reliability analysis. Tech. rep., Frost and Sullivan and Airclaims, August 2004.

[48] GARCIA, D. Convergence and biases of monte carlo estimates of american option prices using a parametric exercise rule. *Journal Of Economic Dynamics and Control 27* (2003), 1855–1879.

[49] GAREY, M. R., AND JOHNSON, D. S. *Computers and Intractability.* Freeman, San Francisco, CA, 1979.

[50] GENDREAU, M., HERTZ, A., AND LAPORTE, G. A tabu search heuristic for the vehicle routing problem. *Management Science 40*, 10 (1994), 1276–1290.

[51] GENDREAU, M., AND POTVIN, J.-Y. A guide to tabu search. Tech. Rep. CRT-2003-23, Centre de Recherche sur les Transports, 2003.

[52] GESKE, R., AND JOHNSON, H. The american put valued analytically. *Journal of Finance 39*, 5 (1984), 1511–1524.

[53] GIBSON, R., AND SCHWARTZ, E. Stochastic convenience yield and the pricing of oil contingent claims. *Journal of Finance 45*, 3 (1990), 959–976.

[54] GLOVER, F. Future paths for integer programming and links to artificial intelligence. *Computers and Operations Research 13* (1986), 533–549.

[55] GLOVER, F., AND LAGUNA, M. *Tabu Search*. Kluwer Academic Publishers, Boston, 1997.

[56] GRANT, D., VORA, G., AND WEEKS, D. Path-dependent options: Extending the monte carlo simulation approach. *Management Science 43*, 11 (1997), 1589–1602.

[57] GURTUNA, O., AND TREPANIER, J. On-orbit satellite servicing: A space-based vehicle routing problem. In *Operations Research in Space and Air*, T. A. Ciriani, Ed. Kluwer, Dordrecht, 2003.

[58] HACKNEY, J., AND DE NEUFVILLE, R. Lifecycle model of alternative fuel vehicles: Emissions, energy, and cost tradeoffs. *Transportation Research Part A 35* (2001), 243–266.

[59] HAMEL, G. *Leading the Revolution: How to Thrive in Turbulent Times by Making Innovation a Way of Life*. Penguin Putnam Inc., New York, NY, 2000.

[60] HAUGH, M. B., AND KOGAN, L. Pricing american options: A duality approach. *Operations Research 52*, 2 (2004), 258–270.

[61] HELVIG, C., ROBINS, G., AND ZELIKOVSKY, A. Moving-target tsp and related problems. In *European Symposium on Algorithms* (1998).

[62] HOFFMANN, P. *Tomorrow's Energy: Hydrogen, Fuel Cells, and the Prospects for a Cleaner Planet*. The MIT Press, Cambridge, MA, 2001.

[63] IBANEZ, A., AND ZAPATERO, F. Monte carlo valuation of american options through computation of the optimal exercise frontier, 2001. Working Paper, Instituto Tecnologico Autonomo de Mexico, Mexico.

[64] ICHOUA, S., GENDREAU, M., AND POTVIN, J.-Y. Vehicle dispatching with time-dependent travel times. *European Journal Of Operational Research 144*, 2 (2003), 379–396.

[65] IEA. World energy investment outlook 2003. Tech. rep., IEA/OECD, 2003.

[66] IEA. World energy outlook 2004. Tech. rep., IEA/OECD, 2004.

[67] JU, N. Pricing an american option by approximating its early exercise boundary as a multi-piece exponential function. *Review of Financial Studies 11* (1998), 627–646.

[68] KAY, J., DAVISON, M., AND RASMUSSEN, H. The early exercise region for bermudan options on multiple underlyings. *Working Paper University of Western Ontario, London, Canada* (2003).

[69] KIM, J. The analytic valuation of american options. *Review of Financial Studies 3*, 4 (1990), 547–572.

[70] KNOX, J. Tabu search performance on the symmetric traveling salesman problem. *Computers and Operations Research 21* (1994), 867–876.

[71] KULATILAKA, N. The value of flexibility: The case of a dual-fuel industrial steam boiler. *Financial Management 22*, 3 (1993), 271–280.

[72] LACOMBE, X. Oos: Satellite die another day. In *1st Bilateral DLR-CSA Workshop on On-Orbit Servicing* (Cologne, Germany, 2002).

[73] LAGUNA, M. A guide to implementing tabu search. *Investigación Operativa 4*, 1 (1994).

[74] LINT, O. R and d as an option on market introduction. *R and D Management 28*, 4 (1998), 279–287.

[75] LIPMAN, T. E., AND DELUCCHI, M. A. Hybrid-electric vehicle design retail and lifecycle cost analysis. Tech. rep., Institute of Transportation Studies, April 2003.

[76] LONGSTAFF, F., AND SCHWARTZ, E. Valuing american options by simulation: A simple least-square approach. *Review of Financial Studies 14* (2001), 113–148.

[77] LUEHRMAN, T. A. Strategy as a portfolio of real options. *Harvard Business Review 76*, 5 (1998), 87–99.

[78] MACLEAN, H. L., AND LAVE, L. B. Evaluating automobile fuel/propulsion system technologies. *Progress in Energy and Combustion Science 29* (2003), 1–69.

[79] MALANDRAKI, C., AND DASKIN, M. Time dependent vehicle routing problems: Formulations, properties and heuristic algorithms. *Transportation Science 26* (1992), 185–200.

[80] MALEK, M., GURUSWAMY, M., M., P., AND OWENS, H. Serial and parallel simulated annealing and tabu search algorithms for the traveling salesman problem. *Annals of Operations Research 21* (1989), 59–84.

[81] MARGRABE, W. The value of an option to exchange one asset for another. *Journal of Finance 33* (1978), 177–87.

[82] MCKEAN, H. A free boundary problem for the heat equation arising from a problem in mathematical economics. *Industrial Management Review 6* (1965), 32–39.

[83] MERTON, R. Theory of rational option pricing. *Bell Journal of Economics and Management Science 4* (1973), 141–183.

[84] MITCHELL, G. R., AND HAMILTON, W. F. Managing r and d as a strategic option. *Research Technology Management 31*, 3 (1988), 15–22.

[85] MOORE, G. A. *Crossing the Chasm: Marketing and Selling Disruptive Products to Mainstream Customers*. HarperCollins, New York, NY, 2002.

[86] MORRIS, J. Darpa's orbital express demonstration set for 2006. *Aviation Week and Space Technology* (April 2002).

[87] MYERS, S. C. Determinants of corporate borrowing. *Journal of Financial Economics 5* (1977), 147–175.

[88] NAVC. Future wheels ii: A survey of expert opinion on the future of transportation fuel cells and fuel cell infrastructure. Tech. rep., Northeast Advanced Vehicle Consortium, 2003.

[89] OGDEN, J. M., WILLIAMS, R. H., AND LARSON, E. D. Societal lifecycle costs of cars with alternative fuels/engines. *Energy Policy 32* (2004), 7–27.

[90] OMAN, H. Energy sources for the world's post-petroleum era. *IEEE AES Systems Magazine 18*, 11 (2003), 35–39.

[91] OSMAN, I. Metastrategy simulated annealing and tabu search algorithms for the vehicle routing problem. *Annals of Operations Research 41* (1993), 421–451.

[92] OSMAN, I., AND LAPORTE, G. Metaheuristics: A bibliography. *Annals of Operations Research 63* (1996), 513–623.

[93] OUESLATI, S. K. *Evaluation of Nested and Parallel Options: Case Study of Ford's Investment in Fuel Cell Technology*. PhD thesis, MIT, 1999.

[94] PADDOCK, J. L., SIEGEL, D. R., AND SMITH, J. L. Option valuation of claims on real assets: The case of offshore petroleum leases. *The Quarterly Journal of Economics 103*, 3 (1988), 479–508.

[95] PADRO, C., AND PUTSCHE, V. Survey of the economics of hydrogen technologies. Tech. Rep. NREL/TP-570-27079, NREL, September 1999.

[96] PAXSON, D. A., Ed. *Real R and D Options*. Butterworth-Heinemann, Oxford, UK, 2003.

[97] REEVES, C. R. *Modern Heuristic Techniques for Combinatorial Problems*. Advanced Topics in Computer Science Series. Blackwell Scientific Publications, Oxford ; Boston, Mass., USA, 1993.

[98] REINELT, G. *The Traveling Salesman: Computational Solutions for Tsp Applications*. Springer Verlag, Berlin, 1994.

[99] RIFKIN, J. *The Hydrogen Economy: The Creation of the Worldwide Energy Web and the Redistribution of Power on Earth*. Penguin Putnam Inc., New York, NY, 2002.

[100] ROBEL, G. Real options and mean-reverting prices, 2001.

[101] ROBERTS, P. *The End of Oil: The Decline of the Petroleum Economy and the Rise of a New Energy Order*. Bloomsbury Publishing, London, UK, 2004.

[102] ROCHAT, Y., AND TAILLARD, E. D. Probabilistic diversification and intensification in local search for vehicle routing. *Journal of Heuristics 1* (1995), 147–167.

[103] SALEH, J. H., LAMASSOURE, E., AND HASTINGS, D. E. Space systems flexibility provided by on-orbit servicing: Part 1. *Journal of Spacecraft and Rockets 39*, 4 (2002), 551–560.

[104] SCHWARTZ, E., AND MOON, M. Evaluating research and development investments. In *Project Flexibility, Agency, and Competition*, M. J. Brennan and L. Trigeorgis, Eds. Oxford University Press, New York, 2000.

[105] SHISHKO, R., EBBELER, D. H., AND FOX, G. Nasa technology assessment using real options valuation. *Systems Engineering 7*, 1 (2003), 1–13.

[106] SMITH, J. E., AND MCCARDLE, K. F. Options in the real world: Lessons learned in evaluating oil and gas investments. *Operations Research 47*, 1 (1999), 1–15.

[107] SODAL, S. Fuel cells in shipping: Higher capital costs and reduced flexibility, 2003. Presented in the 7th Annual International Real Options Conference Washington DC July 10-12.

[108] TAILLARD, E. D. Parallel iterative search methods for vehicle routing problems. *Networks 23*, 8 (1993), 661–673.

[109] TEAL. Teal survey counts 600-610 active satellites currently in orbit. *Spacedaily* (2001).

[110] THOMAS, C., JAMES, B. D., LOMAX, F. D. J., AND KUHN, I. F. J. Fuel options for the fuel cell vehicle: Hydrogen, methanol or gasoline? *International Journal of Hydrogen Energy 25* (2000), 551–567.

[111] TOURINHO, O. A. *The Valuation of Reserves of Natural Resources: An Option Pricing Approach*. Ph.d. thesis, University of California at Berkeley, 1979. available at http://ideas.repec.org/p/ucb/calbrf/94.html.

[112] TRIGEORGIS, L. *Real Options: Managerial Flexibility and Strategy in Resource Allocation*. MIT Press, Cambridge, MA, 1996.

[113] TWISS, B. C. *Managing Technological Innovation*, 4th ed. Longman, London ; New York, 1992.

[114] WALTZ, D. M. *On-Orbit Servicing of Space Systems*. Krieger, Malabar, Fla., 1993.

[115] WEISS, M. A., HEYWOOD, J. B., SCHAFER, A., AND NATARAJAN, V. K. Comparative assessment of fuel cell cars. Tech. Rep. MIT LFEE 2003-001 RP, MIT, 2003.

[116] WILF, H. S. *Algorithms and Complexity*. University of Pennsylvania, Philadelphia, PA - Internet Edition, 1994.

[117] WILMOTT, P., HOWISON, S., AND DEWYNNE, J. *The Mathematics of Financial Derivatives: A Student Introduction*. Cambridge University Press, Cambridge, UK, 1995.

Appendix

A.1 Orbital Mechanics Equations

Period of an Elliptical Orbit

The period of an elliptical orbit is given by:

$T = 2\pi\sqrt{\frac{a^3}{\mu}}$

where μ is a physical constant ($398600.4\ km^3/s^2$) and a is the semi-major axis of the ellipse.

Plane Changes

The ΔV required for a simple plane change can be calculated by:

$\Delta V = 2V_i \sin\left(\frac{\alpha}{2}\right)$

where

$\Delta V =$ the velocity change required to produce a plane change

$V_i =$ the velocity of the servicing platform on the initial orbit at the intersection of the initial and final orbit planes

$\alpha =$ the angle of the plane change

Combined Hohmann Transfer

For a transfer from a circular orbit of radius r_1 to a coplanar, concentric orbit of radius r_2, the semi-major axis of the transfer ellipse is:

$a = \frac{1}{2}(r_1 + r_2)$

Initial velocity of a circular orbit is given by:

$$V_i = \sqrt{\frac{\mu}{r_1}}$$

Velocity at the perigee of the transfer ellipse is:

$$V_p = \sqrt{\frac{2\mu}{r_1} - \frac{\mu}{a}}$$

Therefore, ΔV for the first burn is:

$$\Delta V_1 = \sqrt{V_i^2 + V_p^2 - 2V_iV_p\cos(\alpha_{fraction} \times \alpha)}$$

where $\alpha_{fraction}$ is the optimal portion of the desired plane change α, that yields minimum ΔV_{total} by dividing the amount of plane change between the first burn and the second burn. $\alpha_{fraction}$ is computed using an iterative procedure.

For the second burn (plane change and circularization):

$$V_a = \sqrt{\frac{2\mu}{r_2} - \frac{\mu}{a}}$$

$$V_b = \sqrt{\frac{\mu}{r_2}}$$

$$\Delta V_2 = \sqrt{V_a^2 + V_b^2 - 2V_aV_b\cos[(1 - \alpha_{fraction}) \times \alpha]}$$

The total delta V required for the combined Hohmann transfer is:

$$\Delta V_{total} = \Delta V_1 + \Delta V_2$$

Phasing Maneuvers

Phasing maneuvers are used to equalize the angular phase of the servicing platform and the client satellite when the two are within the same orbital plane. The total ΔV needed for a phasing maneuver can be calculated using the following equation:

$$\Delta V_{total} = 2\left|\sqrt{\frac{2\mu}{a_0} - \frac{\mu}{a_t}} - \sqrt{\frac{\mu}{a_0}}\right|$$

where a_0 and a_t are the semimajor axes of the initial and transfer orbits, respectively. This equation can be transformed as follows:

$$\frac{\Delta V_{total}}{Vi} = 2\left|\sqrt{2 - \frac{a_0}{a_t}} - 1\right|$$

where V_i is the initial velocity.

Note that

$$\frac{a_0}{a_t} = \left(\frac{n_t}{n_0}\right)^{\frac{2}{3}}$$

where n_0 and n_t are the number of orbits per unit time for the initial and transfer orbits, respectively.

What needs to be determined now is the relative positions of the servicing platform and the client satellite. The relative positions determine the ratio $\left(\frac{n_t}{n_0}\right)^{\frac{2}{3}}$ and by substituting this value in the above equations, total ΔV for phasing maneuvers can be calculated.

Fuel Requirements

The fuel needed for an orbital maneuver can be calculated using the *rocket equation*:

$$m_f = m_0[\exp(\frac{1000.\Delta V}{I_{sp}.g}) - 1]$$

where

m_f =Mass (in kilograms) of fuel needed for the maneuver

m_0 =Mass of the servicing platform before the maneuver

I_{sp} =Specific impulse (in seconds) of the propulsion subsystem

g =Physical constant (9.80665 m/s^2) and ΔV is given in km/s.

A.2 Sequences for Table 3.6

Code	Best Solution (ΔV) - Exhaustive	Best Solution (time) - Exhaustive
0	n/a	n/a
1	0,1,0	0,1,0
2	0,1,2,0	0,2,1,0
3	0,1,3,2,0	0,2,3,1,0
4	0,3,2,4,1,0	0,2,3,4,1,0
5	0,5,1,3,2,4, 0	0,2,1,5,4,3,0
6	0,5,1,6,3,2, 4,0	0,6,1,5,4,2,3, 0
7	0,5,1,6,2,4, 7,3,0	0,6,1,5,4,7,2, 3,0
8	0,8,1,6,2,4, 7,3,5,0	0,2,1,5,4,8,3,7,6,0
9	0,5,3,2,4,7,9,1,6,8,0	0,2,1,5,4,6,9,8,3,7,0
10	0,5,2,4,7,10,3,9,1,6,8,0	0,6,9,1,5,4,8,7,2,3,10,0

Code	Best Solution (ΔV) - Tabu Search	Best Solution (time) - Tabu Search
0	n/a	n/a
1	0,1,0	0,1,0
2	0,1,2,0	0,2,1,0
3	0,1,3,2,0	0,2,3,1,0
4	0,3,2,4,1,0	0,2,3,4,1,0
5	0,5,1,3,2,4, 0	0,2,1,5,4,3,0
6	0,2,4,3,6,1,5,0	0,6,4,2,1,3,5,0
7	0,5,1,6,3,2, 4,7,0	0,4,6,1,7,2,3,5,0
8	0,8,1,6,2,4, 7,3,5,0	0,2,1,5,4,8, 3,7,6,0
9	0,5,3,2,4,7,9,1,6,8,0	0,2,1,5,4,6,9,8,3,7,0
10	0,4,7,10,3,2,9,1,6,8,5,0	0,10,3,4,2,8,1,9,5,7,6,0
11	0,3,2,4,7,10,6,1,9,5,8, 11,0	0,2,6,9,1,11,4,8,5,7,3, 10,0
12	0,3,2,4,7,10,5,8 ,9,12,1,6,11,0	0,2,6,1,9,12,11, 4,8,3,10,5,7,0
13	0,2,4,7,10,3,5,8,11, 12,1,6,9,13,0	0,12,1,5,4,8,11, 13,2,3,10,7,6,9,0
14	0,14,2,4,7,3,9,6,12,8, 5,10,13,1,11,0	0,10,14,9,8,1,5,2, 3,7,6,12,11,4,13,0
15	0,1,3,2,4,7,5,8,11,9, 12,6,10,13,15,14,0	0,7,12,5,11,9,8,13,14,4, 15,3,10,2,6,1,0
16	0,1,2,15,4,5,8,6,7,10,9, 12,11,14,16,13,3,0	0,10,3,8,11,13,14,15,6,1,9, 5,4,16,12,7,2,0
17	0,1,3,2,4,5,8,6,7,10,9, 12,11,14,15,16,13,17,0	0,1,5,14,4,8,11,13,12,7,2,3, 10,15,16,6,9,17,0
18	0,18,3,2,4,5,8,6,7,10,9, 12,11,14,15,16,13,1,17,0	0,2,6,12,9,1,11,16,18,7,17, 5,13,8,3,10,14,4,15,0
19	0,19,3,2,4,5,8,6,7,10,9,12, 11,14,15,16,13,17,1,18,0	0,17,1,11,7,18,13,2,14,9,12,3, 10,8,19,6,16,5,4,15,0
20	0,3,20,2,4,6,8,5,7,10,9,12, 11,14,15,16,13,17,1,19,18,0	0,17,10,15,3,20,2,12,11,19,18,13,6, 9,5,7,8,16,14,4,1,0

A.3 Screenshots for TSMC Software Implementations

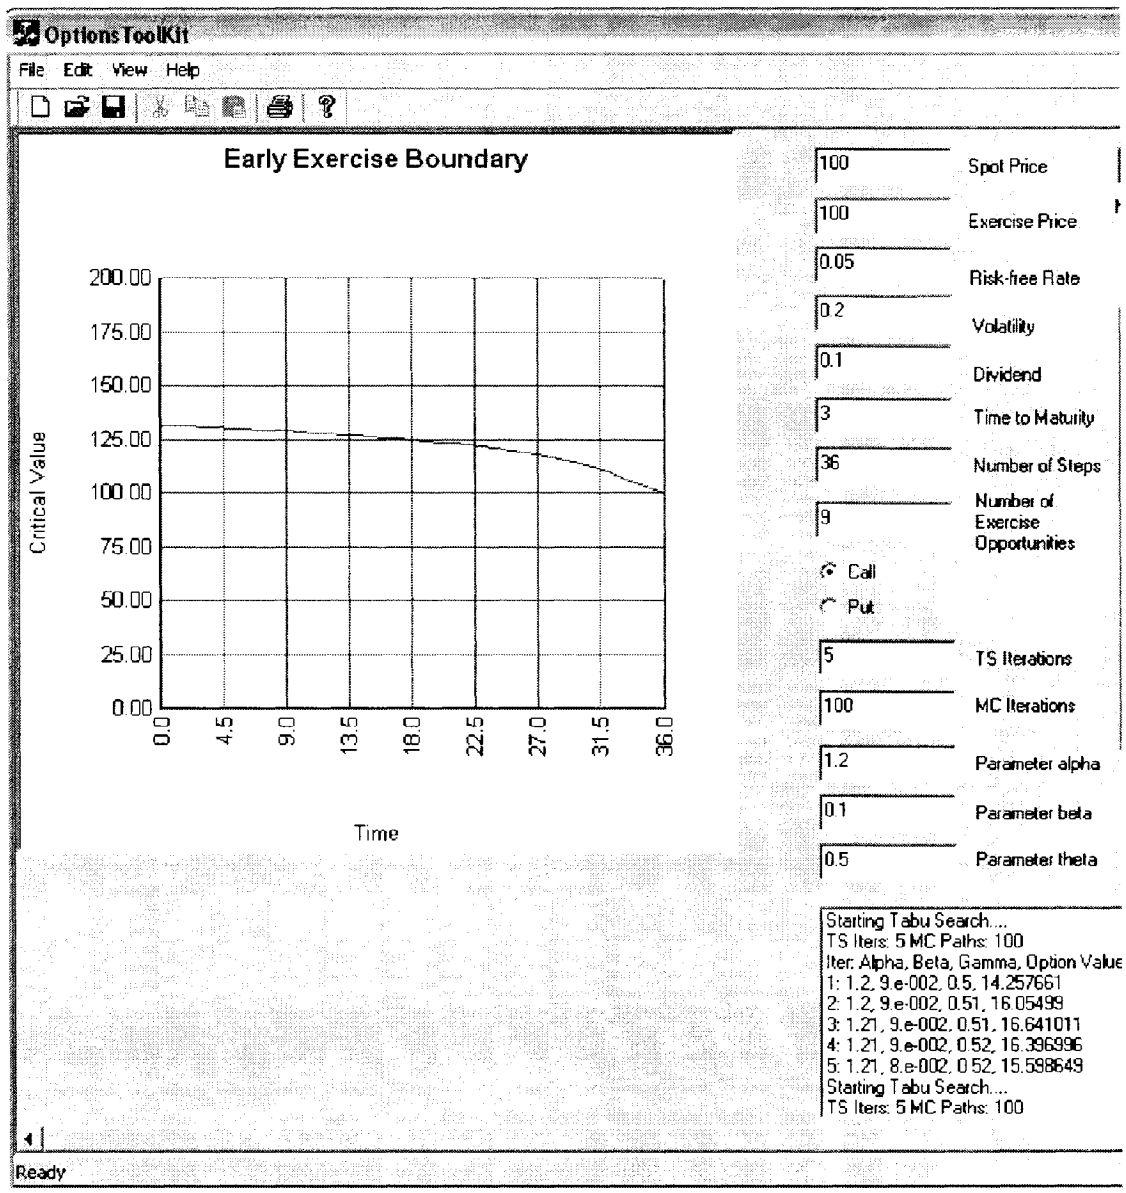

TSMC Software Implementation for Chapter 4

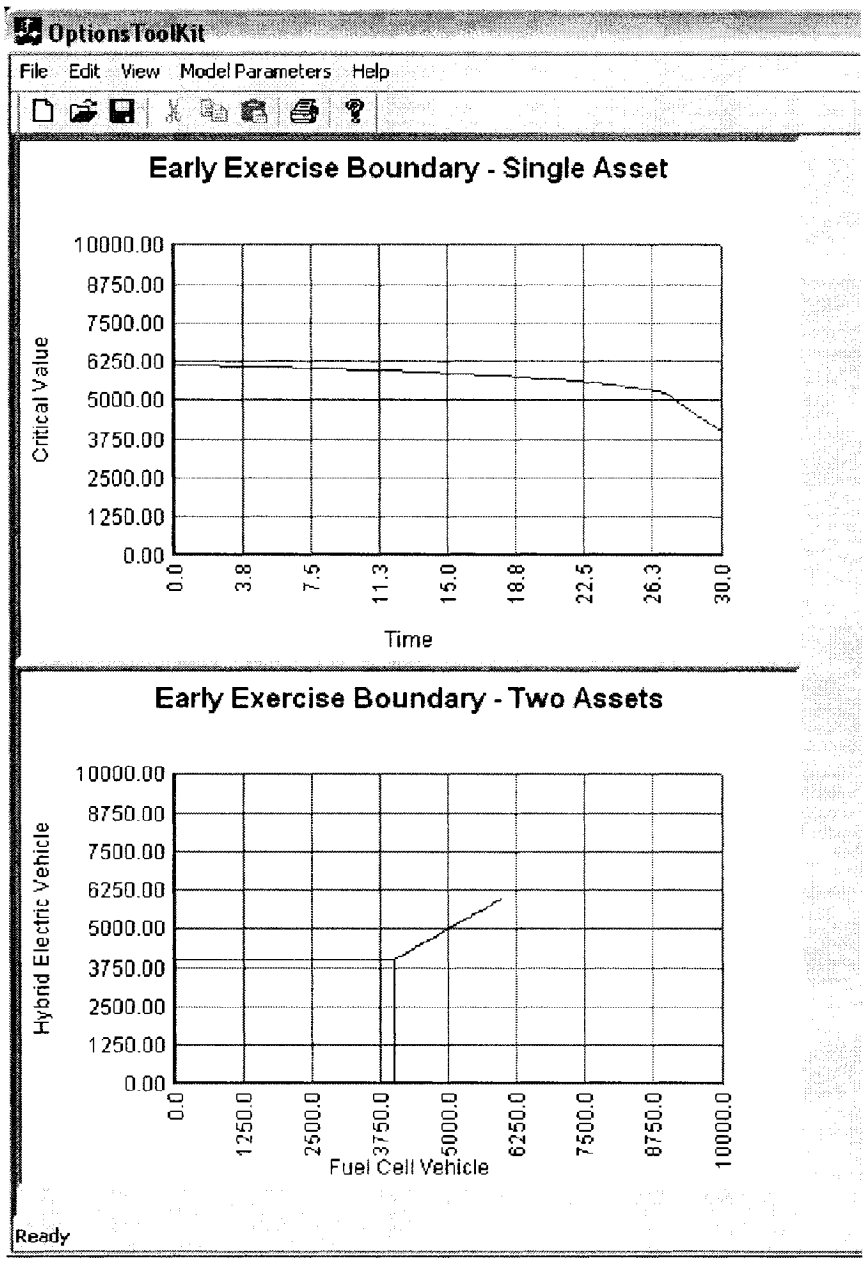

Figure 6.1: TSMC Software Implementation for Chapter 5