Theory and Applications of Compressed Annealing

by

Jeffrey W. Ohlmann

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Industrial & Operations Engineering)
in The University of Michigan
2003

Doctoral Committee:

Professor James C. Bean, Co-Chairperson
Professor Shane G. Henderson, Co-Chairperson
Assistant Professor Dushyant Sharma
Professor Frederick B. Talbot

ACKNOWLEDGEMENTS

From the one-room schoolhouse in the Sandhills of Nebraska to the campus of the University of Michigan, my 20+ years of formal education have been blessed with patient instructors, brilliant educators, and sage mentors.

I would first like to thank my co-advisors, Dr. James Bean and Dr. Shane Henderson, through which all things seem possible. Their advice and encouragement have been an invaluable resources. In addition, their generous financial support of my graduate studies allowed me to focus on research. Simply put, I can not imagine a better pair of mentors. I would also like to thank my committee members, Dr. Dushyant Sharma and Dr. F. Brian Talbot, for their contributions.

I would like to acknowledge the Alfred P. Sloan Foundation and the University of Michigan Trucking Industry Program (UMTIP) for supporting my work involving the trucking industry. I would also like to thank Katta Murty for administering the Graduate Aid in the Area of National Needs Fellowship. In addition, I owe a debt of gratitude to Meghan Cherry for the assistance she provided through the Marian Sarah Parker Scholars Program.

Over the course of my graduate studies at the University of Michigan, I have been fortunate to be surrounded by great people. Without these individuals, my degree would have not been possible. The IOE computer support staff of Chris Konrad, Rod Capps, and Mint Samsur provided essential technical support. The friendly staff of Fran Bourdas, Tina Blay, Gwen Brown, Celia Eidex, Pam Linderman, and

Mary Winter always shielded me from bureaucratic headaches and brightened the department with their smiles. I would also like to thank the IOE faculty for their excellent instruction. In particular, I thank Dr. Lawrence Seiford for his savvy advice and sharp wit.

Through it all, I have enjoyed the camaraderie of a collection of great friends. To Brock, Campbell, Gubbels, and Andy, thanks for keeping life in perspective. One of the best things about my five years at Michigan was the friends I made. I probably would have been wearing a straitjacket by now without the gang: Bailey, Fry, Darby, Scott G., Gries., Huggins, David K., Lisa Maillart, Matt Marshall, Maciek, and BT. Thank you for all the welcome distractions. A special thanks to Darby Grande for her comments on writing style, Chris Grieshaber for the construction of the outstanding figures, and Barrett Thomas for his counsel on the routing literature.

God has blessed me with a supportive and loving family. Thanks Mom, Dad, Deb, Brad, and Neal. Finally, I would like to express my appreciation for my loving wife, Amie. She's the most understanding and caring person I have ever met. She inspires me to be a better person.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

LIST OF APPENDICES

Appendix

CHAPTER I

Introduction

In the arena of operations and logistics, management is faced with problems that concern the efficient allocation of limited resources to meet desired objectives. The decisions in these problems are often modeled as discrete decision variables, due to the indivisibility of activities and resources (equipment, people, etc.). Just a few prominent examples include vehicle routing, capital budgeting, facility location and layout, health care management, telecommunications network design, and airline crew scheduling. The discrete nature of these problems leads to a combinatorial explosion in the number of possible solutions as problem size grows.

An interest in solving large-scale discrete optimization problems has led to the development of a variety of combinatorial optimization techniques. Entire texts have been devoted to various aspects of combinatorial optimization [26, 88, 104, 105, 2]. Hundreds of papers have been published on solution techniques for combinatorial problems ranging from classical examples (such as the assignment problem, knapsack problem, set-covering problem, and the traveling salesman problem) to specialized instances tailored for a particular study.

Early attempts at solving combinatorial problems focused their attention on optimization algorithms that employed sophisticated mathematical constructs. These "exact" methods often use the principles of linear programming, graph theory, dynamic programming, or branch-and-bound to utilize the particular structure of the problems for which they are designed. In many cases, these exact methods prove to be very effective. However, there exists a class of "difficult" problems for which optimal solutions for large-scale instances are unattainable within a reasonable amount of computation time using exact optimization approaches. Indeed, for some types of combinatorial problems, researchers have discovered that the computational effort required to obtain an optimal solution increases exponentially with the problem size.

Problems that have a known polynomial-time algorithm are categorized as class **P**. Unfortunately, many combinatorial problems do not have polynomial-time algorithms. Most of these "hard" problems cannot be proven to have exponential complexity, but rather are classified as **NP** (non-deterministic polynomial). For a formal discussion on computational complexity, we direct the reader to Garey and Johnson [44]. To overcome the computational complexity of **NP** problems, heuristic methods that effectively explore enormous solution spaces at reasonable computational costs have entered the mainstream, in both practice and theory. Heuristics may provide the only usable solutions to very difficult optimization problems for which the current exact algorithms are incapable of providing an optimal solution in reasonable times; when heuristics are used within an exact algorithm, they provide a bound to fix variables and to fathom branches in a search-tree.

The operations research literature is enriched with empirical and theoretical studies on heuristics, ranging from heuristic methods developed to solve particular instances to general purpose metaheuristic approaches. We omit discussion on the legion of specialized heuristics, encouraging the reader to explore the literature relevant to her/his area of interest. For a sampling of general purpose heuristic approaches, we refer the reader to numerous monographs containing discussions on simulated annealing [1, 85, 103, 6], genetic algorithms [65, 51, 95], tabu search [50], neural networks [64], and scatter search [86].

While the primary motivation for the application of heuristics is the computational intractability of large-scale problems, Reeves [114] supplies an additional argument. Heuristics often allow more flexibility in developing the model since they are capable of dealing with more complicated "realistic" objective functions and/or constraints while optimization routines may require simplifying assumptions (such as linearity). Therefore, the implementation of heuristics often allows more model accuracy at the price of obtaining approximate, rather than optimal, solutions to the model. Reeves [114] argue that this tradeoff favors heuristics since the optimization of a simpler, mathematically tractable model does not guarantee the optimization of the underlying real-world problem.

1.1 Research Objectives and Overview

The application of this study focuses on large-scale decision-making problems faced by the trucking industry. Faced by narrowing profit margins due to increased competition, trucking companies are continually looking for ways to reduce costs. In particular, we concentrate on the areas of vehicle replacement and routing.

The complexity and scope of the models that we consider motivate the application of heuristic solution procedures. Chapter II presents analysis of a metaheuristic called compressed annealing that we propose as a problem-solving tool. We describe the dynamics created by the integration of a variable penalty method with a hill-climbing search method. Through a Markov chain representation, we prove that, under appropriate assumptions, compressed annealing converges in probability to the set of global minima. We conclude Chapter II with general details for implementing compressed annealing, including procedures for initializing and calibrating parameters.

In Chapter III, we consider an asset replacement problem complicated by stochastic asset deterioration and budget constraints. A survey of current replacement strategies is conducted and maintenance logs for linehaul trucks are analyzed. We modify a mathematical formulation introduced by Morse [98] to fashion a replacement model that seamlessly assimilates data tracked by information technology to create a decision support system capturing the critical issues faced by trucking companies. Using a longitudinal sample of maintenance logs and information provided in interviews with fleet managers, we construct realistic data sets, the first of their kind in the literature. We employ compressed annealing on these instances of the stochastic fleet replacement problem with budget constraints (SFRPB), and compare the results to a trade cycle heuristic that mimics age-based replacement strategies commonly employed in the trucking industry.

Chapter IV contains a description of the traveling salesman problem with time windows (TSPTW) and its utility within the field of vehicle routing and beyond. We introduce a mathematical formulation of the problem and its associated notation. After describing a penalty method approach, we apply compressed annealing to the TSPTW and evaluate the results with respect to current benchmarks in the literature. We provide closing remarks in Chapter V summarizing the contribution of this thesis. In addition, promising avenues for future research are discussed.

The remainder of Chapter I surveys the three main bodies of literature relevant to the topics discussed throughout this thesis. In particular, we survey the voluminous body of research on simulated annealing, asset replacement, and the traveling salesman problem with time windows.

1.2 Literature Review

1.2.1 Simulated Annealing

Simulated annealing is a stochastic search method in which the ability to "climb hills" is governed by a control parameter called temperature. For high temperatures, simulated annealing is essentially a form of random search, which suffers from the curse of dimensionality. At the other extreme, as temperature approaches zero, simulated annealing becomes a descent method and cannot escape local minima. To capture the benefits of being able to escape local minima and still satisfactorily explore the basins containing them, temperature is initiated at a high value (corresponding to a high probability of accepting transitions to non-improving solutions) and reduced gradually over time. Such a "cooling schedule" complies with the analogy of physical annealing, in which a solid material is heated past its melting point and cooled slowly back to a solid state to assure a pure crystalline structure.

Simulated annealing can be also viewed as a modification of local neighborhood search. Rather than simply accepting only transitions to neighbors that offer an improvement in terms of solution value, simulated annealing's control parameter, temperature, triggers a mechanism to also allow "uphill" moves to avoid the convergence to local optima. At each iteration, simulated annealing selects a neighbor of the current solution. As in local search, if the neighbor solution has a cost less than the current solution's cost, it is accepted as the new "current" solution. If the neighbor solution has a cost greater than the current solution's cost, it is still accepted with a certain probability. This probability is dependent on the difference between the neighbor solution's cost and the current solution's cost as well as the temperature parameter. Acceptance probability decreases both as the temperature decreases and as the magnitude of the difference between the neighbor solution's cost and current solution's cost increases. Thus, as the algorithm proceeds and temperature is periodically reduced, the probability that a nonimproving solution is accepted also decreases.

Inspired by the Metropolis algorithm [94] of statistical mechanics, Cerny [19] and Kirkpatrick et al. [81] introduced simulated annealing as a discrete optimization tool in independent efforts. This work sparked a multitude of research on various aspects of the algorithm. The development of cooling schedules which dictate the rate at which the temperature parameter is reduced is a topic of both theoretical and empirical interest. Geman and Geman [46], Anily and Federgruen [5], Mitra et al. [96], and Johnson and Jacobsen [72] determine various sufficient conditions on the cooling schedule for convergence in probability to a global minimum. Chiang and Chow [23] and Holley [66] provide additional convergence results. Lundy and Mees [90] suggest a cooling schedule and stopping condition designed to produce a solution within ϵ of the global optimum with a given probability. Using results from continuous-time nonhomogeneous Markov chains, Gidas [49] and Hajek [56] determine necessary and sufficient conditions on the cooling schedule for the algorithm to converge in probability to a global minimum. Additionally, Rajasekaran [113] provides a bound on the time within which the algorithm converges with a high probability.

The collection of theoretical results shows that simulated annealing requires solution times that are exponential in problem size in order to guarantee optimal convergence. Empirical testing, however, has shown that simulated annealing can obtain near-optimal solutions within reasonable computation time for many discrete optimization problems. Theoretical research has guided empirical testing that has unearthed numerous findings regarding practical cooling schedules [15], neighborhood structures [118], and acceptance probabilities [73, 133].

In this thesis, we examine an approach that integrates variable penalty methods and simulated annealing. The relaxation of constraints through the use of penalty functions is a standard mathematical programming technique [8]. For heuristics with static penalty multipliers, either elaborate procedures must be performed to determine appropriate multiplier values, or they must be calibrated via extensive experimentation. As Costa and Oliveira [29] note, this can be a major drawback as it can be difficult to determine the correct weighting factors for the different penalty terms. To remedy this problem, Deb [31] develops a penalty term for a genetic algorithm approach that does not depend on a penalty parameter. Hadj-Alouane and Bean [55] present an alternate methodology by integrating a dynamic penalty method within a genetic algorithm approach. Coit et al. [28] and Coit and Smith [27] present an adaptive penalty method based on feedback from the heuristic search performed by a genetic algorithm. Gopalakrishnan et al. [52] use self-adjusting penalties to guide a tabu search heuristic applied to a class of production planning problems.

Similar constrained annealing approaches have been explored in the literature. Theodoracatos and Grimsley [131] develop a variable penalty method within an annealing algorithm to solve a two-dimensional packing problem. In their implementation, the penalty multiplier is defined as a function of the temperature parameter. Morse [98] implements an annealing method that treats penalty multipliers as a control parameter. This approach outperforms genetic algorithms and traditional simulated annealing for large-scale instances of a vehicle replacement problem.

Theoretical study of constrained annealing includes work by Geman [45], who provides convergence results for the application of the Gibbs sampler on an appropriately conditioned state space. Using Dobrushin's contraction technique, Yao [137] extends the work of Geman [45] to provide a sufficient condition on the convergence of constrained simulated annealing. Robini [116] improves this sufficient condition by providing a tight upper bound on the second largest eigenvalue in absolute value of the transition probability matrix associated with the underlying Metropolis chain. Frigerio and Grillo [42] and Del Moral and Miclo [92] also consider annealing with time-dependent energy functions, but under assumptions that imply an upper bound on the penalty multiplier.

1.2.2 Asset Replacement

We delineate the replacement literature along a pair of main identifying aspects: *scope* and *cost structure*. Scope refers to the perspective of the decision-maker, which in turn is dependent on the relationship between the assets. If the management of each asset is independent of the other assets, the scope is a single asset; this scenario is termed serial replacement. On the contrary, if there exists interdependencies between assets, then the scope consists of multiple assets. In general, there are two types of multiple asset interactions, parallel and series. If the assets are economically interdependent (due to economies of scale, resource constraints, etc.) and operate in parallel, this scenario is termed parallel replacement. Series replacement models describe assets that possess operational dependencies, i.e., machines that operate in series.

Cost structure refers to the certainty of the cash flows. If all cash flows are known with certainty at time zero, the replacement model is classified as deterministic. Alternatively, if cash flows resulting from purchase, operating, maintenance, and salvage are uncertain, the replacement model is classified as stochastic.

A vast majority of the replacement literature focuses on serial replacement and variants thereof. The classical treatment of this *defender* versus *challenger* decision assumes that the deterministic cash-flow series of any current challenger be repeated identically regardless of when that challenger might be acquired over an infinite horizon. Terborgh [129] and Alchian [3] present infinite-horizon models that relax the repeatability assumption by considering linear productivity improvement in future challengers under the restriction of constant replacement intervals for current and future challengers. Oakford et al. [101] generalizes a finite-horizon dynamic programming formulation in Wagner [135] that relaxes the repeatability assumption on future challengers.

Additional prominent work regarding deterministic serial replacement includes planning horizon approaches considering improving technology by Sethi and Chand [122], Chand and Sethi [20], and Bean et al. [10]. Hopp and Nair [70] and Nair and Hopp [100] present serial replacement models with stochastic technological breakthroughs. Cheevaprawatdomrong and Smith [21] investigate the effect of accelerating technological improvement on the replacement frequency. Under the assumptions of their model, they establish that replacements are made less frequently as technology improves. Bethuyune [13] analyzes serial replacement with variable utilization and discovers that the economic life of an asset can be prolonged by decreasing the utilization over its lifetime.

We are particularly interested in the serial replacement models that consider stochastic cash flows. Lohmann [89] analyzes uncertain cash flows by considering the deterministic dynamic program in Bean et al. [9] in conjunction with Monte Carlo simulation. For replacement models with uncertain cash flows due to stochastic deterioration, Markov decision processes are a common modeling approach. Derman [33] presents a model in which costs are dependent on the state of the asset, which deteriorates stochastically according to stationary transition probabilities. Assuming that costs are time-invariant, Derman [33] proves the optimality of a stationary *control limit* rule, stating that an asset is replaced if and only if it is in a state above a calculated threshold. *Repair limit* replacement rules, stating that a vehicle is replaced whenever its estimated repair costs exceed a threshold value, are developed in Drinkwater and Hastings [36], Hastings [63] and Mahon and Bailey [91]. For a replacement model with Markovian deterioration and possible technological breakthrough, Hopp and Nair [71] extend the results of Derman [33] by proving the existence of an optimal nonstationary control limit policy. In work related to the stochastic equipment replacement model, Hartman [59] explores a serial replacement model with probabilistic asset utilization in which an asset's *state* is defined by the asset's age and cumulative utilization.

The parallel replacement problem, termed by Vander Veen [93], considers a portfolio of economically interdependent assets operating in parallel. Research on parallel replacement problems has focused almost exclusively on deterministic cash flows. Jones et al. [76], Tang and Tang [128], and Hopp et al. [69] establish replacement rules for a parallel replacement problem in which economic interdependence is induced by a fixed charge for replacing one or more assets. Chen [22] develops efficient solution algorithms for the parallel replacement problem with economies of scale in replacement cost in both the finite- and infinite-horizon cases. Karabakal et al. [78] add capital rationing constraints to convert the serial replacement model in Oakford et al. [101] into a parallel replacement problem, and they present a branch-and-bound algorithm for solving moderately-sized instances optimally. In an extension of this work, Karabakal et al. [79] develop a dual heuristic for solving realistically sized instances of the same problem. Via integer programming, Hartman and Lohmann [60] model parallel replacement under demand and budget constraints with multiple replacement options. Hartman [58] generalizes the replacement rules of Jones et al. [76] for a parallel replacement problem with fixed and variable replacement costs in the presence of demand and budget constraints. Jones and Zydiak [74] analyze fleet design in the presence of time-invariant economic parameters, economies of scale in replacement costs, and dis-economies of scale in maintenance costs. They show that optimal steady-state fleet designs are composed of equal-sized replacement groups. Hartman [61] and Jones and Zydiak [75] continue discussion of the fleet design problem, focusing on the intricacies of cash-flow modeling.

Recent developments in parallel replacement have involved the unification of related decision-making processes. Hartman [57] develops an integer programming formulation to solve a replacement problem in which replacement and utilization decisions are jointly determined. Hartman [62] extends this work into stochastic demand environment, resulting in state-dependent stochastic costs, where the state is defined as a combination of age and utilization. Rajagopalan et al. [112] and Rajagopalan [111] present approaches unifying capacity expansion decisions and parallel asset replacement.

In this thesis, we analyze a fleet replacement problem considering the real-world conditions of stochastic vehicle deterioration, annual budget limits, and time-variant costs. In the context of the replacement literature, this model can be classified as parallel replacement with stochastic costs. To the author's knowledge, there is little research that directly deals with the prominent issues posed by this difficult combinatorial problem. Prior research involving parallel replacement has focused on deterministic cash flows, while work published on replacement with stochastic deterioration has generally concentrated on serial models that do not consider the effect of expenditure limits across a fleet.

1.2.3 Traveling Salesman Problem with Time Windows

The vehicle routing literature encompasses research on numerous applications approached with a variety of techniques. We study one of these prominent problems, the traveling salesman problem with time windows (TSPTW), a problem amenable to the constraint-handling machinery of the compressed annealing heuristic. Solution approaches for the TSPTW range from exact mathematical programming techniques to various heuristic approaches. With respect to the best-known results in the literature, we evaluate the performance of compressed annealing on instances of the TSPTW.

Exact approaches to the TSPTW have focused on dynamic programming techniques. Christofides et al. [25] and Baker [7] present branch-and-bound algorithms that solve problems with up to 50 vertices, but require "moderately tight" time windows and/or little overlap between them. Langevin et al. [87] introduce a two-commodity flow formulation well-suited to handling time windows; they solve instances with up to 40 nodes. Dumas et al. [38] extend earlier dynamic programming approaches by using state space reduction techniques that enable the solution of problems up to 200 customers. In an alternate approach, Pesant et al. [107, 108] utilize constraint programming to solve the TSPTW. Similarly, Focacci et al. [41] embed optimization techniques within a constraint programming approach.

Because of limitations with exact formulations (Savelsburgh [121] proved that even finding a feasible solution to the TSPTW is an NP-complete problem), particularly difficulties associated with wide time windows, there exists a facet of research focusing on heuristic techniques for the TSPTW. Carlton and Barnes [17], solve the TSPTW with a tabu search approach. Gendreau et al. [48] offer a construction and post-optimization heuristic based on a near-optimal TSP heuristic presented by Gendreau et al. [47]. Calvo [16] introduces a heuristic that constructs an initial tour using a unique relaxation to the assignment problem.

In recent years, heuristics have also been shown to be an effective means of solving a closely related problem, vehicle routing with time windows (VRPTW). The VRPTW is concerned with routing a fleet of vehicles when customers have time window constraints. As shown by Calvo [16] and Gendreau et al. [48], TSPTW heuristics are potentially more effective at optimizing the individual routes generated by the

VRPTW heuristics. This has led to many "cluster first, route second" approaches to the VRPTW in which a series of TSPTWs is solved after customers have been assigned to routes.

For the VRPTW, Solomon [123] adapts construction heuristics to handle time windows and conducts an extensive study across a wide range of data sets; a sequential insertion heuristic performed particularly well. To partition the customers in the VRPTW, Thangiah et al. [130] use a genetic algorithm. Garcia et al. [43] introduce a parallel implementation of a tabu search heuristic for the VRPTW. Russell [120] describes a hybrid VRPTW heuristic approach that integrates parallel construction and interchange improvement that effectively reduces vehicle fleet sizes. Chiang and Russell [24] embed simulated annealing for route improvement into this parallel construction algorithm for the VRPTW. Potvin et al. [110] and Taillard et al. [125] present tabu search heuristics, while Potvin and Bengio [109] approach the VRPTW with a genetic algorithm. For a comprehensive account of the implementation of simulated annealing, tabu search, and genetic algorithms on the VRPTW, we refer the reader to Tan et al. [127]. Utilizing constraint programming techniques, Caseau and Laburthe [18] and Rousseau et al. [119] also construct heuristics for VRPTW.

Although the VRPTW is NP-hard, there is still an active field of research in exact solution methods. These exact solution methods are applicable for smaller instances with tight time windows. Although the mathematical programming techniques of exact algorithms are vastly different than the search methods in heuristic approaches, we provide a brief genealogy. Kolen et al. [84] presented the first optimization method for the VRPTW. They apply a branch-and-bound procedure to solve VRPTW problems with up to 15 customers. Desrochers et al. [34] present an exact algorithm for the VRPTW based on a column generation approach of a set partitioning formulation. Kohl and Madsen [83] approach the VRPTW with a Lagrangian relaxation of the constraint set that all customers must be serviced to obtain optimal solutions. Fisher et al. [40] describe an algorithm based on the K-tree relaxation of the VRPTW. Most recently, Kohl et al. [82] apply a branch-and-bound algorithm that utilizes a method to obtain improved lower bounds to solve to optimality several previously unsolved problems.

CHAPTER II

Analysis of Compressed Annealing

A combinatorial optimization problem can be formulated as:

$$
\begin{array}{llll}
\textbf{CP} & \text{minimize} & f(x) & \\
 & \text{subject to:} & g_i(x) \geq b_i; & i = 1, \ldots, m \\
 & & x \in \mathcal{S}, &
\end{array}
$$

where $\mathcal{S}$ is a finite set, and $f(\cdot)$ and $g_i(\cdot)$ are real-valued functions on $\mathcal{S}$. We are particularly interested in instances of **CP** classified as **NP**-hard. The intractability of these problems suggests the application of metaheuristics to find near-optimal solutions. A trait common to many metaheuristic approaches is the requirement of "neighborhood" structures to generate candidate solutions from a current solution [115]. We consider instances of **CP** for which the formation of a neighborhood structure of feasible solutions is impeded by the constraints $\{g_i(x) \geq b_i\}$ for $i = 1, \ldots, m$. We recover well-defined neighborhoods by relaxing the complicating constraints into the objective function with a penalty term.

2.1 Penalty Methods and Annealing

For a solution $x \in \mathcal{S}$, let $\varrho(x)$ be the m-vector whose i^{th} element, $\varrho_i(x)$, is a real-valued, nonnegative function indicating violation of the constraint $g_i(x) \geq b_i$, so that $\varrho_i(x) > 0$ if and only if $b_i > g_i(x)$. In [54], Hadj-Alouane identifies one particular class of such penalty functions, which for $s > 0$ is given by:

$$\varrho(x) = \begin{bmatrix} |\min(0, g_1(x) - b_1)|^s \\ |\min(0, g_2(x) - b_2)|^s \\ \vdots \\ |\min(0, g_m(x) - b_m)|^s \end{bmatrix}.$$

Let θ be an m-vector of nonnegative, scalar penalty multipliers. Taking the dot product of θ and the penalty function $\varrho(x)$, we obtain the penalty term $\theta^t \varrho(x)$. Adding this penalty term to the objective function, we obtain **PP**, a general relaxation of **CP**.

PP minimize $\quad f(x) + \theta^t \varrho(x)$

subject to: $\quad x \in \mathcal{S}$

In the analysis of the compressed annealing algorithm, we make the restriction that all the components of θ are equal, i.e. $\theta_1 = \ldots = \theta_m$, so that all constraints relaxed into the objective function are multiplied by the same value. Under this restriction,

$$\theta^t \varrho(x) = \lambda \sum_{i=1}^{m} \varrho_i(x) = \lambda p(x),$$

where $\lambda = \theta_1 = \ldots = \theta_m$, and $p(x) = \sum_{i=1}^{m} \varrho_i(x)$.

Note that $p(x)$ is a real-valued, nonnegative function indicating violation of the constraints $\{g_i(x) \geq b_i\}$ for $i = 1, \ldots, m$, so that $p(x) > 0$ if and only if x is infeasible. We duly note that the choice of penalty function will have an impact on the empirical performance of the algorithm, and therefore favor penalty functions which maintain strong duality. For the nonnegative, scalar λ, we refer to the function $v(x, \lambda) = f(x) + \lambda p(x)$ as the *auxiliary function*. Then we formulate $\mathbf{RP}(\lambda)$, a relaxation of **CP** as

$$\begin{array}{lll} \mathbf{RP}(\lambda) & \text{minimize} & v(x,\lambda) = f(x) + \lambda p(x) \\ & \text{subject to:} & x \in \mathcal{S} \end{array}$$

We focus on solving **CP** via an implementation of simulated annealing on $\mathbf{RP}(\lambda)$. In addition to temperature, we have another parameter, namely the value of the penalty multiplier λ. Maintaining the physical analogy of simulated annealing, we call this parameter "pressure" [98]. The set $\mathcal{S}$ is finite, so for sufficiently large λ, any optimal solution to the relaxation $\mathbf{RP}(\lambda)$ is optimal for **CP** (we call this property strong duality); see [55]. Unfortunately, for large-scale problems, it is impractical to determine the exact multiplier value at which strong duality first holds. Fixing pressure at a "large" value to avoid converging to infeasible solutions might seem reasonable, but this makes it difficult for the annealing algorithm to move through the solution space. The high penalties mean that infeasible solutions are excessively penalized, and so practically speaking, the search is limited to feasible solutions. Fixing pressure at a "small" value ensures that the annealing algorithm can more easily move through the solution space, but one could converge to an infeasible solution.

In view of these observations, and because computational experience has demonstrated that it is often difficult to determine a "good" value for pressure, we examine a heuristic called compressed annealing [98]. Compressed annealing simultaneously adjusts pressure and temperature within the annealing run. Compressed annealing can be viewed as a hill-climbing algorithm that also alters the height of the "hills" in the solution topography. In this study, we will confine our attention to temperature schedules that are decreasing, and pressure schedules that are increasing.

Compressed annealing has exhibited success in obtaining reasonable solutions to difficult problems that genetic algorithms and traditional simulated annealing have not [98]. Additionally, Theodoracatos and Grimsley [131] also report empirical success using a similar variable penalty method and simulated annealing to solve a two-dimensional packing problem. In both of these applications, parameters are set after extensive experimentation. Our theoretical results provide insight on the trade-off between cooling and compression. We identify how the properties of a solution topology affect the specification of cooling and compression rates. In particular, we define the *dynamic depth* of a local minimum. The approximation of dynamic depth guides the practitioner's selection of annealing parameters and lessens the degree of "blind" experimentation.

2.2 Dynamics

In this section, we describe the dynamic behavior induced by the variable multiplier approach. For each state $x \in \mathcal{S}$, define $N(x) \subseteq \mathcal{S}$ as the static neighborhood of x, and let Γ be a stochastic matrix for generating neighbors.

Assumption 1. *For all $x, y \in \mathcal{S}$, such that $y \in N(x)$, there exists a constant c_0 such that $0 < c_0 \leq \Gamma(x, y) \leq 1$, and furthermore, $\Gamma(x, y) > 0$ if and only if $y \in N(x)$.*

Every state $x \in \mathcal{S}$ is described by the ordered pair $(f(x), p(x))$ consisting of the state's cost and degree of infeasibility. We say that a state x with $f(x) = f$ and $p(x) = p$ is at *level* (f, p). Levels are ordered such that $(f, p) \succ (f', p')$ if and only if either (i) $p > p'$ or (ii) $p = p'$ and $f > f'$. Furthermore, $(f, p) = (f', p')$ if and only if $p = p'$ and $f = f'$. In this manner, we lexicographically order the states in $\mathcal{S}$ with respect to $p(\cdot)$ and then $f(\cdot)$.

Define a solution *topology*, $\Theta = (\mathcal{S}, N, f, p)$, as the collection of states in a solution space connected by a neighborhood structure together with the functions f and p. A solution topology encapsulates all the essential problem parameters. We also define a solution *topography*, $\Theta_\lambda = (\mathcal{S}, N, v(\cdot, \lambda))$. The topography, Θ_λ, utilizes the information contained in the topology, Θ, to describe the relational behavior between states for a particular value of λ. While the topology is static, observe that the topography captures the dynamic impact of pressure since a change in λ alters the value of the auxiliary function, $v(x, \lambda)$, for every x such that $p(x) > 0$.

The ordering of states with respect to $v(\cdot, \lambda)$ may vary with λ. However, since $|\mathcal{S}| < \infty$, for λ sufficiently large, a constant ordering within the solution topography can be attained. Lemma 1 articulates this intuition and its corollary motivates the utilization of **RP**(λ) to solve **CP**.

Lemma 1. *For a solution topography Θ_λ with $|\mathcal{S}| < \infty$, the following are true.*

(a) There exists $\lambda^*_{x,y} \geq 0$ such that if $\lambda > \lambda^*_{x,y}$ and $(f(x), p(x)) \succ (f(y), p(y))$, then

$v(x, \lambda) > v(y, \lambda)$.

(b) There exists $\lambda^* \geq 0$ *such that* $\lambda > \lambda^*$ *if and only if* $v(x, \lambda) > v(y, \lambda)$ *for every pair of states* x, y *with* $(f(x), p(x)) \succ (f(y), p(y))$.

(c) For $\lambda > \lambda^*$, *all states that are local minima for the original problem* **CP** *are local minima for* $\mathbf{RP}(\lambda)$. *(Note: Local minima for* $\mathbf{RP}(\lambda)$ *are not necessarily feasible states.)*

Proof. *Part (a).* Consider $x, y \in \mathcal{S}$ such that $(f(x), p(x)) \succ (f(y), p(y))$. We want to show that for λ large enough,

$$v(x, \lambda) = f(x) + \lambda p(x) > f(y) + \lambda p(y) = v(y, \lambda). \tag{2.1}$$

Since $(f(x), p(x)) \succ (f(y), p(y))$, either (i) $p(x) > p(y)$ or (ii) $p(x) = p(y)$ and $f(x) > f(y)$. First, consider the case when $p(x) > p(y)$. For (2.1) to hold true,

$$\lambda > \max\left\{0, \frac{f(y) - f(x)}{p(x) - p(y)}\right\} = \hat{\lambda}_{x,y}.$$

Next, consider the case when $p(x) = p(y)$ and $f(x) > f(y)$. It is clearly true that (2.1) holds true for all values of $\lambda \geq 0$.

Part (b). ($\Rightarrow$) Let $\bar{\mathcal{S}}$ be the set of pairs of states $x, y \in \mathcal{S}$ such that $(f(x), p(x)) \succ (f(y), p(y))$. By Lemma 1(a), for each pair of states $x, y \in \bar{\mathcal{S}}$, there exists $\hat{\lambda}_{xy} \geq 0$ such that if $\lambda > \hat{\lambda}_{xy}$ then $v(x, \lambda) > v(y, \lambda)$. Therefore, if

$$\lambda > \max_{x, y \in \bar{\mathcal{S}}} \left\{\hat{\lambda}_{xy}\right\} = \lambda^*$$

observe that $v(x, \lambda) > v(y, \lambda)$ for every pair of states x, y with $(f(x), p(x)) \succ (f(y), p(y))$.

($\Leftarrow$) Suppose that $v(x, \lambda) = f(x) + \lambda p(x) > f(y) + \lambda p(y) = v(y, \lambda)$ for every pair of states $x, y \in \bar{\mathcal{S}}$. Solving for λ over every pair $x, y \in \bar{\mathcal{S}}$ and using the fact that λ is nonnegative, we obtain

$$\lambda > \max_{x,y \in \bar{\mathcal{S}}} \left[\max \left\{ 0, \frac{f(y) - f(x)}{p(x) - p(y)} \right\} \right] = \lambda^*.$$

Part (c). For ease of exposition, denote the state space of **CP** as $\mathcal{U}$, i.e.,

$$\mathcal{U} = \{ y : g_i(y) \geq b_i \text{ for } i = 1, \ldots, m; y \in \mathcal{S} \}.$$

Observe that $\mathcal{U} \subset \mathcal{S}$. Consider a state x that is a local minimum of **CP**. By definition of a local minimum of **CP**, $f(x) \leq f(y)$ for all $y \in N(x) \cap \mathcal{U}$.

Since x is a local minimum of **CP**, x must be a feasible solution of **CP**, $x \in \mathcal{U}$. Therefore, x is also feasible for **RP**(λ) and $p(x) = 0$.

To show that x is also a local minimum of **RP**(λ) for $\lambda > \lambda^*$, we must show that $v(x, \lambda) \leq v(z, \lambda)$ for all $z \in N(x)$. First, consider any state $y \in N(x) \cap \mathcal{U}$. Note that $p(y) = 0$ for all such y. Therefore,

$$\begin{aligned} v(x, \lambda) &= f(x) + \lambda p(x) \\ &= f(x) \\ &\leq f(y) \\ &= f(y) + \lambda p(y) \\ &= v(y, \lambda). \end{aligned}$$

Now consider any state $w \in N(x) \cap \mathcal{U}^c$. Since $w \in \mathcal{U}^c$, $g_i(w) < b_i$ for at least one i, and thus $p(w) > 0$. It follows then that $(f(w), p(w)) \succ (f(x), 0)$. From Lemma 1(a), there exists λ^* such that $\lambda > \lambda^*$ implies that $v(x, \lambda) \leq v(w, \lambda)$ for all such w.

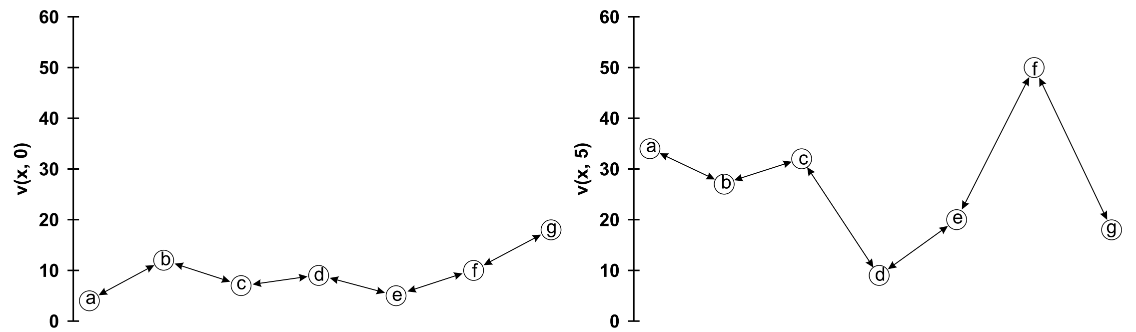

Figure 2.1: The two diagrams illustrate the dynamic nature of a solution topography for $\lambda = 0$ and $\lambda = 5$. States d and g are feasible ($p(d) = p(g) = 0$) and the other states are infeasible to varying degrees ($p(f) > p(a) > p(c) > p(b) > p(e)$). For this example, $\lambda^* = 4.\bar{3}$, so the topography on the left is "volatile" and the topography on the right is "stable."

Corollary 1. *If x is a global minimum of* **CP***, then x is also a global minimum of* **RP**(λ) *for $\lambda > \lambda^*$.*

Proof. The proof is similar to that of Lemma 1(c) and is omitted.

For $\lambda < \lambda^*$, there is a transient period in which the topography is "volatile." That is, the ordering of the states' auxiliary function values relative to each other is not fixed. For $\lambda > \lambda^*$, the topography stabilizes in the sense that the ordering of states with respect to $v(\cdot, \lambda)$ agrees with the relation $\succ$. Note that while a constant ordering of states in $\mathcal{S}$ is achieved for all $\lambda > \lambda^*$, further compression will increase the difference in the auxiliary function values between any particular pair of states x and y if $p(x) > p(y)$. Figure 2.1 illustrates the dynamic solution topography in a small example.

2.3 Convergence of Compressed Annealing

The augmentation of pressure has a significant impact on the annealing algorithm's convergence. We demonstrate that continued compression complicates the convergence of annealing by deepening the "valleys" and therefore making it increasingly difficult for the process to escape local minima.

Ideally, one would perhaps like to construct cooling and compression schedules that minimize the expected time required to find a global optimum. A first-order concern related to this goal is whether the time required to find a global minimum is almost-surely finite or not. In Theorem 3, we show that if the sum of the expected jump probabilities, denoted $\sum_{k=0}^{\infty} \eta(d_k)$, is finite, then there is a positive probability that a global minimum will never be reached (from certain states), i.e., the time to hit a global minimum, T_{global}, is infinite. We can conclude that if $P\{T_{global} = \infty\} = 0$, then $\sum_{k=0}^{\infty} \eta(d_k) = \infty$.

We also prove the converse, i.e., if $\sum_{k=0}^{\infty} \eta(d_k) = \infty$, then $P\{T_{global} = \infty\} = 0$, by supplying a stronger converse. Specifically, in Theorem 1 we show that if $\sum_{k=0}^{\infty} \eta(d_k) = \infty$, then $\lim_{k\to\infty} P\{Y(k) \in \mathcal{G}^*\} = 1$, where $\mathcal{G}^*$ is the set of global minima.

The main analytic result on the compressed annealing algorithm is a set of necessary and sufficient conditions for compressed annealing to converge in probability to the set of global minima. This result is a generalization of that in [56], in the sense that when the penalty multiplier is static or when there are no relaxed constraints, our results reduce to his. Our proof of this result follows a path hewn by Hajek [56], although many aspects of the proof involve nontrivial extensions of Hajek's concepts and results. We chose to adopt Hajek's framework as the basis for our approach because his necessary and sufficient conditions are the strongest in the literature.

Hajek showed that cooling rates that ensure convergence depend on the shape of the auxiliary function, and in particular, the "depth" of the deepest local, nonglobal minimum. This observation is of particular prominence in our work, where we alter the shape of the auxiliary function via compression during the annealing process. Continued compression deepens the "valleys" and therefore makes it increasingly more difficult for the process to escape local minima. Nevertheless, we prove that even if we allow pressure to increase without bound, there still exists a temperature schedule such that the annealing algorithm converges in probability to the set of global minima.

In §2.3.4.1, we show that such a temperature schedule must converge to 0 slower than $O([\ln k]^{-1})$. Therefore, practically speaking, one cannot expect to use compressed annealing to hit the set of global minima with probability 1. One might then consider the less-lofty goal of minimizing the expected time to hit a set of ϵ-optimal solutions, $E[T_\epsilon]$. But from Theorem 3, one can again conclude for some topologies and starting states, if $\sum_{k=0}^{\infty} \eta(d_k) < \infty$, then $P\{T_\epsilon = \infty\} > 0$. Despite these implications, we demonstrate compressed annealing's effectiveness as a heuristic approach on a pair of constrained combinatorial optimization problems (see Chapters III and IV).

The rest of this chapter is organized as follows. We describe a Markov chain representation of compressed annealing and its associated terminology in § 2.3.1, § 2.3.2, and § 2.3.3. In §2.3.4, we state our main result, its implications, and the outline of the proof. The remainder of the chapter is dedicated to the proof of the main result. In particular, §2.3.5 proves a result quantifying compressed annealing's ability to climb out of local, nonglobal minima. Then §2.3.6 lower-bounds the probability of the process being trapped near a local minima. In §2.3.7, we utilize the results of §2.3.5 and §2.3.6 to state a sufficient condition for the process to settle onto states "near" sufficiently deep local minima. Finally, §2.3.8 completes the proof of the main result.

2.3.1 Markov Chain Model and Vernacular

To begin our analysis, we introduce notation related to the control parameters, temperature and pressure. Define the cooling schedule to be a deterministic, decreasing sequence of strictly positive numbers $(\tau_0, \tau_1, \ldots)$ such that $\lim_{t\to\infty} \tau_t = 0$. Additionally, define the compression schedule as a deterministic, increasing sequence of nonnegative numbers $(\lambda_0, \lambda_1, \ldots)$ such that $\lim_{t\to\infty} \lambda_t = \infty$.

Using the compression schedule, we parameterize the auxiliary function of $\mathbf{RP}(\lambda)$, by defining $v_t(x) = f(x) + \lambda_t p(x)$. By Lemma 1 and our definition of a compression schedule, there exists $t^* < \infty$ such that $\lambda_t \leq \lambda^*$ for $0 \leq t \leq t^*$, and $\lambda_t > \lambda^*$ for all $t > t^*$. That is, t^* is the point in the compression schedule at which the topography stabilizes.

We model the compressed annealing algorithm as a time-inhomogeneous, discrete-time Markov chain $(Y(k), T(k) : k \geq 0)$ with state space $\mathcal{S} \times \{0, 1, \ldots\}$. At step k, the current solution of the algorithm is $Y(k) = x$, and $T(k)$ is interpreted as the "clock time." At each step of the algorithm, a candidate solution, y, is *generated* from the mass function, $\Gamma(x, \cdot)$, and y is then *accepted* with a probability dependent on $\tau_{T(k)}$ and the quantity $h_{T(k)} = (v_{T(k)}(y) - v_{T(k)}(x))^+$. If y is accepted, then $Y(k+1) = y$, otherwise $Y(k+1) = x$. In §2.3.2, we explain how $(T(k) : k \geq 0)$ evolves.

Let the acceptance function, $(\eta_t : t \geq 0)$, be a function of temperature, τ_t, such that $0 < \eta_t < 1$, η_t is nonincreasing in t, and $\lim_{t\to\infty} \eta_t = 0$. Typically, $\eta_t = \exp(-1/\tau_t)$. This function governs the probability of accepting transitions in our probabilistic hill-climbing algorithm. The term $\eta_{T(k)}^{h_{T(k)}}$ represents the probability of, at step k, accepting a transition from the current state, x, to a state $y \in N(x)$. Note that the height of the jump between states, $h_{T(k)} = \{f(y) - f(x) + \lambda_{T(k)}[p(y) - p(x)]\}^+$ is a function of the penalty multiplier, thus exhibiting the acceptance probability's relationship with the pressure parameter. For notational simplicity, we henceforth write $\eta_t^{h_t}$ as $\eta(h_t)$.

2.3.2 Structuring the Solution Space

Motivated by stabilization of the solution topography, Θ_λ, for $\lambda > \lambda^*$, we partition $\mathcal{S}$ into sets based on each state's *level*. Let $L(f,p) = \{x \in \mathcal{S} : f(x) = f, p(x) = p\}$ be the set of states at level (f,p). Note that for any $x \in L(f,p)$, $v_t(x) = f + \lambda_t p$, i.e., for any value of λ, all states in the same level have the same auxiliary function value. The order relation $\succ$ partitions $\mathcal{S}$ into $\{L(f_1,p_1), L(f_2,p_2), \ldots, L(f_\ell,p_\ell)\}$ where $(f_1,p_1) \prec (f_2,p_2) \prec \cdots \prec (f_\ell,p_\ell)$, and ℓ is the number of levels in $\mathcal{S}$.

As in [56], we assume that the problem is structured so that the process can only climb up one level at a time. This concept is formalized below.

Assumption 2. *The topology Θ possesses the continuous increase property. That is, if $x \in L(f_i, p_i)$ and $y \in L(f_j, p_j)$ with $j > i+1$, then $y \notin N(x)$.*

If a problem's topology possesses the continuous increase property, then for $x \in L(f_i, p_i)$, $N(x)$ consists only of states y such that $y \in L(f_j, p_j)$ where $j \leq i+1$. To interpret this, consider $Y(k) = x \in L(f_i, p_i)$ at iteration k where $\lambda_{T(k)} > \lambda^*$. Then the "distance" of any uphill transition is $f_{i+1} - f_i + \lambda_{T(k)}(p_{i+1} - p_i)$ (since we only climb one level at a time). Thus, for λ sufficiently large, the probability of accepting an uphill transition is the same for every state on a given level. Note that this property only affects uphill transitions and not downhill transitions that are always accepted.

At first glance, the continuous increase property appears to be quite restrictive. [56] shows that one can introduce artificial states to recover the continuous increase property for the case where all states are feasible. A similar augmentation technique can be applied to our situation.

Suppose that the topology, Θ, does not possess the continuous increase property. The idea is to augment $\mathcal{S}$ with artificial states to recover the continuous increase property and yet maintain the characteristics of the original solution space. We denote the augmented solution space $\hat{\mathcal{S}} = \mathcal{S} \cup \mathcal{A}$, where $\mathcal{A}$ is a set of artificial states complementing the set of real states, $\mathcal{S}$, such that $\mathcal{A} \cap \mathcal{S} = \emptyset$. The appropriate augmentation scheme involves inserting an appropriate artificial state at each level skipped between pairs of neighboring states.

If Θ does not possess the continuous increase property, then there exists at least

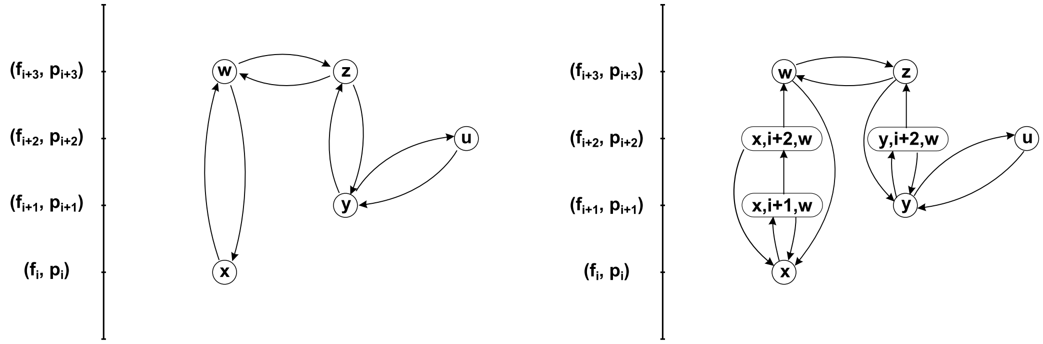

Figure 2.2: Solution topography (for $\lambda > \lambda^*$) before and after state augmentation.

one pair of states, $x \in (f_i, p_i)$ and $w \in (f_{i+l}, p_{i+l})$, such that $w \in N(x)$ and $l > 1$ (see Figure 2.2). For every such pair, we insert artificial states $(x, i+1, w) \in L(f_{i+1}, p_{i+1})$, $(x, i+2, w) \in L(f_{i+2}, p_{i+2})$, $\ldots$, and $(x, i+l-1, w) \in L(f_{i+l-1}, p_{i+l-1})$ into the solution space. The neighborhood structure is altered so that $(x, i+1, w) \in \hat{N}(x)$ is generated with probability $\hat{\Gamma}(x, (x, i+1, w)) = \Gamma(x, w)$ and $\hat{\Gamma}(x, w) = 0$. Additionally, $(x, i+2, w) \in \hat{N}(x, i+1, w)$ with $\hat{\Gamma}((x, i+1, w), (x, i+2, w)) = 1$, $\ldots$, and $w \in \hat{N}(x, i+l-1, w)$ with $\hat{\Gamma}((x, i+l-1, w), w) = 1$.

To keep the augmentation of artificial states from altering the behavior of $(Y(k), T(k))$, we design a Markov chain such that the artificial states have holding times of zero time units while the real states have holding times of one time unit. This is the motivation for the clock time component of our Markov chain. For example, consider the sample path: $Y(k) = x_0$, $Y(k+1) = x_1$, $Y(k+2) = x_2$, $Y(k+3) = x_3$. If $x_0, x_1, x_2, x_3 \in \mathcal{S}$ and $T(k) = t$, then the corresponding sequence of clock times is: $T(k+1) = t+1$, $T(k+2) = t+2$, $T(k+3) = t+3$. However, if $x_0, x_1, x_3 \in \mathcal{S}$ and $x_2 \in \mathcal{A}$ instead, then $T(k+1) = t+1$, $T(k+2) = t+1$,

$T(k+3) = t+2$.

To be precise, the one-step transition probabilities can be defined as follows. For $x \in \mathcal{S}$ and $y \in \hat{\mathcal{S}} \backslash \{x\}$,

$$P\left\{Y(k+1) = y, T(k+1) = s \,|\, Y(k) = x, T(k) = t\right\} =$$
$$\hat{\Gamma}(x,y)\eta([v_t(y) - v_t(x)]^+) I\left\{s - t = I\left\{y \in \mathcal{S}\right\}\right\},$$

and

$$P\left\{Y(k+1) = x, T(k+1) = t+1 \,|\, Y(k) = x, T(k) = t\right\} =$$
$$1 - \sum_{z \neq x} \hat{\Gamma}(x,z)\eta((v_t(z) - v_t(x))^+).$$

Now suppose the algorithm is in state $(x,i,w) \in \mathcal{A}$ at step k. The algorithm generates a neighbor solution, y, with probability $\hat{\Gamma}((x,i,w),y)$ and accepts this state with probability $\eta([v_{T(k)}(y) - v_{T(k)}(x,i,w)]^+)$. If the move to y is rejected, then the process descends to x. More precisely, for $(x,i,w) \in \mathcal{A}$ and $y \in \hat{\mathcal{S}} \backslash \{x\}$,

$$P\left\{Y(k+1) = y, T(k+1) = s \,|\, Y(k) = (x,i,w), T(k) = t\right\} =$$
$$\hat{\Gamma}((x,i,w),y)\eta([v_t(y) - v_t(x,i,w)]^+) I\left\{s - t = I\left\{y \in \mathcal{S}\right\}\right\},$$

and

$$P\left\{Y(k+1) = x, T(k+1) = t+1 \,|\, Y(k) = (x,i,w), T(k) = t\right\} =$$
$$1 - \sum_{z} \hat{\Gamma}((x,i,w),z)\eta([v_t(z) - v_t(x,i,w)]^+).$$

In the augmented solution space, $\hat{\mathcal{S}}$, there is a positive probability of the process transitioning from an artificial state (x,i,w) to the state $x \neq (x,i,w)$ at any time

$t \geq 0$. That is, x is a "neighbor" of (x, i, w) although $\hat{\Gamma}((x, i, w), x) = 0$. To remedy this contradiction, we define neighborhoods for artificial states differently than neighborhoods for real states.

Assumption 3. *For any state $z \in \mathcal{A}$, a state $x \in \hat{N}(z) \subset \hat{\mathcal{S}}$ if and only if either*
(a) $\hat{\Gamma}(z, x) = 1$ (x is the state generated for consideration), or
(b) $P\{Y(k+1) = x, T(k+1) = t+1 \mid Y(k) = z, T(k) = t\} > 0$ (x is the last real state visited).

The above procedure allows us to recover the continuous-increase property while maintaining the transition probabilities between real states. So we henceforth assume, without loss of generality, that the continuous-increase property holds.

2.3.3 Definition of Concepts

In this section, we use insight provided by the constant ordering of the topography, Θ_λ, for $\lambda > \lambda^*$ to define structural concepts. In particular, we extend the Markov chain concepts of accessibility and communication to account for the dynamic solution topography. These definitions extend the framework in [56] to allow for infeasible states.

We say that state y is *accessible at level* (f, p) from state x if (i) $x = y$ and $(f(x), p(x)) \preceq (f, p)$, or (ii) there is a sequence of states $x = x_0, x_1, \ldots, x_s = y$ for some $s \geq 1$ such that $x_{n+1} \in N(x_n)$ for $0 \leq n < s$ and $(f(x_n), p(x_n)) \preceq (f, p)$ for $0 \leq n \leq s$.

A topology $(\mathcal{S}, f, p, N)$ is *level irreducible* if $(\mathcal{S}, N)$ is irreducible and, for any level (f, p) and any two states $x, y \in S$, x is accessible at level (f, p) from y if and only if y is accessible at level (f, p) from x. Note that symmetric neighborhoods are sufficient, but not necessary, for level irreducibility.

We say that a state x is a *local minimum* if there does not exist a state y accessible at level $(f(x), p(x))$ from x such that $(f(y), p(y)) \prec (f(x), p(x))$. For a local minimum x, suppose $(f'_x, p'_x) \succ (f(x), p(x))$ is the lowest level at which a state y with $(f(y), p(y)) \prec (f(x), p(x))$ is accessible. Then, the *static depth* of local minimum x is $S(x) = f'_x - f(x)$, and the *penalty depth* of local minimum x is $P(x) = p'_x - p(x)$. Therefore, each local minimum x is described by a *depth pair* $(S(x), P(x))$.

We compare the depth pairs of local minima x and y via the $\succ$ operator. If $(S(x), P(x)) \succ (S(y), P(y))$, then we say that the local minimum x is "deeper" than the local minimum y. We define the *dynamic depth* of local minimum x at time t as $d_t(x) = S(x) + \lambda_t P(x)$. For $t \geq t^*$, $d_t(x)$ is interpreted as the smallest distance that the process would have to climb from a local minimum x to reach a local minimum z such that $(f(z), p(z)) \prec (f(x), p(x))$.

For a local minimum x, recall $(f'_x, p'_x) \succ (f(x), p(x))$, and thus $(S(x), P(x)) \succ (0, 0)$. For a state y that is not a local minimum, there exists a state z with $(f(z), p(z)) \prec (f(y), p(y))$ that is accessible at level $(f(y), p(y))$ from y. Thus, any state y that is not a local minimum has a depth pair $(S(y), P(y)) = (0, 0)$. On the other hand, if a local minimum x is also a global minimum, then $(S(x), P(x))$ is defined to be $(+\infty, +\infty)$.

A *cup* C is a set of states such that for some level (f, p) and some x such that $(f(x), p(x)) \preceq (f, p)$, $C = \{y : y \text{ is accessible at level } (f, p) \text{ from } x\}$. By level irreducibility, one can take the state "x" in the definition of the cup C to be any element of C.

To facilitate our analysis of cups, we define some cup-specific quantities. Let the *bottom* of a cup C be defined as $B = \{x \in C : (f(x), p(x)) \preceq (f(y), p(y)) \,\forall\, y \in C\}$. Let the *rim* of cup C be defined as $R = \{x \in C : (f(x), p(x)) \succeq (f(y), p(y)) \,\forall\, y \in C\}$. That is, B is the set of states in the lowest level of the cup, and R is the set of states in the highest level of the cup. Let the *froth* of a cup C be $F = \{y : y \notin C \text{ and } y \in N(x) \text{ for some } x \in C\}$. Note that F is empty if and only if $C = \mathcal{S}$. We are particularly interested in the states in R that have neighbors in F. Formally, let this set of "escape states" be denoted $R^e = \{x \in R : N(x) \cap F \neq \emptyset\}$. By the continuous increase property, the process escapes C by jumping from a state in R^e to a state in F. The sets B,R, and F are cup-specific, i.e., they are more accurately denoted $B(C)$, $R(C)$, and $F(C)$. We notationally suppress this dependence, unless discussing more than one cup at a time.

For a cup $C \neq \mathcal{S}$, we define the *static depth* $S = f(y) - f(x)$ and the *penalty depth* $P = p(y) - p(x)$, where $y \in F$ and $x \in B$. Note that for any cup $C \neq \mathcal{S}$, $(S, P) \succ (0, 0)$ implying that either (i) $P > 0$ and S is unrestricted, or (ii) $P = 0$ and $S > 0$. We say that cup C with depth pair (S, P) is deeper than cup C' with depth pair (S', P') if $(S, P) \succ (S', P')$. Figure 2.3 illustrates a cup and related concepts.

We define the *dynamic depth* of cup C at time t as $d_t = S + \lambda_t P$. For $t \geq t^*$, $d_t(C)$ is the distance that the process would have to climb to reach some state $y \in F$ from some state $x \in B$. In accordance with the above definitions, the deepest cup C such

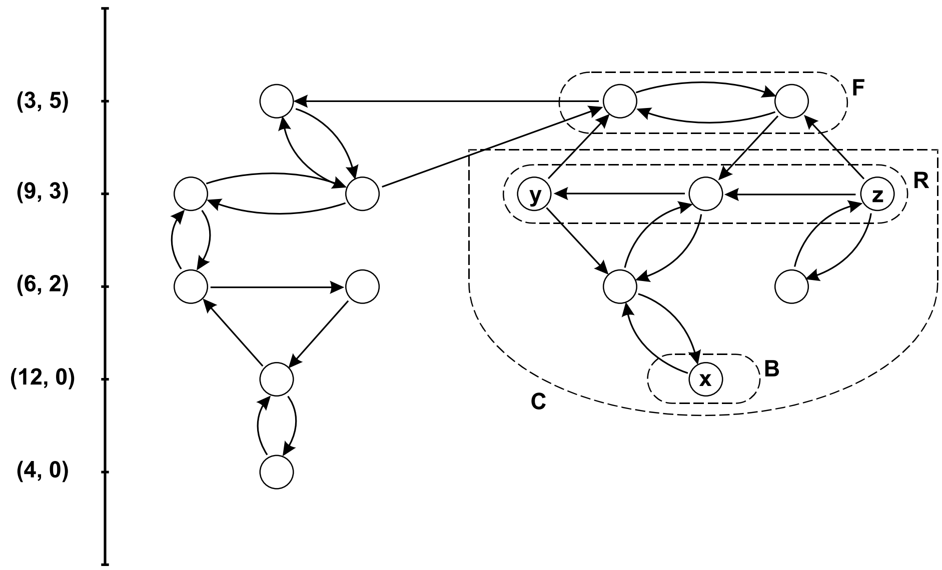

Figure 2.3: A cup C with depth pair $(S, P) = (S(x), P(x)) = (-9, 5) = [(3, 5) - (12, 0)]$. Note that $R^e = \{y, z\}$.

that a local minimum x is in the bottom of C has depth pair $(S, P) = (S(x), P(x))$.

Let the "rim height" $r_t(C)$ be the distance at time t from a state $u \in R$ to $z \in F$, i.e., $r_t = f(z) - f(u) + \lambda_t(p(z) - p(u))$. Similarly, let the "girth," $g_t(C)$, be the distance at time t from a state $x \in B$ to $u \in R$, i.e., $g_t = f(u) - f(x) + \lambda_t(p(u) - p(x))$. In this way, the dynamic depth of cup C at time t can be decomposed into two pieces, $d_t = r_t + g_t = S + \lambda_t P$.

Note that cups are defined independently of λ_t, although the depths of cups depend on λ_t and only are intuitive for times $t \geq t^*$. The deepening of cup C from time t to time $t + 1$, given by $d_{t+1} - d_t = P(\lambda_{t+1} - \lambda_t)$, is governed by its penalty depth and the compression schedule.

2.3.4 Necessary and Sufficient Condition

With the model structure and language in place, we are ready to state our main result.

Theorem 1. *If the topology* Θ *is level irreducible, then the following statements hold.*
(a) For any state x *that is not a local minimum,* $\lim_{k\to\infty} P\{Y(k) = x\} = 0$.
(b) Let B *be the bottom of a cup* C *of static depth* S *and penalty depth* P *so that the states in* B *are local minima of static depth* S *and penalty depth* P*. Then,*

$$\lim_{k\to\infty} P\{Y(k) \in B\} = 0 \tag{2.2}$$

if and only if

$$\sum_{k=0}^{\infty} \eta(d_k^+) = \sum_{k=0}^{\infty} \eta([S + \lambda_k P]^+) = +\infty. \tag{2.3}$$

(c) Let (S^*, P^*) *be the maximum depth of all states which are local, nonglobal minima. Let* $\mathcal{G}^*$ *denote the set of global minima. Then*

$$\lim_{k\to\infty} P\{Y(k) \in \mathcal{G}^*\} = 1 \tag{2.4}$$

if and only if

$$\sum_{k=0}^{\infty} \eta([S^* + \lambda_k P^*]^+) = +\infty \tag{2.5}$$

Part (a) of Theorem 1 states that the compressed annealing chain is unlikely to reside in a state that is on a "valley wall" in the solution topography. Part (b) declares that the compressed annealing algorithm is unlikely to reside in a local minimum with depth pair less than or equal to (S, P) if and only if the system is cooled and compressed slowly enough. We use the "positive part" operator in the sum of escape probabilities to avoid situations where $S + \lambda_k P < 0$, which can only occur for small k. This operator does not play a role as k tends to infinity. Part (c) is a direct consequence of parts (a) and (b).

2.3.4.1 Cooling and Compression Schedules

Suppose the probability of accepting a nonimproving solution is given by the function $\eta_k = \exp(-1/\tau_k)$. Then the necessary and sufficient condition for convergence is

$$\sum_{k=0}^{\infty} \exp(-[S^* + \lambda_k P^*]^+/\tau_k) = \infty. \tag{2.6}$$

[56] demonstrates that a cooling schedule of the form $\tau_k = c/\ln(k+1)$ for $k \geq 1$ satisfies this condition if c is a constant greater than or equal to the depth of the deepest local, nonglobal minimum. Thus, if we consider a compression schedule such that $\lim_{k\to\infty} \lambda_k = \bar{\lambda} \geq \lambda^*$, where λ^* is defined in Lemma 1, Hajek's cooling schedule still satisfies Equation (2.6) if $c \geq S^* + \bar{\lambda} P^*$.

However, if $P^* > 0$ and we allow $\lim_{k\to\infty} \lambda_k = \infty$, the depths of local minima also grow to infinity and therefore Hajek's cooling schedule does not satisfy the necessary and sufficient condition. In this case, one particular set of joint cooling and compression schedules that satisfy the necessary and sufficient condition is

$$\tau_k = \frac{(S^* + \phi) + \lambda_k P^*}{\ln(k+2)}$$

for $k \geq 0$, where λ_k grows to infinity slower than $O(\ln k)$ and $\phi \geq 0$ is appropriately defined to ensure that $\{\tau_k\}$ is decreasing.

While cooling and compression schedules that satisfy the necessary and sufficient condition are often empirically ineffective on practically sized problems, our theoretical structure still contributes to the practice of compressed annealing. Equation (2.6) suggests compression and cooling schedules with rates of change that decrease over time. In addition, it is clear that the performance of compressed annealing is affected by the dynamic depths of local minima, and in particular, the depth pair of the deepest local, nonglobal minimum, (S^*, P^*). This suggests that approximations of depth pairs and the pressure cap (λ^*) supply insight on the impact of compression on the solution topography and further guide the practitioner's specification of temperature and pressure. These insights have proven valuable in empirical work on constrained truck fleet replacement and vehicle routing with time windows. See also the motivation for our results given in the introduction.

2.3.4.2 Outline of Proof

The proof of Theorem 1, given in Section 2.3.8, follows from three theorems we present in Sections 2.3.5, 2.3.6, and 2.3.7. In this section, we describe the key ideas behind these three major steps.

In Section 2.3.5, we establish Theorem 2 which states that under appropriately slow cooling and compression: (a) the sum of jump probabilities over the number of steps required to escape a cup C is bounded in expectation, and (b) there is a positive uniform lower bound on the probability of exiting the cup via any particular state in the cup's froth. To prove Theorem 2, we begin by partitioning a cup C into its rim and its underlying cups. Structuring a cup in this manner accommodates the use of induction. The proof of part (a) rests on the fact that by the induction hypothesis, Y escapes any underlying cup within cup C by entering a state in the rim of C, and on these trips to $R(C)$, Y occasionally hits a state in $R^e(C)$. Once Y is in $R^e(C)$, the process can escape cup C in a single transition, so if Y visits $R^e(C)$

often enough, the process will eventually escape C, and escape quickly enough to ensure the expectation is bounded.

To show part (b) of Theorem 2, we consider some state $\bar{x} \in R^e(C)$ with $\bar{y} \in N(\bar{x}) \cap F$ and recognize that if $P\{Y \text{ escapes } C \text{ through } \bar{x}\} > \delta > 0$, then $P\{Y \text{ enters } \bar{y} \text{ upon exiting } C\}$ is also bounded away from zero. We prove that the $P\{Y \text{ escapes } C \text{ via } \bar{x}\}$ has a positive lower bound by showing that, for some constant c,

$$P\left\{Y \text{ escapes } C \text{ via } \bar{x} \text{ after hitting } \bar{x}\ i \text{ times }\right\} \geq$$
$$cP\left\{Y \text{ escapes } C \text{ from another state in } R^e \text{ after hitting } \bar{x}\ i \text{ times }\right\},$$

and then summing over i. Since every state $y \in F$ is a neighbor of some $x \in R^e$, part (b) follows directly.

Section 2.3.6 develops Theorem 3 which presents a lower bound on the probability of remaining in a cup over a given number of iterations. To achieve this result, we partition a cup C into a set D and a collection of cups $\{C_1, \ldots, C_n\}$. The set D is defined such that for $x \in D$, the bottom of C is accessible at level $(f(x), p(x))$. Alternatively, for $x \in C_i$, $(f, p) \succ (f(x), p(x))$ is the lowest level such that the bottom of C is accessible at (f, p) from x. Viewing the structure of C in this way, we see that Y can escape the cup or hit the bottom only through states in D. We proceed by utilizing a time-homogeneous birth-and-death process (that stochastically dominates (Y, T)) to determine an upper bound on the probability that Y escapes C before hitting the bottom. This upper bound on $P\{Y \text{ escapes } C \text{ before hitting } B\}$ is then used to develop the lower bound on the probability of remaining in the cup over n iterations as stated in Theorem 3.

In Section 2.3.7, Theorem 4 provides the capstone to the proof of Theorem 1. Theorem 4 establishes $\sum_{k=0}^{\infty} \eta(S + \lambda_k P) = +\infty$ as the sufficient condition for convergence in probability to the set of states composing the lower portions of cups deeper than (S, P). Define $E_{S,P}$ as the set of local minima deeper than (S, P). To establish this sufficient condition, we partition the solution space $\mathcal{S}$ into a set U and a collection of cups $\{C_1, \ldots, C_l\}$. The set U is defined such that for any $x \in U$, there exists a state in $E_{S,P}$ that is accessible at level $(f(x), p(x))$. Alternatively, for $x \in C_j$, $(f, p) \succ (f(x), p(x))$ is the lowest level such that a state in $E_{S,P}$ is accessible at (f, p) from x. Furthermore, the depth of any cup in the collection $\{C_1, \ldots, C_l\}$ is at most (S, P).

Due to the partitioning of $\mathcal{S}$, the process hits U upon exiting any cup C_j, and from any state $x \in U$, there exists a path to $E_{S,P}$ that rises no higher than $(f(x), p(x))$. Thus, we use Theorem 2 to show that Y escapes any cup C_j, enters a state in U, and eventually reaches a local minimum deeper than (S, P). In particular, we give a lower bound on the probability of reaching $E_{S,P}$ within a given number of iterations. Applying Theorem 3, we get a lower bound on the probability of remaining "near" $E_{S,P}$ over a given number of iterations. Combining these two observations in a limiting interval argument, we obtain the result of Theorem 4.

2.3.5 Climbing the Increasing Depths

In this section, we analyze the ability of the process to escape deepening cups. For a given cup C, define the step at which the process first escapes C as $W =$

$\inf\{k \geq 0 : Y(k) \in F\}$.

Theorem 2. *There exist* $\epsilon > 0$ *and* $0 < \xi \leq 1$ *depending only on* Θ *and* C *so that for every time* $t_0 \geq t^*$, *every* $x_0 \in C$ *where* C *has depth pair* (S, P), *every* $y_0 \in F$, *and every* $(\eta_t, \lambda_t : t \geq 0)$ *such that* $\eta(r_{t_0}) \leq \xi$ *and* $\sum_{k=0}^{\infty} \eta(d_{t_0+k}) = +\infty$,

$$(a)\ \ E\left[\sum_{k=0}^{W} \eta(d_{T(k)}) \,|\, (Y(0), T(0)) = (x_0, t_0)\right] \leq \frac{1}{\epsilon},\ \textit{and}$$

$$(b)\ \ P\left\{Y(W) = y_0 \,|\, (Y(0), T(0)) = (x_0, t_0)\right\} \geq \epsilon.$$

In [56], an analogous result is proved for the situation where pressure is constant over time. We generalize this result to handle deepening cups.

As in Theorem 3 of [56], Theorem 2 is proved by strong induction via a sequence of lemmas. As shown in Figure 2.4, the states in a cup C can be partitioned into its rim, R, and a collection of underlying cups, $C_1, C_2, \ldots, C_n$. Any of the cups nested within C have a depth pair $(S(C_i), P(C_i)) \prec (S(C), P(C))$. Furthermore, by the continuous increase property, when the process escapes a nested cup C_i, it must visit a state in R.

Induction Hypothesis: Theorem 2 is true for any cup C' with $(S(C'), P(C')) \prec (S(C), P(C))$.

Base Case: Cup C consists only of states at a single level, i.e., $C = R$.

The proof of the base case for Theorem 2 is similar to that of the induction step. We continue with the exposition of the general case and note that all the arguments hold in the base case (often trivially). We need some preliminaries before proceeding to the proof of Theorem 2.

To track the process when it enters the rim of C, we define $\{J_i\}$ such that $J_0 = 0$

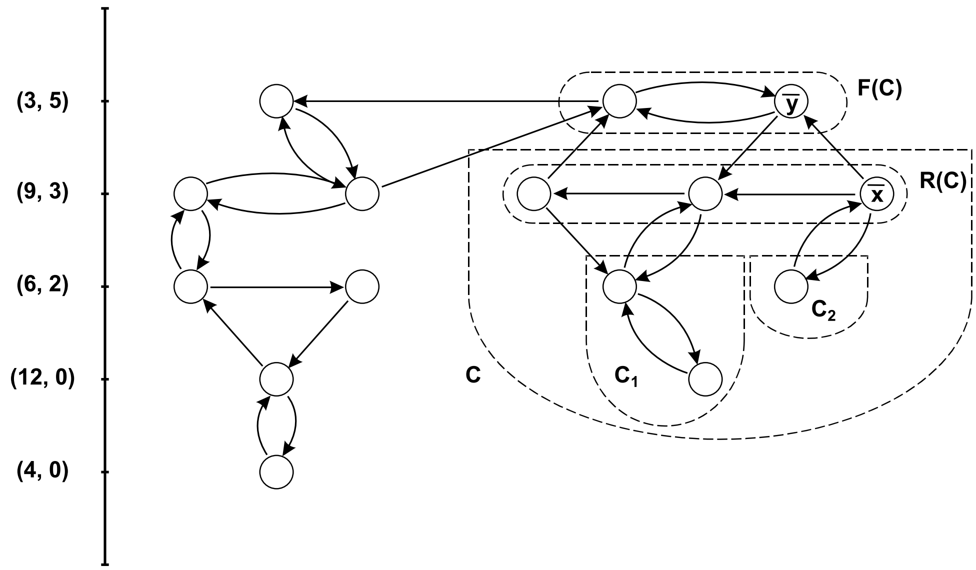

Figure 2.4: Cup C is partitioned into its rim $(R(C))$ and a collection of shallower cups (C_1 and C_2).

and $J_{i+1} = \inf\{k > J_i : Y(k) \in R\} \wedge W$ for $i \geq 0$. Observe that J_i is the iteration at which the process visits R for the i^{th} time or finally escapes the cup.

Since the dynamic depth of any cup C_i at time $t \geq t^*$ is no deeper than $g_t(C)$ which is less than $d_t(C)$, we can apply the induction hypothesis. Therefore, we get ϵ_i and ξ_i for each cup C_i such that if $\eta(r_{t_0}(C_i)) \leq \xi_i$ and $x_0 \in C_i$, then

$$E\left[\sum_{k=0}^{J_1} \eta(g_{T(k)})\right] \leq \frac{1}{\epsilon_i}.$$

(For notational convenience, we define all probabilities and expectations with respect to the probability measure induced when $(Y(0), T(0)) = (x_0, t_0)$ for $x_0 \in C_i$ and $t_0 \geq t^*$.

Furthermore, for every y such that $y \in R \cap N(x)$ for $x \in C_i$,

$$P\{Y(J_1) = y\} \geq \epsilon_i.$$

Therefore, we can set $\hat{\epsilon} = \min_i \epsilon_i$ and $\hat{\xi} = \min_i \xi_i$, so that if $\eta(r_{t_0}(C_i)) \leq \hat{\xi}$ and $x_0 \in C_i$ for some $i = 1, \ldots, n$, then

$$E\left[\sum_{k=0}^{J_1} \eta(g_{T(k)})\right] \leq \frac{1}{\hat{\epsilon}}$$

and for every y such that $y \in R \cap N(x)$ for $x \in C_i$,

$$P\left\{Y(J_1) = y\right\} \geq \hat{\epsilon}.$$

Let $K_0 = 0$ and define $K_{i+1} = (\inf\left\{k > K_i : Y(k) \in R^e\right\} \wedge W)$ for $i \geq 0$. Iteration K_i is the step at which the process visits R^e for the i^{th} time or finally escapes the cup.

We describe the history of the process up to step J_k through the σ-algebra $\mathcal{F}_k$ generated by $\{(Y(i) : 0 \leq i \leq J_k), T(0)\}$. Note that

$$\begin{aligned} \mathcal{F}_k &= \sigma\left\{(Y(i) : 0 \leq i \leq J_k), T(0)\right\} \\ &= \sigma\left\{(Y(i), T(i)) : 0 \leq i \leq J_k\right\}. \end{aligned}$$

That is, given the sequence of states and the initial clock time, we can compute the clock time at each step in the entire sequence.

In Lemma 2, we describe a bound on the probability of transitioning between two states in the rim of the cup, given that they can communicate without visiting other states in the rim.

Lemma 2. *Under the conditions of Theorem 2, there exists $\bar{\epsilon} > 0$ depending only on Θ and C so that, for $x, y \in R$ that communicate without visiting another state in R, $P\left\{Y(J_{k+1}) = y \mid \mathcal{F}_k\right\} \geq \bar{\epsilon}$ on the event $\{Y(J_k) = x\}$.*

Proof. For scenarios such that $\{Y(J_k) = x\}$, there are two possible ways for $\{Y(J_{k+1}) = y\}$ to occur. The trivial case is when $y \in N(x)$, since then $P\{Y(J_{k+1}) = y \mid \mathcal{F}_k\} \geq c_1$ on the event $\{Y(J_k) = x\}$.

The second case is when $y \notin N(x)$. In this case, we use the fact that Y must first visit a state in R upon exiting any of the cups nested within C. In order for $\{Y(J_{k+1}) = y, Y(J_k) = x\}$ to occur if $y \notin N(x)$, the process has to descend into some cup C_i and hit y upon first jump out. By assumption, the process can transition from x to y without visiting another state in R, so there must exist a path $x = x_0, x_1, \ldots, x_s = y$ such that $x_1, \ldots, x_{s-1} \in C_i$ and $x_{n+1} \in N(x_n)$ for $0 \leq n < s$. From the induction hypothesis of Theorem 2, the probability of escaping C_i by jumping to y is bounded below, i.e., $P\{Y(J_{k+1}) = y\} \geq \hat{\epsilon}$ on the event $\{Y(J_k + 1) = x_1 \in C_i\}$. Combining this with $P\{Y(J_k + 1) = x_1 \mid \mathcal{F}_k\} = \Gamma(x, x_1) \geq c_1$ on the event $\{Y(J_k) = x\}$, we obtain $P\{Y(J_{k+1}) = y \mid \mathcal{F}_k\} \geq c_1\hat{\epsilon}$ on the event $\{Y(J_k) = x\}$. Setting $\bar{\epsilon} \leq c_1\hat{\epsilon}$, we obtain the desired result.

Lemma 3 states that once in a cup, the expected number of visits to states in R until reaching a state in R^e is bounded above by a constant. Let $M = \inf\{j > 0 : Y(J_j) \in R^e \cup F\}$. The definition of M includes the set F to account for the case in which $Y(0) = x_0 \in R^e$ and $Y(1) \in F$.

Lemma 3. *Under the conditions of Theorem 2, for any $x_0 \in C$, $E[M] \leq c_1$ for some constant c_1 depending only on Θ and C.*

Proof. Without loss of generality, assume $Y(0) = x_0 \in C_i$ for some i. Note that it does not matter how long the process wanders around in the cup C beneath the rim R as long as we eventually hit a state in R since M is only incremented by visits to R. From the induction hypothesis of Theorem 2, we know that if the process is in any state $x \in C_i$ at any time $t \geq t^*$, then with probability one it will eventually hit a state in R. To see this, observe that $P\{Y \text{ gets stuck below } R\} \leq \lim_{n\to\infty}\{(1-\hat{\epsilon}_0)(1-\hat{\epsilon}_1)\cdots(1-\hat{\epsilon}_n)\} = 0$, where $\hat{\epsilon}_n$ is the lower bound of hitting some state $y \in R \cap N(x)$ for $x \in C_i$ after the n^{th} time we dip into the underlying cups. Now we show that on one of these "trips" to R, the process will hit a state in R^e.

By definition of a cup C and its rim R, any two states in C are accessible to each other at $(f(R), p(R))$. Therefore, there exists a path from any state $w \in R$ to $\bar{x} \in R^e \subseteq R$ that visits the set R no more than $|R|$ times and does not climb higher than level $(f(R), p(R))$. Denote the sequence of states in R visited on this path by $\{w = s_0, s_1, \ldots, s_m = \bar{x}\}$, where $m \leq |R|$.

Define $M_k(\bar{x}) = \inf\{j > 0 : Y(J_{k+j}) \in \{\bar{x}\} \cup F\}$. From Lemma 2, we know that $P\{Y(J_{k+1}) = s_{i+1} \mid \mathcal{F}_k\} \geq \bar{\epsilon}$ on the event $\{Y(J_k) = s_i\}$ for $0 \leq i < m$, independent of k. Therefore, applying Lemma 2 along the entire path from w to $\bar{x}$, we obtain $P\{M_k(\bar{x}) \leq |R| \mid \mathcal{F}_k\} \geq \bar{\epsilon}^{|R|} > 0$. Consequently, $P\{M_k(\bar{x}) > |R| \mid \mathcal{F}_k\} \leq 1 - \bar{\epsilon}^{|R|} < 1$.

Through a geometric trials argument, observe that

$$
\begin{aligned}
&P\{M_0(\bar{x}) > (n+1)|R|\} \\
&= P\left\{Y(J_0) \neq \bar{x}, \ldots, Y(J_{(n+1)|R|}) \neq \bar{x}\right\} \\
&= \sum_{w \in R\setminus\{\bar{x}\}} P\left\{Y(J_0) \neq \bar{x}, \ldots, Y(J_{(n)|R|}) = w, \ldots, Y(J_{(n+1)|R|}) \neq \bar{x}\right\} \\
&= \sum_{w \in R\setminus\{\bar{x}\}} P\left\{Y(J_{n|R|+1}) \neq \bar{x}, \ldots, Y(J_{(n+1)|R|}) \neq \bar{x} \,|\, Y(J_{(n)|R|}) = w, \ldots, Y(J_0) \neq \bar{x}\right\} \\
&\quad P\left\{Y(J_{(n)|R|}) = w, \ldots, Y(J_0) \neq \bar{x}\right\} \\
&\leq (1 - \bar{\epsilon}^{|R|}) \sum_{w \in R\setminus\{\bar{x}\}} P\left\{Y(J_{(n)|R|}) = w, \ldots, Y(J_0) \neq \bar{x}\right\} \\
&= (1 - \bar{\epsilon}^{|R|}) P\{M_0(\bar{x}) > n|R|\}.
\end{aligned}
$$

Repeating this argument n more times yields

$$
P\{M_0(\bar{x}) > (n+1)|R| \,|\, \mathcal{F}_0\} \leq (1 - \bar{\epsilon}^{|R|})^{n+1}.
$$

By definition, $M_0(\bar{x}) \geq M$. Hence,

$$
\begin{aligned}
E[M] &\leq E[M_0(\bar{x})] \\
&= \sum_{m=0}^{\infty} P\{M_0(\bar{x}) > m\} \\
&= \sum_{k=0}^{\infty} \sum_{r=0}^{|R|-1} P\{M_0(\bar{x}) > k|R| + r\} \\
&\leq \sum_{k=0}^{\infty} \sum_{r=0}^{|R|-1} P\{M_0(\bar{x}) > k|R|\} \\
&\leq |R| \sum_{k=0}^{\infty} (1 - \bar{\epsilon}^{|R|})^k \\
&= \frac{|R|}{\bar{\epsilon}^{|R|}} \\
&\equiv D_3.
\end{aligned}
$$

Lemma 4 essentially states that the expected jump probabilities accumulated between entrances into the cup's rim until hitting an "escape state" is bounded above by a constant.

Lemma 4. *Under the conditions of Theorem 2, there exists a constant c_2 depending only on Θ and C such that*

$$E\left[\sum_{s=J_k+1}^{J_{k+1}} \eta(g_{T(s)}) I\left\{s \leq K_1\right\} \mid \mathcal{F}_k\right] \leq c_2.$$

Proof. Note that $Y(J_k+1)$ is the state to which the system transitions immediately after visiting R for the k^{th} time, while $Y(J_{k+1})$ is the state of the process when in $R \cup F$ for the $(k+1)$st time. Considering the two scenarios, $Y(J_k+1) \notin R \cup F$ or $Y(J_k+1) = Y(J_{k+1}) \in R \cup F$, the result follows from the induction hypothesis.

$$\begin{aligned}
& E\left[\sum_{s=J_k+1}^{J_{k+1}} \eta(g_{T(s)}) I_{\{s \leq K_1\}} \mid \mathcal{F}_k\right] \\
&= E\left[\eta(g_{J_{k+1}}) I_{\{J_{k+1} \leq K_1\}} I_{\{Y(J_k+1) \in R \cup F\}} \mid \mathcal{F}_k\right] + \\
&\quad E\left[I_{\{Y(J_k+1) \notin R \cup F\}} \sum_{s=J_k+1}^{J_{k+1}} \eta(g_{T(s)}) I_{\{s \leq K_1\}} \mid \mathcal{F}_k\right] \\
&\leq P\left\{Y(J_k+1) \in R \cup F \mid \mathcal{F}_k\right\} + \frac{1}{\epsilon} P\left\{Y(J_k+1) \notin R \cup F \mid \mathcal{F}_k\right\} \\
&\leq 1 + \frac{1}{\epsilon} \\
&\equiv c_2.
\end{aligned}$$

Lemma 5 states that the expected jump probabilities accumulated until reaching a state in the cup's rim that has a neighbor outside the cup is bounded by a constant times the probability of accepting a jump from the rim of the cup to the froth at time t_0.

Lemma 5. *Under the conditions of Theorem 2, there is a constant c_3 depending only on Θ and C such that for every $x_0 \in C$, and every $t_0 \geq t^*$ with $\eta(r_{t_0}) \leq \xi$,*

$$E\left[\sum_{s=0}^{K_1} \eta(d_{T(s)})\right] \leq c_3 \eta(r_{t_0}).$$

Proof. Consider the stochastic process $(Z_k : k \geq 0)$, where $Z_0 = 0$ and for $k > 0$

$$Z_k = M \wedge k - \frac{1}{c_2} \sum_{s=0}^{J_k} \eta(g_{T(s)}) I\left\{s \leq K_1\right\}.$$

With an application of Lemma 4 for the case when $k < M$, and the fact that $K_1 = J_M$ for $k \geq M$, it is straightforward to show that $(Z_k : k \geq 0)$ is a submartingale with respect to $(\mathcal{F}_k : k \geq 0)$. That is,

$$\begin{aligned}
&E\left[Z_{k+1} \mid \mathcal{F}_k\right] \\
&= E\left[M \wedge (k+1) - \frac{1}{c_2} \sum_{s=0}^{J_{k+1}} \eta(g_{T(s)}) I\left\{s \leq K_1\right\} \mid \mathcal{F}_k\right] \\
&= E\left[M \wedge (k+1) - \frac{1}{c_2} \sum_{s=0}^{J_k} \eta(g_{T(s)}) I\left\{s \leq K_1\right\} - \frac{1}{c_2} \sum_{s=J_k+1}^{J_{k+1}} \eta(g_{T(s)}) I\left\{s \leq K_1\right\} \mid \mathcal{F}_k\right] \\
&= \begin{cases} Z_k + 1 - \frac{1}{c_2} E\left[\sum_{s=J_k+1}^{J_{k+1}} \eta(g_{T(s)}) I\left\{s \leq K_1\right\} \mid \mathcal{F}_k\right] & \text{on the event } \{k < M\} \\ Z_k - \frac{1}{c_2} E\left[\sum_{s=J_k+1}^{J_{k+1}} \eta(g_{T(s)}) I\left\{s \leq K_1\right\} \mid \mathcal{F}_k\right] & \text{on the event } \{k \geq M\} \end{cases} \\
&\geq Z_k \text{ for all } k.
\end{aligned}$$

Therefore, for all k, $E\left[Z_k\right] \geq E\left[Z_0\right] = 0$, which implies

$$E\left[\frac{1}{c_2} \sum_{s=0}^{J_k} \eta(g_{T(s)}) I\left\{s \leq K_1\right\}\right] \leq E[M \wedge k] \leq E[M] \; \forall \; k.$$

Note that $X_k = \frac{1}{c_2} \sum_{s=0}^{J_k} \eta(g_{T(s)}) I\left\{s \leq K_1\right\}$ for $k = 0, 1, 2, \ldots$ is a sequence of non-negative random variables, non-decreasing in k. Additionally, $\lim_{k \to \infty} X_k =$

$\frac{1}{c_2}\sum_{s=0}^{K_1}\eta(g_{T(s)})$ almost surely. By the monotone convergence theorem,

$$\begin{aligned} & E\left[\frac{1}{c_2}\sum_{s=0}^{K_1}\eta(g_{T(s)})\right] \\ = \; & E\left[\lim_{k\to\infty}\frac{1}{c_2}\sum_{s=0}^{J_k}\eta(g_{T(s)})I\left\{s\le K_1\right\}\right] \\ = \; & \lim_{k\to\infty} E\left[\frac{1}{c_2}\sum_{s=0}^{J_k}\eta(g_{T(s)})I\left\{s\le K_1\right\}\right] \\ \le \; & \limsup_{k\to\infty} E\left[\frac{1}{c_2}\sum_{s=0}^{J_k}\eta(g_{T(s)})I\left\{s\le K_1\right\}\right] \\ \le \; & E\left[M\right]. \end{aligned}$$

Applying Lemma 3,

$$E\left[\sum_{s=0}^{K_1}\eta(g_{T(s)})\right] \le c_2 c_1 \equiv c_3. \tag{2.7}$$

Now observe that since $\eta(r_t)$ is decreasing for all $t\ge t^*$, and $T(0)=t_0\ge t^*$,

$$\frac{\eta(d_{T(s)})}{\eta(r_{t_0})} = \frac{\eta(r_{T(s)})\eta(g_{T(s)})}{\eta(r_{t_0})} \le \eta(g_{T(s)}).$$

The result now follows from (2.7).

In Lemma 6, we show that the expected jump probabilities accumulated over several visits to the set R^e is bounded above by a constant.

Lemma 6. *Under the conditions of Theorem 2, for* $1\le i\le j<+\infty$

$$E\left[\sum_{s=K_i+1}^{K_j}\eta(d_{T(s)})\,|\,Y(K_{i-1}),T(K_{i-1})\right] \le \frac{c_3}{c_0}.$$

Proof. The proof is by reverse induction. The result is trivial if $i=j$. So suppose it is true for $i+1$ with $1\le i+1\le j$. From Lemma 5 we have that

$$E\left[\sum_{s=K_i+1}^{K_{i+1}}\eta(d_{T(s)})\,|\,Y(K_i),T(K_i)\right] \le c_3\eta(r_{T(K_i)})$$

on the event that $\{Y(K_i) \in C\}$ and is 0 otherwise. The inductive hypothesis gives that

$$E\left[\sum_{s=K_{i+1}+1}^{K_j} \eta(d_{T(s)}) \mid Y(K_i), T(K_i)\right] \leq \frac{c_3}{c_0}$$

on the event $Y(K_i) \in C$ and is 0 otherwise. Thus

$$\begin{aligned}
& E\left[\sum_{s=K_i+1}^{K_j} \eta(d_{T(s)}) \mid Y(K_{i-1}), T(K_{i-1})\right] \\
&= E\left[E\left[\sum_{s=K_i+1}^{K_j} \eta(d_{T(s)}) \mid Y(K_i), T(K_i), Y(K_{i-1}), T(K_{i-1})\right] \mid Y(K_{i-1}), T(K_{i-1})\right] \\
&\leq E\left[\left(c_3\eta(r_{T(K_i)}) + \frac{c_3}{c_0}\right) I(Y(K_i) \in C) \mid Y(K_{i-1}), T(K_{i-1})\right] \\
&\leq E\left[\left(c_3\eta(r_{T(K_{i-1})}) + \frac{c_3}{c_0}\right) I(Y(K_i) \in C) \mid Y(K_{i-1}), T(K_{i-1})\right] \\
&\leq \frac{c_3}{c_0} P\{Y(K_i) \in C \mid Y(K_{i-1}), T(K_{i-1})\} + c_3\eta(r_{T(K_{i-1})})) \\
&\leq \frac{c_3}{c_0}(1 - c_0\eta(r_{T(K_{i-1})})) + c_3\eta(r_{T(K_{i-1})})).
\end{aligned}$$

The final inequality follows since $P(Y(K_i) \notin C \mid Y(K_{i-1}), T(K_{i-1}))$ is the probability of directly jumping out of the cup C from $Y(K_{i-1})$, which is bounded below by $c_0 r_{T(K_{i-1})}$.

Proof of Theorem 2(a). Setting $i = 1$ in Lemma 6 and letting j tend to infinity, we apply the monotone convergence theorem to get

$$\begin{aligned}
E\left[\sum_{s=K_1+1}^{W} \eta(d_{T(s)})\right] &= E\left[\lim_{j\to\infty} \sum_{s=K_1+1}^{K_j} \eta(d_{T(s)})\right] \\
&= \lim_{j\to\infty} E\left[\sum_{s=K_1+1}^{K_j} \eta(d_{T(s)})\right] \\
&\leq \limsup_{j\to\infty} E\left[\sum_{s=K_1+1}^{K_j} \eta(d_{T(s)})\right] \\
&\leq \frac{c_3}{c_0}.
\end{aligned}$$

Together with Lemma 5, we get

$$\begin{aligned} E\left[\sum_{s=0}^{W} \eta(d_{T(s)})\right] &= E\left[\sum_{s=0}^{K_1} \eta(d_{T(s)})\right] + E\left[\sum_{s=K_1+1}^{W} \eta(d_{T(s)})\right] \\ &\leq c_3\eta(r_{t_0}) + \frac{c_3}{c_0} \\ &\leq c_3 + \frac{c_3}{c_0}. \end{aligned}$$

□

Now we proceed to the proof of Theorem 2(b). Let $\bar{y} \in F$ be fixed. Choose a state $\bar{x} \in R^e \subset R$ with $\bar{y} \in N(\bar{x}) \cap F$. Define $L^* = \inf\{l \geq 0 : Y(J_l) \in F\}$. We track the process entrances into $\bar{x}$ by letting $L_0 = 0$ and $L_{i+1} = \inf\{l > L_i : Y(J_l) = \bar{x}\} \wedge L^*$ for $i \geq 0$. Note that J_{L_i} is the iteration at which the process enters $\bar{x}$ for the i^{th} time or, if $L_i = L^*$, the iteration at which the process hits a state in F. Observe that the events $\{J_{L_i} = W\}$ and $\{L_i = L^*\}$ are identical, implying that $Y(W) = Y(J_{L^*})$.

Lemma 7 shows that the number of times Y visits states in $R\backslash\{\bar{x}\}$ between the i^{th} and $i+1^{\text{st}}$ visit to $\{\bar{x}\} \cup F$ is bounded above by a constant. The proof is similar to that of Lemma 3 and omitted.

Lemma 7. *Under the conditions of Theorem 2, for any* $x_0 \in C$,

$$E\left[(L_{i+1} - L_i - 1)^+ \mid \mathcal{F}_{L_i}\right] \leq c_4$$

for some constant c_4 *depending only on* Θ *and* C.

Proof. Observe that $(L_{i+1} - L_i - 1)^+ = M_{L_i}(\bar{x}) = \inf\{j > 0 : Y(J_{L_i+j}) \in \{\bar{x}\} \cup F\}$ where $M_k(\bar{x})$ is defined in Lemma 3. The result follows from logic similar to the proof of Lemma 3.

Define the events $(H_i : i \geq 0)$ by $H_0 = \{J_{L_1} = W\}$ and for $i \geq 1$, $H_i = \{J_{L_i} + 1 < W = J_{L_{i+1}}\}$. The event H_0 occurs if the process escapes the cup without ever visiting $\bar{x}$. The event H_i occurs when the process visits $\bar{x}$ i times before exiting the cup, but escapes the cup through another state in R^e. Define the events $(G_i : i \geq 0)$ by $G_i = \{J_{L_i} + 1 = W\}$. The event G_0 occurs if $Y(0) = x_0 = \bar{x}$ and $Y(1) \in F$. For $i \geq 1$, the event G_i occurs when the process escapes the cup immediately after visiting $\bar{x}$ for the i^{th} time. In Lemma 8 below, we give bounds on the probabilities of these events.

Lemma 8. *Under the conditions of Theorem 2, the following inequalities hold for c_4 as defined in Lemma 7 and some constant $c_5 > 0$ depending only on Θ and C:*

$$P\{H_0\} \leq c_4 \eta(r_{t_0}), \tag{2.8}$$

and

$$P\{H_i\} \leq c_5 P\{G_i\} \tag{2.9}$$

for $i \geq 1$.

Proof. Consider the stochastic process $(\Upsilon_k : k \geq 0)$, where $\Upsilon_0 = 0$ and for $k \geq 1$,

$$\Upsilon_k = I\{Y(J_k) \in F\} - \sum_{j=0}^{(k \wedge L^*)-1} \eta(r_{T(J_j)}). \tag{2.10}$$

Observing that $P\{Y(J_{k+1}) \in F \mid \mathcal{F}_k\} \leq \eta(r_{T(J_k)})$ on the event $\{k < L^*\}$ and $Y(J_k) \in F$ for all $k \geq L^*$, it is straightforward to show that $(\Upsilon_k : k \geq 0)$ is a supermartingale with respect to $\mathcal{F}_k$. That is,

$$
\begin{aligned}
E\left[\Upsilon_{k+1} \,|\, \mathcal{F}_k\right] &= E\left[I\left\{Y(J_{k+1}) \in F\right\} | \mathcal{F}_k\right] - E\left[\sum_{j=0}^{(k+1\wedge L^*)-1} \eta(r_{T(J_j)}) \,|\, \mathcal{F}_k\right] \\
&\leq \begin{cases} \eta(r_{T(J_k)}) - \sum_{j=0}^{k} \eta(r_{T(J_j)}) & \text{on the event } \{k < L^*\} \\ 1 - \sum_{j=0}^{L^*-1} \eta(r_{T(J_j)}) & \text{on the event } \{k \geq L^*\} \end{cases} \\
&= \begin{cases} I\left\{Y(J_k) \in F\right\} - \sum_{j=0}^{(k\wedge L^*)-1} \eta(r_{T(J_j)}) & \text{on the event } \{k < L^*\} \\ I\left\{Y(J_k) \in F\right\} - \sum_{j=0}^{(k\wedge L^*)-1} \eta(r_{T(J_j)}) & \text{on the event } \{k \geq L^*\} \end{cases} \\
&= \Upsilon_k
\end{aligned}
$$

Note that since $I\left\{Y(J_k) \in F\right\} = I\left\{J_k = W\right\}$, the event H_i occurs if and only if $I\left\{Y(J_{L_{i+1}}) \in F\right\} - I\left\{Y(J_{L_i+1}) \in F\right\} = 1 - 0 = 1$. From the optional sampling theorem for supermartingales,

$$
\begin{aligned}
&E\left[\Upsilon_{L_{i+1}} - \Upsilon_{L_i+1} \,|\, \mathcal{F}_{L_i}\right] \\
&= E\left[I\left\{Y(J_{L_{i+1}}) \in F\right\} - I\left\{Y(J_{L_i+1}) \in F\right\} - \sum_{j=(L_i+1\wedge L^*)}^{(L_{i+1}\wedge L^*)-1} \eta(r_{T(J_j)}) \,|\, \mathcal{F}_{L_i}\right] \\
&\leq 0.
\end{aligned}
$$

Noting that $\eta(r_k)$ is decreasing in k, we get

$$
\begin{aligned}
P\left\{H_i \,|\, \mathcal{F}_{L_i}\right\} &= E\left[I\left\{Y(J_{L_{i+1}}) \in F\right\} - I\left\{Y(J_{L_i+1}) \in F\right\} | \mathcal{F}_{L_i}\right] \\
&\leq E\left[\sum_{j=((L_i+1)\wedge L^*)}^{(L_{i+1}\wedge L^*)-1} \eta(r_{T(J_j)}) \,|\, \mathcal{F}_{L_i}\right] \\
&\leq c_4 \eta(r_{T(J_{L_i})})
\end{aligned}
$$

where c_4 is defined in Lemma 7. For $i = 0$, this implies (2.8).

For $i \geq 1$, note that

$$P\{G_i \mid \mathcal{F}_{L_i}\} = I\{Y(J_{L_i}) = \bar{x}\} \sum_{z \in N(\bar{x}) \cap F} \Gamma(\bar{x}, z) \eta(r_{T(J_{L_i})}) \geq c_0 \eta(r_{T(J_{L_i})}) I\{Y(J_{L_i}) = \bar{x}\}.$$

Combining the bounds on G_i and H_i, we obtain

$$P\{G_i \mid \mathcal{F}_{L_i}\} \geq \frac{c_0}{c_4} I\{Y(J_{L_i}) = \bar{x}\} P\{H_i \mid \mathcal{F}_{L_i}\}.$$

Observe that $P\{H_i | \mathcal{F}_{L_i}\} = 0$ off the event $\{Y(J_{L_i}) = \bar{x}\}$. That is, when $Y(J_{L_i}) \neq \bar{x}$, then $Y(J_{L_i}) \in F$, and thus $J_{L_i} + 1 > W$. Therefore, we can drop the $I\{Y(J_{L_i}) = \bar{x}\}$ term and take expectations to obtain (2.9), where $c_5 = \frac{c_4}{c_0}$.

Proof of Theorem 2(b). Observe that if $Y(J_{L^*-1}) = \bar{x}$, then the process escapes the cup C by jumping from $\bar{x} \in R^e$. For the event $\{Y(J_{L^*-1}) = \bar{x}\}$ to occur, $\{J_{L_i} + 1 = W\}$ for some i. Therefore, by Lemma 8, we get

$$\begin{aligned} P\{Y(J_{L^*-1}) = \bar{x}\} &= \sum_{i=0}^{\infty} P\{G_i\} \\ &\geq \frac{1}{c_5} \sum_{i=1}^{\infty} P\{H_i\} \\ &\geq \frac{1}{c_5} (1 - P\{H_0\} - P\{Y(J_{L^*-1}) = \bar{x}\}) \\ &\geq \frac{1}{c_5} (1 - c_4 \eta(r_{t_0}) - P\{Y(J_{L^*-1}) = \bar{x}\}). \end{aligned}$$

Combining like terms, we obtain

$$P\{Y(J_{L^*-1}) = \bar{x}\} \geq \frac{1 - c_4 \eta(r_{t_0})}{1 + c_5}.$$

The probability of hitting $\bar{y}$ given we exit the cup via $\bar{x}$ is given by

$$\begin{aligned} P\{Y(W) = \bar{y} \mid Y(J_{L^*-1}) = \bar{x}\} &= \frac{\Gamma(\bar{x}, \bar{y})}{\sum_{w \in N(\bar{x}) \cap F} \Gamma(\bar{x}, w)} \\ &\geq c_0. \end{aligned}$$

We conclude that

$$\begin{aligned}
P\{Y(W)=\bar{y}\} &\geq P\{Y(W)=\bar{y}, Y(J_{L^*-1})=\bar{x}\} \\
&= P\{Y(W)=\bar{y} \mid Y(J_{L^*-1})=\bar{x}\} P\{Y(J_{L^*-1})=\bar{x}\} \\
&\geq \frac{c_0(1-c_4\eta(r_{t_0}))}{(1+c_5)}.
\end{aligned}$$

Setting

$$\epsilon = \min\left\{\frac{c_0}{c_3(1+c_0)}, \frac{c_0(1-c_4\eta(r_{t_0}))}{(1+c_5)}\right\}$$

with the caveat that $\eta(r_{t_0}) \leq \xi \leq \frac{1}{c_4}$, we complete the proof of Theorem 2 by the principle of induction. □

2.3.6 Sinking to the Bottom

In this section, we analyze the behavior of the process once it has reached a local minimum. In particular, we are interested in the probability of the process remaining in a cup C with depth pair (S, P) for q iterations. The arguments are based on an analogous result in [56] for the continuous-time case with static cup depths.

Theorem 3. *There exist* $0 < \bar{\xi} \leq 1$ *and* $c_6 > 0$ *depending only on* Θ *such that for any cup* C *with bottom* B*, any* $x_0 \in B$ *and any* $t_0 \geq t^*$ *with* $\max\{\eta_{t_0}, \eta(d_{t_0})\} \leq \bar{\xi}$*,*

$$P\{Y(k) \in C \text{ for } 0 \leq k \leq q \mid (Y(0), T(0)) = (x_0, t_0)\} \geq \exp\left(-c_6 \sum_{k=0}^{q-1} \eta(d_{T(k)})\right).$$

Since there are only a finite number of cups in $\mathcal{S}$, it suffices to prove Theorem 3 for an arbitrary cup C. Recall that the highest level in a cup C is $(f(y), p(y))$, for any $y \in R$, the rim of the cup. For such y, if $(f(y), p(y)) = (f_\ell, p_\ell)$, then the cup C is the entire state space $\mathcal{S}$. Theorem 3 is then trivial.

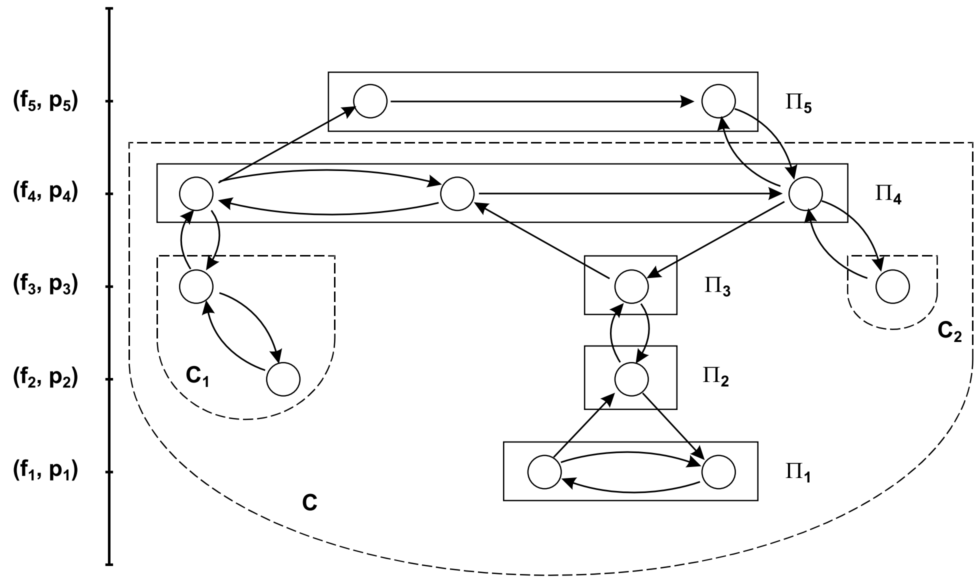

Figure 2.5: Cup C is partitioned into $D = \cup_{i=1}^{4} \Pi_i$ and a collection of smaller cups (C_1 and C_2).

So, assume that the highest level in C is $(f_r, p_r) \prec (f_\ell, p_\ell)$. For notational ease we further assume, without loss of generality, that $(f(x), p(x)) = (f_1, p_1)$. The dynamic depth of the cup C at time $t \geq t^*$ is then given by $d_t = (f_r - f_1) + \lambda_t(p_r - p_1)$.

Given a state $x \in C$, let $(G(x), Q(x))$ be the lowest level at which a state in B is accessible from x. As shown in Figure 2.5, we can partition the cup C into the set D and cups $C_1, \ldots, C_m$, where $D = \{x \in C : (f(x), p(x)) = (G(x), Q(x))\}$. That is, the set D contains the states in C that can reach the bottom of the cup without "climbing" to a higher level. In fact, if $x \in C$ and $(f(x), p(x)) \prec (G(x), Q(x))$, then $x \in C_i$ for some i, and C_i is the set of all states that are accessible from x at a level strictly lower than $(G(x), Q(x))$.

Define $\Pi_{r+1} = F$, and $\Pi_i = \{x \in D : (f(x), p(x)) = (f_i, p_i)\}$ for $i = 1, \ldots, r$. The set Π_i exclusively includes states in the set D at level (f_i, p_i), i.e., $\Pi_i = D \cap L(f_i, p_i)$

for $i \leq r$. Furthermore, let the distance between adjacent levels (f_{j+1}, p_{j+1}) and (f_j, p_j) at iteration k be denoted by $\Delta_j(T(k)) = (f_{j+1} - f_j + \lambda_{T(k)}(p_{j+1} - p_j))$ for $1 \leq j \leq r$.

Define $\delta = \min\{m : Y(m) \in \Pi_1 \cup \cdots \cup \Pi_{j-1} \cup \Pi_{j+1}\}$. Given $Y(0) = x \in \Pi_j$, δ denotes the first iteration that the process changes levels within D. Denote the number of iterations that the process spends in Π_j before changing levels in D as $\vartheta = |\{n : 0 \leq n \leq \delta \text{ and } Y(n) \in \Pi_j\}|$. We now present a couple of lemmas dealing with these new constructs.

Lemma 9. *For $2 \leq j \leq r$, let $\bar{x} \in \Pi_j$. Under the conditions of Theorem 3, for any $\bar{t} \geq t^*$, $E[\vartheta \mid (Y(0), T(0)) = (\bar{x}, \bar{t})] \leq c_7$ for some constant c_7 depending only on Θ and C.*

Proof. The partitioning of the cup C into sets $\{D, C_1, \ldots, C_m\}$ instills the property that Y visits a state in D upon escaping a cup C_i. By definition, for every state in Π_j, there exists a state in $\Pi_1 = B$ that is accessible at level (f_j, p_j). Therefore, we can construct a path from any state $\bar{x} \in \Pi_j$ to some state $y \in (\Pi_1 \cup \cdots \cup \Pi_{j-1})$ that visits the set Π_j no more than $|\Pi_j|$ times and does not climb higher than level (f_j, p_j). By construction, each state along the path between $\bar{x}$ to y must be either in Π_j or C_i for some $i = 1, \ldots, n$. Denote the sequence of states on Π_j by $\{\bar{x} = s_0, s_1, \ldots, s_{m-1} = z\}$, where $s_m = y \in N(z)$ and $m \leq |\Pi_j|$.

To analyze the behavior of the process along the path between $\bar{x}$ and y, we define $X_0 = 0$ and $X_{k+1} = \inf\{n > X_k : Y(n) \in \Pi_j\} \wedge \delta$ for $k \geq 0$. Observe that X_k is the iteration upon which the process is in Π_j for the k^{th} time until changing levels within D. Let $\mathcal{H}_k$ as the σ-algebra generated by $(Y(i) : 0 \leq i \leq X_k, T(0))$.

Along the sequence $\{\bar{x} = s_0, s_1, \ldots, s_{m-1}\}$, there are two ways that a transition can be made from s_i to s_{i+1} for $i \leq m-2$. If $s_{i+1} \in N(s_i)$, then $P\{Y(X_{k+1}) = s_{i+1} | \mathcal{H}_k\} \geq c_1$ on the event $\{Y(X_k) = s_i\}$, independent of k. If $s_{i+1} \notin N(s_i)$, then the transition from s_i to s_{i+1} must be made by descending into some C_i and climbing out via s_{i+1}. Note that upon descending into C_i for some i, there are two possibilities, either the process gets stuck in cup C_i and never escape, or it eventually climbs out. Therefore,

$$P\{Y(X_{k+1}) = s_{i+1}, X_{k+1} < \infty | \mathcal{H}_k\} \geq \epsilon \tag{2.11}$$

on the event $\{Y(X_k) = s_i, X_k < \infty\}$ for some $\epsilon > 0$, independent of k. Note that we can choose ϵ to ensure that $\epsilon < c_1$.

Consider a Markov chain on the state space $\Pi_j \cup C_i$, where C_i is an absorbing state for some i. Let P be a transition matrix defined by

$$\begin{aligned} P(x,y) &= \begin{cases} \epsilon & \text{if (2.11) holds for } x, y \in \Pi_j \\ 0 & \text{otherwise,} \end{cases} \\ P(x, C_i) &= 1 - \sum_{y \in \Pi_j} P(x,y), \text{ for } x \in \Pi_j, \\ P(C_i, C_i) &= 1. \end{aligned}$$

By definition, ϑ is not incremented during iterations spent in C_i. Therefore, ϑ is the time to absorption, known to be finite for such Markov chains [80].

If the process climbs out of every C_i along the path from $\bar{x}$ to y, we can apply the bounds on the transition probabilities to determine that $P\{\vartheta > k + |\Pi_j| | \mathcal{H}_k\} \leq 1 - \epsilon^{|\Pi_j|} < 1$. A standard geometric trials argument (similar to Lemma 3) shows that ϑ has a geometrically decaying tail so that $E[\vartheta | (Y(0), T(0)) = (\bar{x}, \bar{t})] \leq c_7$.

Lemma 10. *For $2 \leq j \leq r$, let $\bar{x} \in \Pi_j$. Under the conditions of Theorem 3, for any $\bar{t} \geq t^*$,*

$$P\{\delta < \infty, Y(\delta) \in \Pi_{j+1} \mid (Y(0), T(0)) = (\bar{x}, \bar{t})\} \leq c_7 \eta(\Delta_j(\bar{t})),$$

where c_7 is defined in Lemma 9.

Proof. Consider the stochastic process $(\Psi_k : k \geq 0)$, where $\Psi_0 = 0$ and for $k \geq 1$,

$$\Psi_k = I\{\delta \leq k, Y(k \wedge \delta) \in \Pi_{j+1}\} - \sum_{s=0}^{(k\wedge\delta)-1} \eta(\Delta_j(T(s))) I\{Y(s) \in \Pi_j\}.$$

Note that $P\{Y(k+1) \in \Pi_{j+1} \mid \mathcal{O}_k\} \leq \eta(\Delta_j(T(k))) I\{Y(k) \in \Pi_j\}$ on the event $\{k < \delta\}$, where $\mathcal{O}_k = \sigma(Y(i) : 0 \leq i \leq k, T(0))$. Using this observation, it is straightforward to show that $(\Psi_k : k \geq 0)$ is a supermartingale with respect to the filtration $\{\mathcal{O}_k : k \geq 0\}$.

As $k \to \infty$, $\Psi_k \to I\{\delta < \infty, Y(\delta) \in \Pi_{j+1}\} - \sum_{s=0}^{\delta-1} \eta(\Delta_j(T(s))) I\{Y(s) \in \Pi_j\}$ almost surely. In addition,

$$\begin{aligned} |\Psi_k| &= |I\{\delta \leq k, Y(k \wedge \delta) \in \Pi_{j+1}\} - \sum_{s=0}^{(k\wedge\delta)-1} \eta(\Delta_j(T(s))) I\{Y(s) \in \Pi_j\}| \\ &\leq 1 + \sum_{s=0}^{(k\wedge\delta)-1} I\{Y(s) \in \Pi_j\} \\ &= 1 + \vartheta. \end{aligned}$$

Since $E[1 + \vartheta] \leq 1 + D_{10} < \infty$ by Lemma 9, we can apply the dominated convergence theorem to get

$$\begin{aligned} & P\{\delta < \infty, Y(\delta) \in \Pi_{j+1} \mid (Y(0), T(0)) = (\bar{x}, \bar{t})\} \\ &\leq E\left[\sum_{s=0}^{\delta-1} \eta(\Delta_j(T(s))) I\{Y(s) \in \Pi_j\} \mid (Y(0), T(0)) = (\bar{x}, \bar{t})\right] \\ &\leq c_7 \eta(\Delta_j(\bar{t})). \end{aligned}$$

We stochastically bound the level process in our setting with a time-homogeneous, discrete-time Markov chain. A bound on a certain probability will then follow. We first give the discrete-time Markov chain result, which is essentially identical to Lemma 4.2 in [56], and so the proof is omitted.

For cup C, define $\Delta(t) = \min\{\Delta_i(t) : 0 \le i \le r\}$ for some fixed clock time $t \ge t^*$. Choose $\bar{\xi}$ with $0 < \bar{\xi} < 1$ so small that $c_7 \bar{\xi}^{\Delta(t)} \le \frac{1}{2}$. That is, $\bar{\xi} \le (2c_7)^{-\frac{1}{\Delta(t)}}$.

Lemma 11. *Fix $t \ge t^*$ as above, and consider a time-homogeneous discrete time Markov chain, Z, with state space $\{1, 2, \ldots, r+1\}$, and one-step transition probabilities*

$$p_{ij} = \begin{cases} c_7 \rho^{\Delta_i(t)} & \text{if } j = i+1 \\ 1 - c_7 \rho^{\Delta_i(t)} & \text{if } j = i-1 \\ 0 & \text{otherwise,} \end{cases}$$

for $2 \le i \le r$, c_7 as defined in Lemma 9, and $0 < \rho < 1$. Assume that states 1 and $r+1$ are absorbing. Define q_i for $2 \le i \le r$ by

$$q_i = P\{Z \text{ hits } r+1 \text{ before it hits } i-1 \mid Z_0 = i\}.$$

Then, for any $\rho \le \bar{\xi}$,

$$q_2 \le (2c_7)^{r-1} \rho^{(f_{r+1} - f_2 + \lambda_t(p_{r+1} - p_2))}.$$

We can now bound the probability that, starting one level above the bottom, the (Y, T) process escapes the cup before hitting the cup bottom.

Lemma 12. *Fix $\bar{t} \ge t^*$ so that $\eta_{\bar{t}} \le \bar{\xi}$ and suppose that $\bar{x} \in \Pi_2$. Then, under the conditions of Theorem 3,*

$$\begin{aligned} & P\{\, Y \text{ visits } F \text{ before visiting } B \mid (Y(0), T(0)) = (\bar{x}, \bar{t})\} \\ \le\ & (2c_7)^{r-1} \eta_{\bar{t}}(f_{r+1} - f_2 + \lambda_{\bar{t}}(p_{r+1} - p_2)) \end{aligned}$$

Proof. Define $V_0 = 0$ and

$$V_{i+1} = \inf\{n \geq V_i : Y(n) \in D \text{ and } (f(Y(n)), p(Y(n))) \neq (f(Y(V_i)), p(Y(V_i)))\} \wedge W.$$

Note that V_i is the iteration that the process changes levels within the set D for the i^{th} time. We focus on $(V_i : i \geq 0)$ since the process can only escape the cup or hit the bottom through a state in D.

Define a function ζ on $\mathcal{S}$ such that for $x \in L(f_i, p_i)$, $\zeta(x) = i$. Consider the process $(\zeta(Y(V_i)) : i \geq 0)$. For example, $\zeta(Y(V_0)) = \zeta(\bar{x}) = 2$. Using Lemma 10, we determine that

$$P\{\zeta(Y(V_{i+1})) = j + 1 \,|\, \zeta(Y(V_i)) = j, T(V_i) = \bar{t}\} \leq c_7 \eta(\Delta_j(\bar{t})).$$

Setting $\rho \geq \eta_{\bar{t}}$ in Lemma 11, we see that $(Z_k : k \geq 0)$ stochastically dominates $(\zeta(Y(V_i)) : \; i \geq 0)$. We conclude that $(\zeta(Y(V_i)) : i \geq 0)$ has a smaller chance of escaping the cup before hitting the bottom than $(Z_i, i \geq 0)$.

Proof of Theorem 3. Assume that $Y(0) = x_0 \in B$ and $T(0) = t_0 \geq t^*$. All probabilities and expectations in this proof are defined with respect to the probability measure induced by $(Y(0), T(0)) = (x_0, t_0)$.

Define $S_0 = 0$ and $S_{i+1} = \inf\{j > S_i : Y(j) \in B\} \wedge W$ for $i \geq 0$. Notice that S_i is the iteration at which the process returns to B for the i^{th} time (until eventually escaping the cup). Also, let $N^* = \inf\{i : Y(S_i) \in F\}$ so that N^* is the number of times that Y visits B before escaping the cup. Let $\mathcal{G}_i$ denote the σ-algebra generated by $(Y(j) : 0 \leq j \leq S_i, T(0))$.

From Lemma 12, we see that on the event $\{Y(S_n) \in B\}$,

$$\begin{aligned} P\{N^* = n+1 \,|\, \mathcal{G}_n\} &\leq \eta(\Delta_1(T(S_n)))(2c_7)^{r-1}\eta((f_{r+1} - f_2 + \lambda_{T(S_n+1)}(p_{r+1} - p_2))) \\ &\leq (2c_7)^{r-1}\eta(d_{T(S_n)}). \end{aligned}$$

Since the event $\{N^* > n\}$ can be determined from the σ-algebra $\mathcal{G}_n$, we have that

$$\begin{aligned} P\{N^* > n+1 \,|\, \mathcal{G}_n\} &= (1 - P\{N^* = n+1 \,|\, \mathcal{G}_n\})\, I\{N^* > n\} \\ &\geq \left(1 - (2c_7)^{r-1}\eta(d_{T(S_n)})\right) I\{N^* > n\} \\ &\geq \left(1 - (2c_7)^{r-1}\eta(d_{t_0+n})\right) I\{N^* > n\}, \end{aligned}$$

where the last inequality holds since the bottom of any cup is composed entirely of "real" states implying that $T(S_n) \geq T(0) + n$. Taking expectations,

$$P\{N^* > n+1\} \geq \left(1 - (2c_7)^{r-1}\eta(d_{t_0+n})\right) P\{N^* > n.\} \tag{2.12}$$

Applying (2.12) recursively (and assuming that $\eta(d_{t_0}) < (2c_7)^{-(r-1)}$ so that each term is positive), we get

$$P\{N^* > n+1\} \geq \prod_{l=0}^{n} \left(1 - (2c_7)^{r-1}\eta(d_{t_0+l})\right).$$

Suppose that $\bar{\xi}$ is defined sufficiently small so that $\eta(d_{t_0}) \leq \bar{\xi}$ implies $(2c_7)^{r-1}\eta(d_{t_0}) \leq x$ where x is the minimal positive solution to $1 - x = \exp(-2x)$. For such $\eta(d_{t_0})$, $1 - (2c_7)^{r-1}\eta(d_{t_0}) \geq \exp\left(-2(2c_7)^{r-1}\eta(d_{t_0})\right)$. With this assumption on $\eta(d_{t_0})$ and the observation that $\{N^* > q\} \subset \{Y(k) \in C \text{ for } 0 \leq k \leq q\}$, we obtain

$$
\begin{aligned}
P\{Y(k) \in C \text{ for } 0 \le k \le q\} &\ge P\{N^* > q\} \\
&\ge \prod_{k=0}^{q-1} \left(1 - (2c_7)^{r-1}\eta(d_{t_0+k})\right) \\
&\ge \prod_{k=0}^{q-1} \exp\left(-2(2c_7)^{r-1}\eta(d_{t_0+k})\right) \\
&= \exp\left(-2(2c_7)^{r-1}\sum_{k=0}^{q-1}\eta(d_{t_0+k})\right) \\
&\ge \exp\left(-2(2c_7)^{r-1}\sum_{k=0}^{q-1}\eta(d_{T(k)})\right).
\end{aligned}
$$

By setting $c_6 \ge 2(2c_7)^{r-1}$ and $\bar{\xi} \le \min\{(2c_7)^{-\frac{1}{\Delta(t^*)}}, (.79681213)(2c_7)^{-(r-1)}\}$, Theorem 3 holds for an arbitrary cup C. Repeating for all other cups, we can choose $\bar{\xi}$ and c_6 to depend only on Θ and obtain the desired result. □

2.3.7 Settling into Deep Cups

In this section, we present a final theorem that helps complete the proof of Theorem 1. For fixed (S, P), define $E_{S,P}$ as the set of local minimum with depth pairs greater than (S, P). Formally stated, $E_{S,P} = \{x \mid x$ is a local minimum such that $(S(x), P(x)) \succ (S, P)\ \}$. Additionally, define $O_{S,P} = \{y \mid y$ is accessible from some $x \in E_{S,P}$ at level $(f' = S + f(x), p' = P + p(x))$.

Theorem 4. *Let (S, P) be given, and consider a cup C with depth pair (S, P), so that $d_{T(k)} = S + \lambda_{T(k)}P$. If $\sum_{k=0}^{\infty} \eta(S + \lambda_{T(k)}P) = +\infty$, then $\lim_{k\to\infty} P\{Y(k) \in O_{S,P}\} = 1$.*

To prove Theorem 4, we partition the state space $\mathcal{S}$ into the sets $\{U, C_1, \ldots, C_l\}$ so that $U = \{z \in \mathcal{S} : E_{S,P}$ is accessible at $(f(z), p(z))$ from $z\}$, and $\{C_1, \ldots, C_l\}$ are

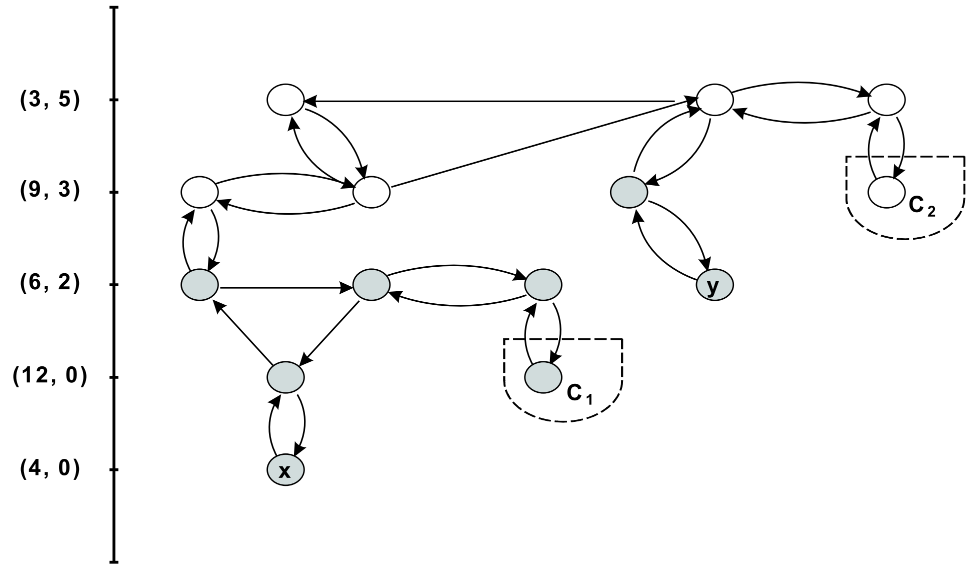

Figure 2.6: Let $(S, P) = (-4, 3)$. This diagram illustrates the partition of $\mathcal{S}$ into U and a set of cups, $\{C_1, \ldots, C_l\}$, such that $(S(C_i), P(C_i)) \prec (-4, 3)$ for $i = 1, \ldots, l$. Cups are enclosed with dashed lines, and all other states are in U. The shaded nodes represent states in the set $O_{(-4,3)}$, and x and y compose $E_{(-4,3)}$.

cups with $(S(C_i), P(C_i)) \prec (S, P)$ for all $i = 1, \ldots, l$. Figure 2.6 illustrates our new perspective of the solution space.

We will invoke the results of Theorems 2 and 3 to prove Theorem 4. Set ξ equal to the minimum of $\bar{\xi}$ as in Theorem 3 and $\min_{i=1,\ldots,l}\{\xi_i\}$ as in Theorem 2, where l indexes the (finite number of) cups.

Let $A_0 = 0$ and $A_{i+1} = \inf\{k > A_i : Y(k) \in U\}$ for $i \geq 0$. Then A_i denotes the iteration when the process is in U for the i^{th} time. Further, define $\alpha = \inf\{i : Y(A_i) \in E_{S,P}\}$ and $\beta = \inf\{i \geq \alpha : Y(A_i) \notin O_{S,P}\}$. Note that α is the number of times that $Y(k)$ visits U by the time it visits the set $E_{S,P}$ for the first time. The related term β is the number of times that $Y(k)$ visits U between first visiting $E_{S,P}$

and climbing to a state $z \notin O_{S,P}$.

We now state three lemmas that are proved using the results of Theorem 2. In Lemma 13, we bound the expected accumulated jump probabilities until escaping a cup C_i. Lemma 14 shows that the expected number of times that Y hits the set U before hitting the bottom of a cup with a depth pair greater than (S, P) is finite. The results of Lemmas 13 and 14 are then combined to bound the expected accumulated jump probabilities until hitting a state in $E_{S,P}$.

Unless noted otherwise, we define all probabilities and expectations with respect to the probability measure induced when $(Y(0), T(0)) = (x_0, t_0)$ for $x_0 \in \mathcal{S}$ and $t_0 \geq t^*$.

Lemma 13. *Under the conditions of Theorem 2, there exists $c_8 > 0$ depending only on Θ such that if $T(A_i + 1) \geq t_0$ and $\eta(r_{t_0}) \leq \xi$, then*

$$E\left[\sum_{k=A_i+1}^{A_{i+1}} \eta(S + \lambda_{T(k)} P) \mid Y(A_i + 1), T(A_i + 1)\right] \leq c_8.$$

Proof. There are two cases to consider, $Y(A_i+1) = Y(A_{i+1}) \in U$ or $Y(A_i+1) \notin U$. The result follows directly from Theorem 2(a) and the fact that Y cannot jump directly from one of the cups to another without visiting a state in U.

Lemma 14. *Under the conditions of Theorem 2, there exists a constant $c_9 > 0$ depending only on Θ such that $E[\alpha] \leq c_9$.*

Proof. For each $x \in U$, there exists $n = n(x)$, and a sequence of distinct states, $\{x = s_0, \ldots, s_n = z\}$, such that $s_j \in U$ for $0 \leq j \leq n$, $(f(s_0), p(s_0)) \succeq (f(s_1), p(s_1)) \succeq \cdots (f(s_n), p(s_n))$, and $s_n \in E_{S,P}$. Along this path, there are two ways that the process can move from s_j to s_{j+1}, either (i) Y moves from s_j to s_{j+1} by direct transition or (ii) Y descends into a cup and climbs out via s_{j+1}. By Theorem 2(b), we see that the probabilities of such (potentially multistep) transitions are bounded below. Therefore, a standard geometric trials argument shows that α has a geometrically decaying tail and so $E[\alpha] \leq c_9$.

Lemma 15. *Under the conditions of Theorem 2, there is a constant c_{10} depending only on Θ such that for every $x_0 \in \mathcal{S}$, and every $t_0 \geq t^*$ with $\eta(r_{t_0}) \leq \xi$,*

$$E\left[\sum_{k=0}^{A_\alpha} \eta(d_{T(k)})\right] \leq c_{10}.$$

Proof. The proof is similar to that of Lemma 5, and uses Lemma 13 and Lemma 14.

We need one more preparatory lemma before proving the main result of this section.

Lemma 16. *For $T(0) = t_0$, $\sum_{k=0}^{\infty} \eta(d_{T(k)}) = \infty$ if and only if $\sum_{k=0}^{\infty} \eta(d_{t_0+k}) = \infty$.*

Proof. First suppose $\sum_{k=0}^{\infty} \eta(d_{T(k)}) = \infty$. Note that the most consecutive artificial states that one may visit is $\ell - 1$ since any state in (f_1, p_1) and (f_ℓ, p_ℓ) must be real. Therefore, the clock will be paused for at most $\ell - 1$ consecutive iterations at a time, and so

$$(\ell - 1)\left[\sum_{k=0}^{\infty} \eta(d_{t_0+k})\right] \geq \sum_{k=0}^{\infty} \eta(d_{T(k)}) = \infty.$$

The converse is immediate since $T(k) \leq t_0 + k$, and so $\eta_{T(k)} \geq \eta_{t_0+k}$.

Proof of Theorem 4. Unless specified otherwise, all probabilities and expectations in this proof are defined with respect to the probability measure induced when

$(Y(0), T(0)) = (x_0, t_0)$ for $x_0 \in \mathcal{S}$ and $t_0 \geq t^*$. Notice that

$$\begin{aligned} P\{A_\alpha > q\} \sum_{l=0}^{q} \eta(d_{t_0+l}) &= E\left[I\{q < A_\alpha\} \sum_{l=0}^{q} \eta(d_{T(0)+l}) \right] \\ &\leq E\left[\sum_{l=0}^{A_\alpha} \eta(d_{T(0)+l}) \right] \\ &\leq E\left[\sum_{l=0}^{A_\alpha} \eta(d_{T(l)}) \right] \\ &\leq c_{10}. \end{aligned}$$

Thus,

$$P\{A_\alpha \leq q\} \geq 1 - \frac{c_{10}}{\sum_{l=0}^{q} \eta(d_{t_0+l})}.$$

Recognize that $Y(A_\alpha)$ must be in the bottom of a cup $\bar{C}$ with depth pair $(S(\bar{C}), P(\bar{C})) \succeq (S + \gamma_S, P + \gamma_P)$ where $(\gamma_S, \gamma_P) = \min\{(S(C') - S, P(C') - P) : C'$ is a cup with depth pair $(S(C'), P(C')) \succ (S, P)\}$.

We can relate the dynamic depth of $\bar{C}$ at iteration k to (S, P) by

$$\begin{aligned} \bar{d}_{T(k)} &= S(\bar{C}) + \lambda_{T(k)} P(\bar{C}) \\ &\geq S + \lambda_{T(k)} P + \gamma_S + \lambda_{T(k)} \gamma_P \\ &= d_{T(k)} + \gamma_S + \lambda_{T(k)} \gamma_P. \end{aligned}$$

Conditioning on A_α and applying Theorem 3, we find that

$$\begin{aligned} P\{A_\beta > q\} &\geq \exp\left(-c_6 \sum_{k=0}^{q-1} \eta(\bar{d}_{T(k)}) \right) \\ &= \exp\left(-c_6 c_{11} \sum_{k=0}^{q-1} \eta(\bar{d}_{t_0+k}) \right) \\ &\geq \exp\left(-c_6 c_{11} \eta(\gamma_S + \lambda_{t_0} \gamma_P) \sum_{k=0}^{q} \eta(d_{t_0+k}) \right), \end{aligned}$$

for some constant $c_{11} \geq 1$.

Since $Y(q) \in O_{S,P}$ if $A_\alpha \leq q < A_\beta$, we have

$$\begin{aligned} P\{Y(q) \in O_{S,P}\} &\geq P\{A_\alpha \leq q < A_\beta\} \\ &\geq \exp\left(-c_6 c_{11} \eta(\gamma_S + \lambda_{t_0}\gamma_P) \sum_{k=0}^{q} \eta(d_{t_0+k})\right) - \frac{c_{10}}{\sum_{k=0}^{q} \eta(d_{t_0+k})} \end{aligned}$$

where we have used the inequality $P\{A \cap B\} \geq P\{A\} + P\{B\} - 1$.

We now show that this inequality gives $\lim_{q\to\infty} P\{Y(q) \in O_{S,P}\} = 1$. Let $\epsilon > 0$ be arbitrary except $\frac{c_{10}}{\epsilon} > 1$. Now choose t_3 such that if $t_0 > t_3$ then

$$\exp\left(-c_6 c_{11} \eta(\gamma_S + \lambda_{t_0}\gamma_P)\left(\frac{3c_{10}}{\epsilon}\right)\right) \geq 1 - \frac{\epsilon}{2}.$$

Choose $q^*(t_3)$ such that

$$\sum_{k=0}^{q^*(t_3)} \eta(d_{t_3+k}) \geq \frac{2c_{10}}{\epsilon},$$

which is possible by Lemma 16. Now, for all $q > q^*(t_3)$, there exists $t_1(q), t_2(q)$ such that (i) $t_3 \leq t_1(q) < t_2(q)$, (ii) $\frac{2c_{10}}{\epsilon} \leq \sum_{k=0}^{q} \eta(d_{t_0+k}) \leq \frac{3c_{10}}{\epsilon}$ for all $t_0 \in [t_1(q), t_2(q)]$, and (iii) $t_2(q) - t_1(q) > 1$. Then, for $t_0 \in [t_1(q), t_2(q)]$,

$$\begin{aligned} P\{Y(q) \in O_{S,P}\} &\geq \exp\left(-c_6 c_{11} \eta(\gamma_S + \lambda_{t_0}\gamma_P) \sum_{k=0}^{q} \eta(d_{t_0+k})\right) - \frac{c_{10}}{\sum_{k=0}^{q} \eta(d_{t_0+k})} \\ &\geq \exp\left(-c_6 c_{11} \eta(\gamma_S + \lambda_{t_0}\gamma_P)\left(\frac{3c_{10}}{\epsilon}\right)\right) - \frac{c_{10}}{\sum_{k=0}^{q} \eta(d_{t_0+k})} \\ &\geq \left(1 - \frac{\epsilon}{2}\right) - \frac{\epsilon}{2} \\ &= 1 - \epsilon. \end{aligned}$$

This probability bound is conditioned on $(Y(0), T(0)) = (x_0, t_0)$ where $t_0 \in [t_1(q), t_2(q)]$. So when $T(0) = 0$, let $\rho = \min\{k : T(k) \geq t_1(q)\}$. Since $t_2(q) -$

$t_1(q) > 1$, $T(\rho) < t_2(q)$. Now all probabilities and expectations are induced by $(Y(0), T(0)) = (x_0, 0)$. We find that

$$
\begin{aligned}
P\{Y(q) \in O_{S,P}\} &= E[P\{Y(q) \in O_{S,P} \mid (Y(\rho), T(\rho))\}] \\
&\geq E[(1-\epsilon)] \\
&= 1-\epsilon.
\end{aligned}
$$

Since ϵ was arbitrary, the result follows. □

2.3.8 Proof of Main Result

In this section, we conclude the proof of Theorem 1 by assembling results from Theorems 3 and 4. With the help of an additional lemma, we modify the final argument in [56] to accommodate the concept of levels.

Lemma 17. *If* $\sum_{k=0}^{\infty} \eta(S + \lambda_{T(k)} P) < +\infty$ *and* $\sum_{k=0}^{\infty} \eta(\bar{S} + \lambda_{T(k)} \bar{P}) = +\infty$, *then* $(S, P) \succ (\bar{S}, \bar{P})$.

Proof. We prove by contradiction by assuming $(S, P) \preceq (\bar{S}, \bar{P})$. Then, either (i) $P < \bar{P}$ and nothing can be said about the relative sizes of S and $\bar{S}$, or (ii) $P = \bar{P}$ and $S \leq \bar{S}$. In either case, by our construction of a compression schedule, there exists $k^* \geq 0$ such that $\eta(S + \lambda_{T(k)} P) \geq \eta(\bar{S} + \lambda_{T(k)} \bar{P})$ for $k \geq k^*$. Therefore,

$$
\begin{aligned}
\sum_{k=0}^{\infty} \eta(S + \lambda_{T(k)} P) &\geq \sum_{k=k^*}^{\infty} \eta(S + \lambda_{T(k)} P) \\
&\geq \sum_{k=k^*}^{\infty} \eta(\bar{S} + \lambda_{T(k)} \bar{P}) \\
&= \infty
\end{aligned}
$$

which is a contradiction.

Proof of Theorem 1. When $S = 0$ and $P = 0$, $O_{S,P}$ is the set of all local minima and the condition of Theorem 4 holds. Thus Theorem 4 implies part (a) of Theorem 1.

We will next prove the "if" half of part(b) of Theorem 1. Let $B_{S,P}$ denote the bottom of a cup C with depth pair (S, P). Suppose that the states in $B_{S,P}$ are local minima with static depth S and penalty depth P, i.e., C is the largest cup containing the states in $B_{S,P}$ as local minima. We must show that $P\{Y(k) \in B_{S,P}\}$ has limit zero as k tends to infinity if $\sum_{k=0}^{\infty} \eta(S + \lambda_{T(k)}P) = \sum_{k}^{\infty} \eta(d_{T(k)}) = \infty$. In view of Theorem 4, it is sufficient to prove that $B_{S,P} \cap O_{S,P} = \emptyset$. For the sake of contradiction, suppose there exists $x \in B_{S,P} \cap O_{S,P}$. Then there is a state $y \in E_{S,P}$ so that (i) y is a local minimum with $(S(y), P(y)) \succ (S, P)$, (ii) x is accessible from y at $(S + f(y), P + p(y))$, and (iii) y is accessible from x at $(S + f(y), P + p(y))$.

<u>Case 1</u>: $(f(x), p(x)) \succ (f(y), p(y))$

Since $x \in B_{S,P}$, $(f' = S + f(x), p' = P + p(x))$ is the lowest level such that a state z, with $(f(z), p(z)) \prec (f(x), p(x))$, is accessible from x. However, by (iii), we know that y is accessible from x at $(S + f(y), P + p(y))$. This contradicts the fact that x is a local minimum with $(S(x), P(x)) = (S, P)$, since $(f(y), p(y)) \prec (f(x), p(x))$, and so $(S + f(y), P + p(y)) \prec (S + f(x), P + p(x))$.

<u>Case 2</u>: $(f(x), p(x)) \preceq (f(y), p(y))$

Since $x \in B_{S,P}$, there exists a state z, with $(f(z), p(z)) \prec (f(x), p(x))$, accessible from x at $(S + f(x), P + p(x))$. Since $(f(x), p(x)) \preceq (f(y), p(y))$, z is also accessible from x at $(S + f(y), P + p(y))$. Combining this observation with (ii), note that z is accessible from y at $(S+f(y), P+p(y))$. Furthermore, (i) implies that $(f'', p'') \succ (S+f(y), P+p(y))$ is the lowest level such that a state w, with $(f(w), p(w)) \prec (f(y), p(y))$, is accessible from y. Therefore, we have reached a contradiction of (i) because z is accessible from y at $(S+f(y), P+p(y))$ and $(f(z), p(z)) \prec (f(x), p(x)) \preceq (f(y), p(y))$.

We obtain a contradiction in both cases, so we have proved that $B_{S,P}$ and $O_{S,P}$ are disjoint. This completes the proof of the "if" half of part (b) of Theorem 1.

To prove the "only if" half of part (b) of Theorem 1, we will prove the contrapositive. Again let $B_{S,P}$ denote the bottom of a cup C with depth pair (S, P), and suppose that the states in $B_{S,P}$ are local minima of static depth S and penalty depth P. Assume that η_t and λ_t are given such that $\sum_{k=0}^{\infty} \eta(S + \lambda_{T(k)} P) < +\infty$. We want to prove that $P\{Y(k) \in B_{S,P}\}$ does not converge to zero as k tends to infinity.

Since $\sum_{k=0}^{\infty} \eta(S + \lambda_{T(k)} P) < +\infty$, we know from Lemma 17 that $(S, P) \succ (\bar{S}, \bar{P})$, where $(\bar{S}, \bar{P})$ are the largest value pair such that $\sum_{k=0}^{\infty} \eta(\bar{S} + \lambda_{T(k)} \bar{P}) = +\infty$. Select a state $y \in B_{S,P}$ and let $\bar{C} = \{x : x \text{ is accessible from } y \text{ at level } (\bar{S} + f(y), \bar{P} + p(y))\}$. Let $\bar{B}$ denote the bottom of cup $\bar{C}$, and let $x_0 \in \bar{B}$.

Observing that $\sum_{k=0}^{\infty} \eta(d_{T(k)}(\bar{C})) < +\infty$ since $(S(\bar{C}), P(\bar{C})) \succ (\bar{S}, \bar{P})$, Theorem 3 implies that

$$\liminf_{k \to \infty} P\{Y(k) \in \bar{C}\} > 0. \tag{2.13}$$

If $z \in \bar{C} - \bar{B}$, then z is not a local minimum or z is a local minimum with depth pair $(S(z), P(z)) \prec (\bar{S}, \bar{P})$. Using this observation, part (a) of Theorem 1, and the "if" half of part (b) of Theorem 1,

$$\lim_{k \to \infty} P\{Y(k) = z\} = 0 \text{ for } z \in \bar{C} - \bar{B}. \tag{2.14}$$

From (2.13) and (2.14), we conclude that $\liminf_{k\to\infty} P\{Y(k) \in \bar{B}\} > 0$, and since $\bar{B} \subset B$, this inequality is true with $\bar{B}$ replaced by B. This completes the “only if” half of part (b) of Theorem 1. □

2.3.9 Summary and Future Theoretical Study

We have developed a necessary and sufficient condition for the convergence in probability of simulated annealing in the presence of complicating constraints. Through careful construction of the concept of levels, we are able to maintain the monotonicity of the long-run transition probabilities and therefore modify the mechanics in [56] to achieve the desired result. We have rigorously established that, to converge to the set of global minima, slower cooling than proposed in [56] is required. Therefore, it is unlikely that we can view compressed annealing as a global optimization method, but rather as a heuristic implemented with accelerated cooling and compression schedules.

Compressed annealing augments the traditional simulated annealing algorithm by implementing an approach that varies the values of penalty multipliers throughout the execution of the algorithm. In this manner, compressed annealing can be viewed as a primal-dual variant of simulated annealing. We assume that the values of each penalty multiplier are the same, masking the duality of the problem. Future research is necessary to explore the effect of having distinct values for each penalty multiplier.

The results presented in this chapter shed light on the performance of the algorithm, but more work is needed. One future endeavor is to create a class of problems for which we can derive analytical results on probability of convergence and rate of convergence in order to garner further insight on the algorithm's behavior. Future research may include the analysis of annealing behavior during the transient period of compression ($\lambda < \lambda^*$) and its impact on the quality of heuristic solutions in practical applications.

2.4 Implementation of Compressed Annealing

To implement the compressed annealing algorithm on a particular application, the practitioner is faced with a number of decisions. Some of these decisions are application-specific, and must be addressed on a problem-by-problem basis. Other issues, however, are generic in the sense that there exist general procedures that guide their resolution. In this section, we describe this generic parameter determination. We reserve discussion for examples of the application-specific decisions until Chapters III and IV.

We provide a general procedural description of compressed annealing in Table 2.1. General issues that must be addressed include the setting of initial temperature and pressure, the manner and rate of changing temperature and pressure, the number of iterations at each temperature/pressure setting, and the algorithm termination criterion.

Cooling and Compression

From our theoretical work on the convergence of compressed annealing, we discover that joint cooling and compression schedules should have decreasing derivatives to ensure convergence to the set of global minima. While the theoretical rates of

Table 2.1: Outline of Compressed Annealing

Generate initial solution, x.
Initialize best solution found, x_{best}, so that $f(x_{best}) = \infty$ and $p(x_{best}) = 0$
Let $k = 0$.
Set initial temperature and pressure, τ_k and λ_k.
Set *iteration limit*, i.e., number of iterations at each temperature/pressure.
Repeat:
Let *counter* $= 0$.
Repeat:
Increment *counter* by 1.
Randomly generate y, a neighbor solution of x.
With probability $\exp\left(\frac{-(v(y,\lambda_k)-v(x,\lambda_k))^+}{\tau_k}\right)$, let $x = y$.
If $p(x) = 0$ and $f(x) < f(x_{best})$, let $x_{best} = x$.
Until *counter* == *iteration limit*.
Increment k by 1.
Compute τ_k and λ_k.
Until termination criterion satisfied.

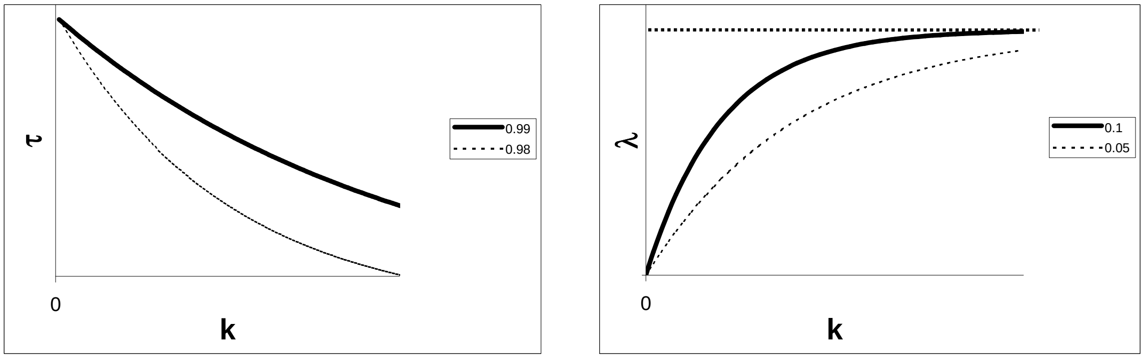

Figure 2.7: Demonstration of practical cooling and compression schedules.

cooling and compression are much too slow to be practical, they supply insight on appropriate ‘shapes” for simultaneous cooling and compression schedules (see Figure 2.7). Combining this intuition with observations documented in the literature, we implement a geometric cooling schedule [35, 81, 134] and a limited exponential compression schedule. For parameters $0 \leq \beta \leq 1$, $\gamma \geq 0$, and $\hat{\lambda} \geq 0$ (an approximation of the pressure cap), these schedules are formally defined by

$$\begin{aligned}\tau_{k+1} &= \beta\tau_k, \text{ and} \\ \lambda_{k+1} &= \hat{\lambda}\left(1 - \frac{(\hat{\lambda} - \lambda_0)}{\hat{\lambda}} e^{-\gamma k}\right).\end{aligned}$$

Values of the cooling parameter, β, typically range from 0.80 to 0.99 and values of the compression parameter, γ, usually vary from 0.01 to 0.1.

To apply these schedules, initial values of temperature and pressure still need to be determined, as well as an approximation for the pressure cap. The limited exponential form of compression allows the convenience of simply setting $\lambda_0 = 0$, but more care must be taken in setting τ_0 and $\hat{\lambda}$. The initial value of temperature,

τ_0, must be selected so that early in the algorithm, the probability of accepting uphill transitions is close to 1; this allows the algorithm sufficient mobility to search the solution space. However, setting the temperature prohibitively high results in long computation times or poor convergence. Setting the initial temperature takes on increased importance in the presence of pressure, as setting τ_0 excessively high wastes the benefit of searching a "relaxed" topography in the sense that the search is random rather than guided by a tendency to go downhill.

As an initial step in parameter initialization, we generate $\mathcal{R}$, a sample of $2n$ solutions, by randomly generating n pairs of neighbor solutions. We specify an appropriate initial value of temperature by adapting techniques from Laarhoven and Aarts [134] and Dowsland [35] to utilize information provided by $\mathcal{R}$. First, we specify χ_0, the percentage of proposed uphill transitions that we require to be accepted at τ_0. Computing $\overline{|\Delta v|}$, the average absolute difference in objective function over the n sample transitions composing $\mathcal{R}$, we determine the initial temperature as

$$\tau_0 = \frac{\overline{|\Delta v|}}{\ln\left(\frac{1}{\chi_0}\right)}.$$

At this value of initial temperature, the actual acceptance ratio over the first loop of iterations of compressed annealing is monitored. If the actual acceptance ratio is less than χ_0, then τ_0 is reset at 1.5 times its current value and re-evaluated over a loop of iterations. This procedure is continued until the observed acceptance ratio for a loop of iterations equals or exceeds χ_0.

As shown by the theoretical analysis in §2.2, there exists a pressure cap, λ^*, beyond which further compression serves only to exaggerate the solution topography's features. Therefore, an ideal practical compression schedule would gradually increase λ from an initial value of zero to λ^*, allowing the algorithm to explore the solution space via solutions infeasible in terms of the relaxed constraints. Unfortunately, determining a tight upper bound on λ^* using only the limited information from the sample $\mathcal{R}$ is difficult. Nonetheless, we present an approach that, while not guaranteeing an upper bound on λ^*, can be experimentally calibrated to determine an approximation of the pressure cap.

To approximate the pressure cap, we introduce an additional parameter, $0 \leq \kappa \leq 1$ to determine our estimate, $\hat{\lambda}$. The value of κ represents the percentage of the objective function value that is composed of the penalty term when $\lambda = \hat{\lambda}$. Our pressure cap approximation is given by

$$\hat{\lambda} = \max_{x \in \mathcal{R}} \left\{ \frac{f(x)}{p(x)} \frac{\kappa}{1-\kappa} \right\},$$

where values of κ ranging from 0.75 to $0.\overline{99}$ have demonstrated computational promise.

Iterations Per Temperature/Pressure Setting

Theoretical research on simulated annealing suggests that the system should be allowed to converge to its stationary distribution at each temperature setting. Unfortunately, the number of iterations necessary to approach the stationary distribution is exponential in problem size [35]. In practice, the length of the Markov chain at each temperature is usually related to the size of the neighborhood structure or even the solution space. Bonomi and Lutton [15] set the number of iterations at the same temperature to a value depending polynomially on the size of the problem. An alternate approach determines the length of the k^{th} Markov chain by not allowing a temperature reduction until a minimum number of transitions have been accepted or a maximum number of iterations has been eclipsed. In this manner, Kirkpatrick et al. [81] let the length of the k^{th} Markov chain be dependent on k.

In the applications that we consider in Chapters III and IV, we use the size of the problem as a general rule-of-thumb in determining the number of iterations per temperature. We fine-tune this parameter through experimental testing.

Termination Criterion

There have been numerous stopping conditions reported in the literature. Bonomi and Lutton [15] fix the number of temperature values for which the algorithm is executed. Johnson et al. [73] terminate the algorithm when the percentage of accepted moves drops below a threshold for a number of iterations. In the applications that we consider in Chapters III and IV, we implement a hybrid of these two approaches by monitoring the mobility of the algorithm while also requiring a minimum number of iterations. We fine-tune this approach through experimental experience.

CHAPTER III

Stochastic Fleet Replacement with Budget Constraints

In the fiercely competitive trucking industry, a firm gains a strategic advantage by by implementing a replacement policy that minimizes its fleet management costs while providing necessary capacity. Fleet managers must balance the operating and maintenance costs of their current fleet of defenders against purchasing costs of potential challengers when determining a replacement strategy for a fleet of vehicles. Complicating this decision-making process is the stochastic deterioration of a vehicle's physical condition and changes in purchasing, operating, and maintenance costs over time due to the advent of new technologies. Additionally, controls on capital expenditures may constrain replacement decisions. Budget reductions for capital expenditures generally increase fleet maintenance expenses as assets are forced to remain in service for longer periods. Thus, there is heightened emphasis on evaluating the fleet over the planning horizon and timing replacement decisions to minimize overall costs.

Feedback from various fleet managers validates these issues. One fleet manager quipped, “Trucks are a lot like children. Across a fleet, trucks’ maintenance behavior varies as broadly as the temperaments of children on a playground.” Another fleet manager noted that there are “quantum leaps” in the maintenance condition of trucks as they age; a truck may be in near-perfect mechanical condition in one year, but poor condition the very next. As described by a fleet manager, “The key is to replace these vehicles before the persistent maintenance problems hit.” Fleet managers also concur that constricting budgets often impede replacement plans [37]. Management of various companies expresses interest in improving or validating the current fleet replacement policies, which are dictated by finance and guided by intuition.

By appropriately modifying the mathematical formulation in Morse [98], we study fleet replacement under the effect of stochastic deterioration, annual budget limits for capital expenditures, and time-variant costs induced by technological change. We quantify the interactions of these three real-world issues by utilizing data provided by a trucking company. In the replacement literature, there is separate treatment of stochastic deterioration [33], budget constraints [78, 79, 58], and even joint stochastic deterioration and technological breakthrough [71]. We extend the analysis to simultaneously address the interaction effects of these three conditions. To solve the complicated model, we employ compressed annealing and evaluate its effectiveness with respect to a lower bound and common industry practice.

The rest of this chapter is organized as follows. In §3.1, we provide a brief overview of the typical replacement strategies employed by trucking companies. We present a mathematical formulation addressing prominent replacement issues facing fleet managers in §3.2. We study the applicability of this model by analyzing actual truck maintenance logs in §3.3. In §3.4, we discuss the implementation of compressed annealing and analyze computational results.

3.1 Replacement Strategies in Practice

We begin with a survey of current fleet management practices. We visited a pair of less-than-truckload (LTL) trucking companies and gathered information regarding their fleets. LTL fleets are composed both of linehaul and pick-up and delivery tractors. Pick-up and delivery is a local operation that is usually confined to an area the size of a city, or in some cases, to regions of a city [132]. All newly picked-up loads are brought to a central terminal where loads are consolidated for the linehaul operation. Linehaul trucking is the high mileage operation (over 60,000 miles/year) of moving freight to locations around the country. Generally, there are more than 30 miles between stopping and starting in linehaul operations.

Our study focuses solely on the fleet of linehaul tractors, which we will treat as independent from the fleet of pickup and delivery tractors. This is a simplification, since linehaul tractors commonly transition into pickup and delivery service upon retirement from linehaul service. However, by examining a fleet of vehicles, homogeneous in service type, we simplify our analysis and hope to build insight for future research that may include the pickup and delivery operation.

Table 3.1 illustrates a stark contrast in the replacement policies for the two trucking companies we visited. Company A considers only one challenger, new tractors of a single brand, when deciding to replace the defender. Company B considers many

Table 3.1: Linehaul Fleet Management Summary

	Company A	Company B
replacement options	new (single brand)	new, used, leased (multiple brands)
planning horizon	10 years	5 years
life-cycle	10 years	6 years
fleet size	98	164
avg. annual mileage	110,000	120,000

challengers including new, used, and leased tractors across a variety of brands such as International, Sterling, Freightliner, Volvo White, and Mack.

Company B attempts to maintain a uniform age distribution across its fleet so that its age-based replacement policy results in a smooth financing requirement. However, age-based policies may be ineffective in the presence of stochastic deterioration, as age alone is often not an adequate indicator of cost. For example, a scenario could develop in which a x-year old truck with a poor maintenance record should be replaced before a $(x + 1)$-year old truck with an average maintenance record.

Company A adheres to a longer life-cycle based on the premise that its premier preventive maintenance program is the key to lowering costs over the life of the vehicle. For example, an auto-lube system is installed in every tractor. Additionally, a tractor is comprehensively repaired whenever it goes in for major work. For instance, if a tractor is having its brakes replaced and the mechanics notice that the clutch is substantially worn, they will replace the clutch also to avoid an additional maintenance visit. In contrast, there are companies that run extremely short life-cycles (three years) and cut back on preventive maintenance to curb costs.

The lack of uniformity in the replacement policies of firms in the same competitive market is intriguing [136]. The disparate approaches may be due to the varied consideration of factors that affect replacement decisions. Fleet managers list a defender's age, cumulative utilization, fuel economy, and maintenance history as well as driver feedback among the factors considered in replacement decisions. Additional factors include the purchase price, warranty service, and new technology associated with the challenger(s). However, there is no comprehensive framework that quantifies the effect of the various factors on replacement policies. Replacement decisions are often decided by the fleet manager's "gut feeling." By encapsulating the key difficulties facing fleet managers, particularly stochastic deterioration and budget constraints, we provide a decision support framework for fleet asset management. With it, we can help determine effective fleet management strategies.

3.2 Model Description

Our fleet replacement model explicitly considers stochastic deterioration, annual budget limits for capital expenditures, and time-variant costs. Markov decision processes (MDPs) are a common technique to model such sequential decision-making problems under uncertainty. However, if we attempt to model the entire fleet with a single MDP, we suffer from the curse of dimensionality as the state space grows exponentially with the number of vehicles. Alternatively, if we model the fleet as a collection of single vehicle MDPs, we cannot capture the interdependence caused by the budget constraints.

The stochastic fleet replacement problem with budgets (SFRPB) can be accurately modeled as a nonhomogeneous Markov decision process with side constraints [99]. Following the framework of Morse [98], we approach the SFRPB by first modeling each individual vehicle with a nonhomogeneous Markov decision process. Then we transform the dynamic programming formulation for individual vehicles into a mixed integer program that models the decisions for the entire fleet subject to constraints reflecting annual budgets.

3.2.1 Dynamic Programming Formulation

We define a vehicle's *state*, i, as the vector containing all the pertinent information about its behavior. In our formulation, we characterize a vehicle by its age, i_n, and its maintenance condition, i_m. Hence, a vehicle's state is described by $i = (i_n, i_m)$.

At the beginning of each period, the decision-maker evaluates a fleet of vehicles to determine appropriate replacement actions. We consider a simple binary "keep-or-replace" decision, but in general the replacement options can include purchasing used vehicles, leasing etc. In the current model, we do not consider maintenance as a decision variable. The degree of maintenance cost is induced by the vehicle's physical condition. We list the assumptions of our mathematical formulation modeling the replacement problem below.

Assumptions

- Replacement decisions are evaluated based on a vehicle's state, i, at the beginning of each period.

- Capital purchases and sales occur at the beginning of the period and are associated with the beginning state, i, of a vehicle.

- At the end of each period t, the beginning state, i, of a vehicle stochastically transitions into the ending state, j.
- Operating and maintenance costs occur at the end of the period and are associated with the ending state, j, of a vehicle.
- The ending state of a vehicle in period t is equal to the beginning state of a vehicle in period $t + 1$.
- Operating costs are vary with a vehicle's age (i_n), but are independent of its maintenance condition (i_m).
- Maintenance costs are dependent on a vehicle's state, $i = (i_n, i_m)$.
- Technological advance creates constant increase in purchase price of challengers and constant decrease in operating cost of challengers.
- Technological change does not affect transition probabilities or maintenance costs.
- Replacement decisions evaluated across a finite forecast horizon.

Modifying the formulation in Morse [98], we tailor a replacement model to conform to data acquired from linehaul tractor fleets. We have the following parameters:

$\pi(p, a, t, i)$ = Capital expenditure if action a is selected for asset p in state i at the beginning of period t.

$P(p,a,t,i,j)$	$=$	Probability that asset p will transition to state j at the end of period t if action a is selected for asset p in state i at the beginning of period t.
$O(p,t,j)$	$=$	Operating cost for asset p in state j at end of period t .
$M(p,t,j)$	$=$	Maintenance cost for asset p in state j at end of period t .
$Z(p,t,j)$	$=$	$O(p,t,j) + M(p,t,j)$.
$S(p,t,i)$	$=$	Salvage value for asset p in state i at beginning of period t .
$D(p,a,t,i)$	$=$	net capital acquisition cost if action a is selected for asset p in state i at beginning of period t.
	$=$	$\max\{0, \pi(p,a,t,i) - S(p,t,i)\}$.
$A(p)$	$=$	Set of possible actions for asset p.
$S(p)$	$=$	Set of possible states for asset p.
α	$=$	Discount factor per period.
$B(t)$	$=$	Budget for capital expenditures in period t.
H	$=$	length of planning horizon
P	$=$	Number of assets in fleet.

To formally express the formulation, we utilize the following notation:

$v(p,t,i)$	$=$	Minimum expected cost for asset p from actions in periods t through H if asset p is in state i at beginning of period t.

$$y(p,a,t,i) = \begin{cases} 1 & \text{if action } a \text{ is performed} \\ 0 & \text{otherwise} \end{cases}$$

$q(p,t,i)$ = Probability that asset p is in state i at the beginning of period t.

$= \sum_j \sum_a q(p,t-1,j)y(p,a,t-1,j)P(p,a,t-1,j,i).$

Y = Matrix of action choices, $y(p,a,t,i)$, for $p = 1,\ldots,P$, $a \in A(p)$, $t = 0,\ldots,H-1$, and $i \in S(p)$.

$E(t,Y)$ = Fleet capital expenditures in period t given action choices Y.

If the initial state of each asset is perfectly observable, then $q(p,0,\cdot) = e_i$ for some i, where e_i is a vector of zeroes with a one in the i^{th} position. Otherwise, $q(p,0,\cdot)$ is the probability distribution reflecting the fleet manager's uncertainty about asset p's initial state.

The minimum discounted expected cost for a single asset p is determined by solving the recursive relation in (3.1) for $\{v(p,0,i)\}_{i=1,\ldots,S(p)}$.

$$v(p,t,i) = \min_{a\in A(p)} \left[D(p,a,t,i) + \alpha \sum_{j\in S(p)} P(p,a,t,i,j)\left(Z(p,t,j) + v(p,t+1,j)\right)\right]$$
$$\text{for all } i \in S(p),\ t = 0,\ldots H-1 \tag{3.1}$$

$$v(p,H,i) = -S(p,H,i),\ \text{for all } i \in S(p).$$

3.2.2 Fleet Replacement with Budget Constraints

Solving the dynamic program in (3.1) for all $p = 1,\ldots,P$ establishes the optimal replacement policy for a fleet of economically independent assets. However, we consider the presence of budget constraints, which intertwine the replacement decisions regarding the assets. Therefore, solving the P independent dynamic programs may result in a fleet replacement policy that violates one or more budget constraint over the forecast horizon.

In the presence of budget constraints, the economic life-cycles of various assets will be altered to accommodate the financing situation. Morse [98] discusses various approaches for modeling the capital rationing effect of the budgets in the presence of stochastically deteriorating assets. However, beyond time zero, we do not know the maintenance condition of the defenders with certainty. Constraining capital expenditure so that budgets will not be violated in any possible combination of states across the fleet overly constrains the problem so that it may be infeasible or very costly. In this "fat" formulation [77], the probability of some infeasible scenarios may be very small.

We adopt the constrained expected expenditure approach of Morse [98]. In this case, expenditures for any particular scenario in a period may be slightly over or under budget, but on average, they will be within budget. This approach reflects the softness of budgets in future periods as management guidelines over a planning horizon rather than hard constraints.

The set of constraints that force expected expenditure to remain within budget are formulated as:

$$E(t,Y) = \sum_{p=1}^{P} \sum_{i \in S(p)} \sum_{a \in A(p)} q(p,t,i)\pi(p,a,t,i)y(p,a,t,i) \quad \text{for all } t,$$

$$B(t) - E(t,Y) \geq 0 \qquad \text{for all } t.$$

Observe that the budget constraint includes the product $q(p,t,i)y(p,a,t,i)$. By the recursive definition of $q(p,t,i)$, it is a function of $y(p,\cdot,t,i)$'s from earlier periods. Hence, the budget constraints contain products of decision variables.

With the explicit inclusion of budget constraints, the stochastic fleet replacement problem is no longer solvable using the recursive structure of dynamic programming. Borrowing the framework developed by Morse [98] and utilizing the penalty function described by Hadj-Alouane [55] to relax the budget constraints, we construct an equivalent nonlinear mixed integer program (**MIP**) parameterized by penalty multipliers, where $M \gg 0$ and $s > 0$.

MIP minimize $\sum_p \sum_i q(p,0,i)v(p,0,i) + \sum_{t=0}^{H-1} \lambda_t \left[\min(0, B(t) - E(t,Y))\right]^s$

subject to:

$$v(p,t,i) \geq D(p,a,t,i) + \alpha \sum_j P(p,a,t,i,j)\left[Z(p,t,j) + v(p,t+1,j)\right] - M[1 - y(p,a,t,i)]$$

for all $p, a \in A(p), i \in S(p),\ t = 0,\ldots,H-2$;

$$v(p,t,i) \leq D(p,a,t,i) + \alpha \sum_j P(p,a,t,i,j)\left[Z(p,t,j) + v(p,t+1,j)\right] + M[1 - y(p,a,t,i)]$$

for all $p, a \in A(p), i \in S(p),\ t = 0,\ldots,H-2$;

$$v(p,H-1,i) \geq D(p,a,H-1,i) + \alpha \sum_j P(p,a,H-1,i,j)\left[Z(p,H-1,j) - S(p,H,j)\right] - M[1 - y(p,a,H-1,i)]$$

for all $p, a \in A(p), i \in S(p)$;

$$v(p,H-1,i) \leq D(p,a,H-1,i) + \alpha \sum_j P(p,a,H-1,i,j)\left[Z(p,H-1,j) - S(p,H,j)\right] + M[1 - y(p,a,H-1,i)]$$

for all $p, a \in A(p), i \in S(p)$;

$$\sum_{a \in A(p)} y(p,a,t,i) = 1 \text{ for all } p, t, i \in S(p).$$

In the remainder of Chapter III, we study the applicability of this replacement model to the trucking industry. Through an application of compressed annealing, we solve instances of **MIP** from the trucking industry and report results in § 3.4.4.

3.3 Data Analysis

In this section, we analyze a longitudinal data sample provided by a trucking company. We study the aggregation of maintenance costs to form categories of

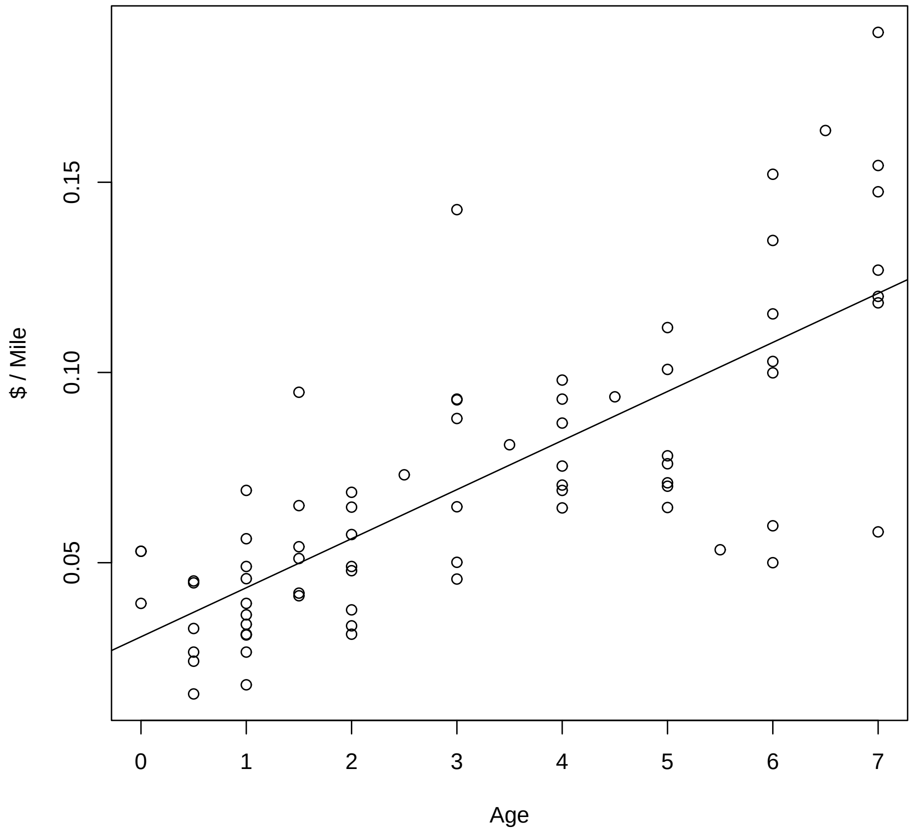

Figure 3.1: Regression of cost per mile on age only.

maintenance condition. To quantify the deterioration behavior, we then generate a set of transition probabilities based on this aggregation scheme.

We use annual maintenance cost per mile (annual maintenance cost / annual mileage) as our cost metric to implicitly account for any variation in vehicle utilization. Since it is commonly assumed in replacement models that maintenance costs vary deterministically with age, we test an initial regression model with age as the only predictor of annual maintenance cost per mile (see Figure 3.1). The resulting fit indicates a weak correlation ($R^2 = 0.5674$), suggesting that age alone is not sufficient to forecast maintenance costs.

Hartman [59] examines an approach that focuses on the effect of utilization (cumulative and annual) on maintenance costs to model stochastic cash flows. To determine whether such a model is appropriate for our fleet management problem, we perform some preliminary statistical analysis. From the plots in Figure 3.2, we see that there is not a strong relationship between total annual maintenance cost and cumulative, annual, or previous year's utilization. Furthermore, a regression model predicting total annual maintenance costs with a combination of age, annual mileage, cumulative mileage, and previous year's mileage results in a poor fit ($R^2 = 0.4775$). Since utilization measures fail to accurately explain the variation in total annual maintenance costs, we pursue a model that predicts annual maintenance cost per mile by classifying the physical condition of the assets.

3.3.1 Aggregation of Maintenance Costs

We observe from the data that annual maintenance cost per mile varies randomly over the life of a tractor. As in Derman [33] and later in Hopp and Nair [71], we model the stochastic cost structure using a finite state space to aggregate the asset maintenance costs.

While the number of maintenance indices to describe a vehicle is flexible, we categorize the annual maintenance cost per mile of operation into three categories: low ($i_m = 1$), medium ($i_m = 2$), and high ($i_m = 3$). This three-tiered classification approach, suggested by fleet managers, is easily implementable and intuitive. In practice, the best number of categories is determined by balancing aggregation error [11] and the accuracy of the transition probability structure. We analyze the

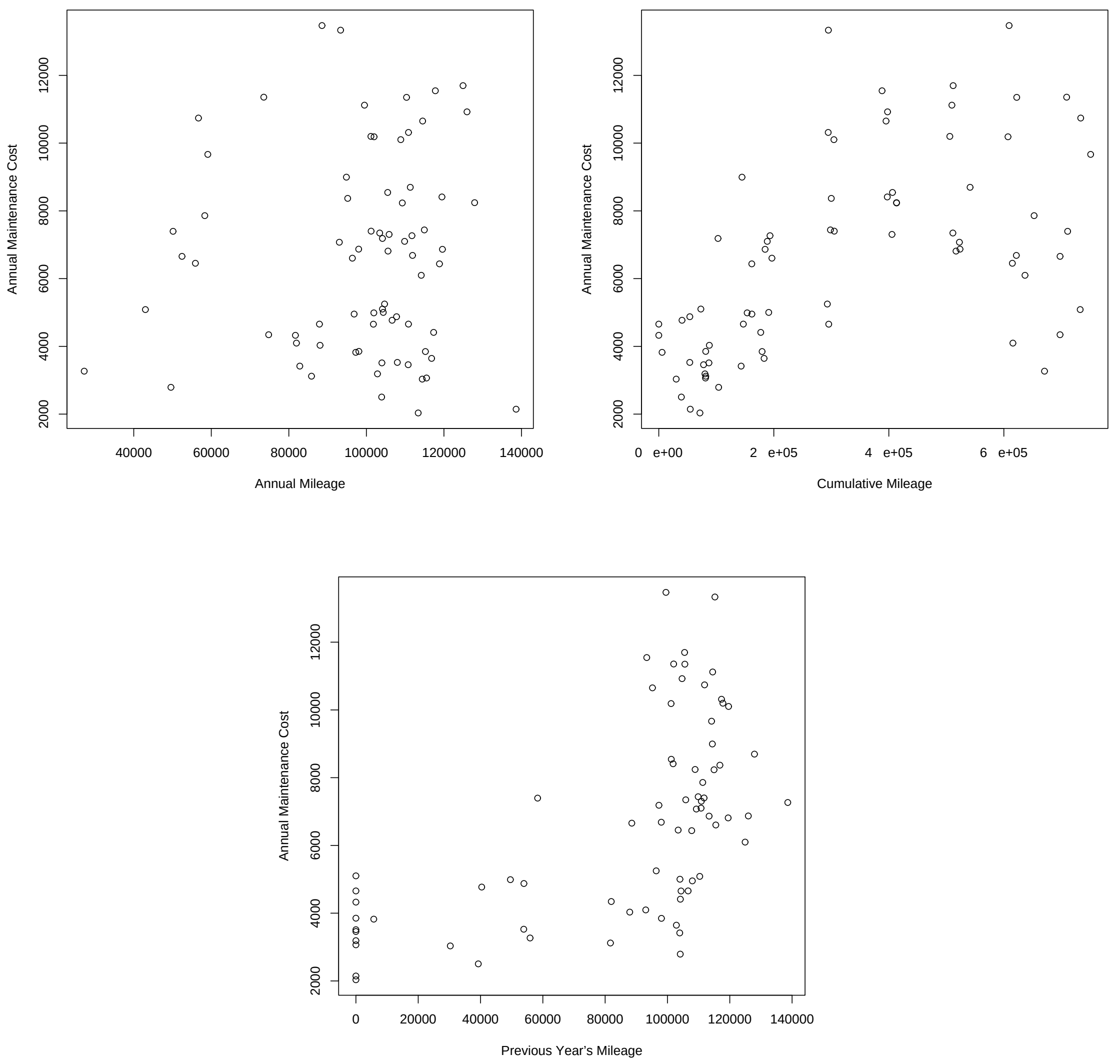

Figure 3.2: Demonstrating the lack of correlation between utilization measures and annual maintenance cost.

Table 3.4: Maintenance Categorization

Category	$/Mile Range
Low	(0.00) 0.01 to 0.06
Medium	0.06 to 0.11
High	0.11 to 0.16 (∞)

aggregation error induced by the three-category approach in Appendix B. Further investigation of the balancing between aggregation error and accuracy of transition probabilities is a possible area of future research in fleet replacement modeling.

In the sample data, an overwhelming majority of the cost per mile ranged from $0.01 to $0.16 per mile. We break this range into three equal intervals for the purposes of aggregation (see Table 3.4). Using this aggregation scheme, Table 3.5 illustrates the evolution of a particular tractor's maintenance category. Over the life of the tractor, not only is there deterioration (transition from low to medium from year 1 to year 2), but also improvement (transition from medium to low from year 5 to year 6). This deterioration and improvement is due to the varying effects of age, driving conditions, maintenance, etc.

When fleet managers make replacement decisions at the beginning of a period, a vehicle's future maintenance costs/category are not known with any certainty. Under the assumption that the evolution of a vehicle's maintenance category is Markovian, we utilize a vehicle's maintenance category at the beginning of the period to determine the expected maintenance category at the end of the period. We demonstrate the validity of the Markov assumption in Appendix A. In the next section, we develop a probability structure that governs the transitional behavior of a vehicle's

Table 3.5: Example of Vehicle-Year Classification

Tractor #62		
Year	$ / Mile	Category
1	0.03	Low
2	0.07	Medium
3	0.05	Low
4	0.09	Medium
5	0.07	Medium
6	0.06	Low
7	0.19	High

maintenance condition over its lifetime.

3.3.2 Development of Transition Probabilities

Through the construction of transition probabilities, we quantify the probabilistic dependence of a vehicle's future maintenance condition on its age and current maintenance condition. Data suggests that transition probabilities depend on a tractor's age. We assume that transition probabilities are not affected by technological change. The impact of technological change on asset deterioration is relatively unknown and could be an area of future research.

We estimate the transition probabilities with smoothed frequency curves (depicted in Figure 3.3) developed by categorizing the data as in Table 3.5. The data and specific functions defining these frequency curves are described in Appendix C.

A tractor's expected annual maintenance cost per mile is calculated according to the appropriate probability distribution, e.g.,

$$E\,[\text{ maintenance \$/mile }] = lp_{ml}(6) + mp_{mm}(6) + hp_{mh}(6),$$

where l, m, and h are \$/mile estimates determined according to the appropriate regression parameters in Table B.1, and $p_{ml}(6)$ represents the probability of a six-year old tractor in medium condition at the beginning of the period shifting to low condition at the end of the period (see Figure 3.3). Table 3.6 displays the associated transition matrices across the entire state space.

3.3.3 Generation of Problem Sets

Utilizing data provided by trucking companies, we construct a collection of problem sets for the SFRPB. In earlier work, Morse [98] generated artificial, time-invariant data sets based on the maintenance and depreciation structure developed in [124]. However, these artificial data sets do not totally capture conditions faced by fleet managers. Therefore, we generate data sets with time-variant parameters considering stochastic deterioration. In addition to the transition probabilities, we need information regarding possible replacement actions and their associated capital expenditures, operating and maintenance costs, salvage values, discount rate, and the length of the planning horizon.

Action Space and Associated Capital Expenditures

We define the action space for each asset p as:

$$A(p) = \begin{cases} 1 & \text{if ``keep''} \\ 2 & \text{if ``replace''} \end{cases}$$

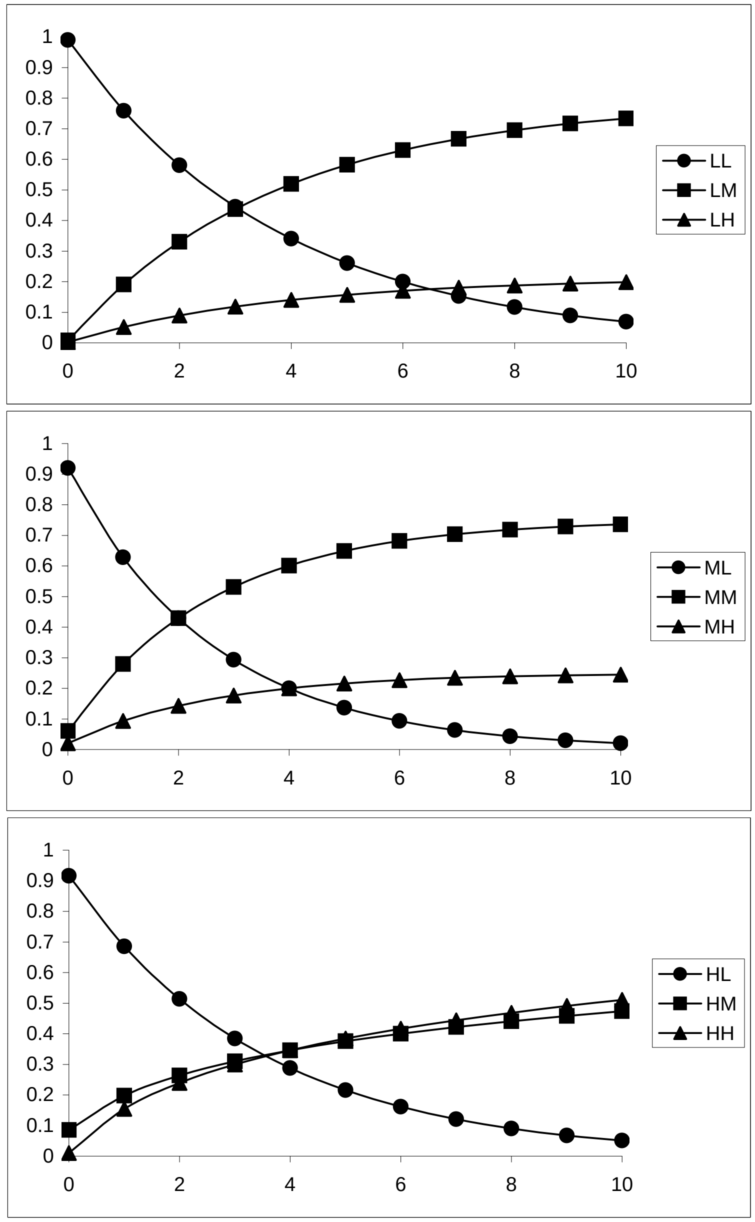

Figure 3.3: Evolution of transition probabilities with asset age.

Table 3.6: Transition Probabilities

(i_n, i_m)	Future State		
Current State	(1,1)	(1,2)	(1,3)
(0,1)	0.99	0.01	0.00
(0,2)	0.93	0.06	0.01
(0,3)	0.88	0.11	0.01

(i_n, i_m)	Future State		
Current State	(2,1)	(2,2)	(2,3)
(1,1)	0.76	0.19	0.05
(1,2)	0.69	0.23	0.08
(1,3)	0.61	0.24	0.15

(i_n, i_m)	Future State		
Current State	(3,1)	(3,2)	(3,3)
(2,1)	0.58	0.33	0.09
(2,2)	0.51	0.36	0.13
(2,3)	0.44	0.32	0.24

(i_n, i_m)	Future State		
Current State	(4,1)	(4,2)	(4,3)
(3,1)	0.44	0.44	0.12
(3,2)	0.39	0.45	0.16
(3,3)	0.31	0.38	0.31

(i_n, i_m)	Future State		
Current State	(5,1)	(5,2)	(5,3)
(4,1)	0.34	0.52	0.14
(4,2)	0.29	0.52	0.19
(4,3)	0.23	0.42	0.35

(i_n, i_m)	Future State		
Current State	(6,1)	(6,2)	(6,3)
(5,1)	0.26	0.58	0.16
(5,2)	0.22	0.57	0.21
(5,3)	0.16	0.45	0.39

(i_n, i_m)	Future State		
Current State	(7,1)	(7,2)	(7,3)
(6,1)	0.20	0.63	0.17
(6,2)	0.16	0.62	0.22
(6,3)	0.11	0.47	0.42

(i_n, i_m)	Future State		
Current State	(8,1)	(8,2)	(8,3)
(7,1)	0.15	0.67	0.18
(7,2)	0.12	0.65	0.23
(7,3)	0.07	0.49	0.44

Table 3.6: *continued*

(i_n, i_m)	Future State		
Current State	(9,1)	(9,2)	(9,3)
(8,1)	0.12	0.70	0.18
(8,2)	0.09	0.67	0.24
(8,3)	0.04	0.50	0.46

(i_n, i_m)	Future State		
Current State	(10,1)	(10,2)	(10,3)
(9,1)	0.09	0.72	0.19
(9,2)	0.07	0.68	0.25
(9,3)	0.03	0.51	0.46

(i_n, i_m)	Future State		
Current State	(11,1)	(11,2)	(11,3)
(10,1)	0.07	0.73	0.20
(10,2)	0.05	0.70	0.25
(10,3)	0.01	0.51	0.48

Table 3.7: Capital Expenditures (in 1000s of 1992 Dollars)

a	$\pi(p, a, t, i)$
1	0
2	$57.983(1.02^t)$

The "replace" option implies salvaging the current tractor and purchasing a new tractor. We assume that technological advance causes a constant annual 2% increase in purchase cost [59]. Capital expenditure information is summarized in Table 3.7.

Maintenance costs

Maintenance costs per mile vary with age and maintenance condition based on a regression fit to data (see Table B.1). To obtain the total annual maintenance cost, we multiply the appropriate maintenance cost per mile by the average annual utilization (100,000 miles). We assume that the advent of new technology has no affect on maintenance costs. Results are summarized in Table 3.8.

Table 3.8: Maintenance Costs, $M(p, \cdot, i)$, in 1000s of 1992 Dollars

	i_m		
i_n	1	2	3
0	3.6262	6.8561	12.0156
1	3.9417	7.1716	12.3311
2	4.2572	7.4871	12.6466
3	4.5727	7.8026	12.9621
4	4.8882	8.1181	13.2776
5	5.2037	8.4336	13.5931
6	5.5192	8.7491	13.9086
7	5.8347	9.0646	14.2241
8	6.1502	9.3801	14.5396
9	6.4657	9.6956	14.8551
10	100	100	100

Operating costs

Operating costs include fuel, oil, and insurance. We focus on fuel expenses, which compose the major portion of operating costs [37]. We assume a tractor's fuel efficiency (miles per gallon) varies linearly with age, but is independent of maintenance condition. Regression analysis on an initial sample of fuel economy data suggests that maintenance condition does not have a significant influence on fuel efficiency. We estimate fuel efficiency with the equation $mpg(i) = 7.669 - 0.215i_n$, where $i = (i_n, i_m)$. This regression line explains 66% of the variance ($R^2 = 0.6648$). Future research will investigate other factors affecting fuel economy, such as driver ability.

We calculate baseline operating costs for each state i with the equation

$$O_0(i) = \frac{\text{average annual mileage } \times \text{ fuel price}}{mpg(i)},$$

where the average annual mileage is 100,000 miles and the fuel price is \$1.21/gallon. To model the effect of new technology, we introduce time-variance by assuming an annual 5% decrease in challenger operating costs [59]. These time-variant costs, $O(p, t, i)$, are calculated for all t according to the following formula:

$$O(p, t, i) = \frac{O_0(i)}{1.05^{t-i_n}}.$$

Salvage Values

We modify a formula in Hartman [59] to consider age and condition as well as improvements in technology to calculate salvage values:

$$S(p, t, i) = (1 - \chi)\pi(p, 2, 0, i)(1.02^{t-i_n})(1 - \gamma i_n - \upsilon(i_m - 1)).$$

where χ is the percent the asset degrades immediately upon purchase, γ is the percent the asset's value degrades each year due to age, and υ is the percent the asset's value may degrade each year due to deterioration in condition. Based on information on the used truck market, we set $\chi = 8.13\%$, $\gamma = 7.69\%$ and $\upsilon = 7.69\%$, so that a 10-year old tractor in high maintenance condition originally purchased for \$57,983 scraps for \$3,361.

Discount Rate

Since we are using cost data presented in constant dollars, i.e., we have already removed the effects of inflation, we use the inflation-free interest rate given by $i' =$

$\frac{i-f}{1+f}$, where i is the market interest rate and f is the inflation rate [106]. Over the period of data collection (1992-2000), the national average annual inflation rate is approximately 3%. We estimate the market interest rate for the trucking companies under consideration to be approximately 10%. These figures imply an inflation-free interest rate of 6.796%, which corresponds to an annual discount factor of $\alpha = 0.93\overline{63}$.

Planning Horizon

In general, nonhomogeneous Markov decision processes can neither be stated in finite information nor solved in finite time. In addition, since the model makes choices using expectations, it biases them with information on scenarios that do not occur. These issues can be remedied by using a rolling horizon technique [4] commonly used to solve in nonhomogeneous Markov decision processes. Rolling horizon updating can be implemented by simply recording a vehicle's age and tracking its maintenance condition on a periodic basis. Following this procedure, we determine the decisions that are to be implemented immediately by modeling their impact on possible future actions.

For deterministic problems, Bean et al. [9] provide empirical evidence that finite horizons of approximately twice the length of the maximum physical age of the asset minimize end-of-the-horizon effects and assure consistent time zero decisions. Hopp et al. [68] extend these results to stochastic problems, discuss various definitions of optimality, and demonstrate the multiplicative interaction of weak ergodicity and the discount factor. This relationship is used [67, 12]to identify appropriate forecast horizons for nonhomogeneous Markov decision processes.

Due to the compounding, diminishing influence contributed by the coefficient of weak ergodicity, the forecast horizon for a stochastic problem is typically shorter than a horizon for a comparable deterministic problem. We leave the explicit computation of forecast horizons for constrained Markov decision processes as a topic of future research; stopping rules for this class of problems have not been established and may not be well defined. Presently, we use the theoretical insight from Bean et al. [9] and Hopp et al. [68], feedback from fleet managers (see Table 3.1), and the fact that the maximum asset life in our study is 10 years to set the planning horizon at 15 years, i.e., $H = 15$.

3.4 Computational Results

In this section, we describe solution approaches for the mathematical program **MIP** described in § 3.2. We describe the general structure of a solution and its interpretation. Armed with this knowledge of the problem, we discuss implementation details of compressed annealing. Results from the compressed annealing algorithm are compared to a trade cycle approach that mimics an age-based replacement strategy commonly used in practice. To evaluate the quality of the solutions, we generate lower bounds on the optimal objective function value using a multiplier adjustment method developed by Morse [98].

3.4.1 Solution Representation

A fleet replacement plan consists of a collection of P individual replacement policies. An individual asset's replacement policy proposes an action for each of the

Table 3.9: Replacement Policy for Asset p in Period t

		i_m		
		1	2	3
	0	$a_{0,1}(t)$	$a_{0,2}(t)$	$a_{0,3}(t)$
	1	$a_{1,1}(t)$	$a_{1,2}(t)$	$a_{1,3}(t)$
i_n	2	$a_{2,1}(t)$	$a_{2,2}(\text{t})$	$a_{2,3}(t)$
	$\vdots$	$\vdots$	$\vdots$	$\vdots$
	10	$a_{10,1}(t)$	$a_{10,2}(t)$	$a_{10,3}(t)$

$(N+1) \times M$ possible states, where N is the maximum life of an asset and M is the number of maintenance conditions. We depict the replacement policy for a single asset p in period t with a $(N+1) \times M$ matrix, as displayed in Table 3.9. Each entry in Table 3.9, $a_{i_n,i_m}(t)$, represents the replacement action for a single asset p when in state $i = (i_n, i_m)$ at the beginning of period t. We relate this to the notation of **MIP** by noting that for an asset p, $y(p, a, t, i) = I\{a = a_{i_n,i_m}(t)\}$.

We consider instances with fleets of 100 assets ($P = 100$), each with a maximum life of 10 years ($N = 10$) and categorized into three maintenance conditions ($M = 3$) over a horizon of 15 years ($H = 15$). For the binary "keep" or "replace" decision, this results in $2^{(49,500)}$ different possible fleet replacement policies.

3.4.2 Trade Cycle Approach

Current replacement strategies for most trucking firms rely upon trade cycles [97]. Trade cycles give rule-of-thumb replacement rules based on years of service, mileage, and/or maintenance history. These rules ignore the stochastic nature of maintenance costs and do not explicitly consider budgets or the variation in costs over time.

Table 3.10: Age-r Trade Cycle for Asset p in Period t

		i_m = 1	2	3
i_n	0	1	1	1
	⋮	⋮	⋮	⋮
	r - 1	1	1	1
	r	2	2	2
	⋮	⋮	⋮	⋮
	10	2	2	2

For computational comparison, we implement an age-based trade cycle approach. To construct an age-r trade cycle, all assets at least r-years old are scheduled for replacement during each period in the planning horizon. For every asset in each period across the planning horizon, an age-r trade cycle implies the individual replacement policy displayed in Table 3.10.

Depending on the tightness of the budgets, the fleet replacement policy implied by the age-r trade cycle may not be feasible. To regain feasibility over the planning horizon, we iteratively modify the trade cycle in a fashion consistent with industrial practice. The procedure, outlined in Table 3.11, is a crude method designed to modify an age-r trade cycle so that the youngest and best-maintained vehicles are kept. Since the distribution of an asset's state in period t is determined by the actions selected in periods 0 through $t-1$ and this procedure obtains feasibility systematically from the beginning of the horizon to the end, it is guaranteed to obtain a feasible solution (if one exists).

Table 3.11: Budget-Feasible Age-r Trade Cycle

Generate age-r trade cycle, Y.
For $k = 0$ to $H - 1$:
Let $i = 0$ and $j = 1$.
While $E(k, Y) > B(k)$:
Modify Y by setting $a_{r+i,j}(k) = 1$ for $p = 1, \ldots, P$.
Increment j by 1.
Re-calculate $E(k, Y)$.
If $i = 10 - r$ and $j = 4$, stop since no feasible policy exists.
If $j = 4$, increment i by 1 and let $j = 1$.

3.4.3 Compressed Annealing Approach

To implement compressed annealing, we need to calibrate the annealing parameters and develop a neighborhood structure to generate candidate solutions. In particular, the neighborhood we describe utilizes the structure of the replacement problem to provide an efficient, yet thorough, search of the solution space.

A primary consideration in selecting a neighbor-generating mechanism is computational efficiency. Given a current replacement policy, we observe that we only need to recalculate the costs and expenditures associated with assets for which decisions have been changed. Therefore, computationally inexpensive neighborhoods will only change the decisions for a few assets.

We generate neighbors for a current fleet replacement policy by making changes in the replacement policies of n randomly-selected assets. Once we have selected an asset p, we randomly select a period, t, and a state, $i = (x, y)$. First, consider the case where $a_{x,y}(t) = 2$. For asset p, we set $a_{i_n,i_m}(t) = 1$ for all (i_n, i_m) pairs such that

Table 3.12: Neighbor Policy Generation

i_n \ i_m	1	2	3
0	1	[2]	2
1	[2]	1	1
2	1	~~2~~ 1	2
⋮	⋮	⋮	⋮
10	1	2	1

$\Rightarrow$

i_n \ i_m	1	2	3
0	1	1	2
1	1	1	1
2	1	1	2
⋮	⋮	⋮	⋮
10	1	2	1

both $i_n \leq x$ and $i_m \leq y$. Similarly, if $a_{x,y}(t) = 1$, we generate a neighbor policy by setting $a_{i_n,i_m}(t) = 2$ for all (i_n, i_m) pairs such that both $i_n \geq x$ and $i_m \geq y$ for asset p. Table 3.12 illustrates the neighborhood scheme where, for asset p in period t, we select $i = (2, 2)$.

This neighbor-generation scheme is motivated by preliminary results using more general neighborhoods and the optimal threshold policies established in Derman [33] for a single stochastically deteriorating asset. While uniform thresholds will not be optimal for the SFRPB in general, we find that inducing threshold-like policies for individual vehicle-years results in quality heuristic solutions within the compressed annealing algorithm. We leave theoretical study of this neighborhood structure as a topic for future research.

To effectively apply compressed annealing to solve the SFRPB, we utilize the experimental design approached of Coy et al. [30] to determine appropriate values for a variety of parameters. While the algorithm performs well over a relatively wide range of parameters, we suggest a robust set in Table 3.13. Compressed annealing runs terminate when the best solution found has not been updated within the last 50

Table 3.13: Compressed Annealing Parameters for SFRPB

Parameter	Value
cooling coefficient (β)	.95
initial acceptance ratio (χ_0)	.95
compression coefficient (γ)	.02
pressure cap ratio (κ)	.99
iterations per temperature	5000
neighborhood size (n)	5

temperature/pressure changes; maximum run time allowed is 10,000 CPU seconds. We use these parameter settings for all the results reported in § 3.4.4.

3.4.4 Computational Comparison

We conduct computational experiments to evaluate the effectiveness of the compressed annealing algorithm versus the trade cycle approach. We vary the budget level, expressed as the number of new asset purchases allowed (ϵ), to assess the impact of constricting financial constraints on the life-cycle costs and the performance of the solution approaches.

We consider a fleet of 100 linehaul tractors over a time horizon of 15 years; the initial age and maintenance condition for each of these tractors is randomly generated. We test instances for values of ϵ varying from 14 to 35 annual replacements. Instances with $\epsilon < 14$ are infeasible; the initial condition of the fleet forces the operation of vehicles beyond their retirement age in order to keep expected capital expenditures within the restrictive annual budgets. Instances with $\epsilon > 35$ are relatively loosely constrained; even trade cycles with rotations as short as three years can be implemented.

For each value of ϵ, we run the three algorithms (coded in C++) on a PC Pentium IV 1.9 GHz (256 MB of RAM). Compressed annealing's performance is the average of five runs from randomly generated initial solutions; computation time for compressed annealing ranged from 3000 CPU seconds to 10,000 CPU seconds. The multiplier adjustment method was allowed to run for 10,000 CPU seconds. The running time for the trade cycle approach was dependent on the budget level, but all runs for $14 \leq \epsilon \leq 35$ took 57 CPU seconds or less.

Table 3.14: Fleet Replacement, $(P = 100, H = 15)$

Budget	Lower	Compressed Annealing			Trade Cycle	
ϵ	Bound	Avg. Value	% Above LB	% Below UB	Value	Cycle
35	24,268	24,633	1.5	0.1	24,651	3
34	24,572	24,673	0.4	0.9	24,893	4
33	24,572	24,741	0.7	0.6	24,893	4
32	24,572	24,779	0.8	0.8	24,980	4
31	24,572	24,833	1.1	0.6	24,980	4
30	24,572	24,918	1.4	0.3	24,994	4
29	24,572	24,969	1.6	0.1	24,994	4
28	24,946	25,003	0.2	0.1	25,033	5
27	24,946	25,131	0.7	1.0	25,390	5
26	25,012	25,220	0.8	0.8	25,418	5

Table 3.14: *continued*

Budget ϵ	Lower Bound	Compressed Annealing			Trade Cycle	
		Avg. Value	% Above LB	% Below UB	Value	Cycle
25	25,012	25,303	1.2	0.5	25,418	5
24	25,012	25,398	1.5	0.3	25,484	5
23	25,012	25,500	2.0	0.5	25,632	5
22	25,012	25,610	2.4	0.1	25,632	5
21	25,538	25,777	0.9	0.4	25,870	6
20	25,538	26,018	1.9	0.1	26,045	6
19	25,538	26,275	2.9	1.5	26,672	6
18	25,538	26,378	3.3	1.2	26,713	6
17	25,943	26,534	2.3	0.7	26,723	7
16	26,138	26,677	2.1	0.6	26,836	8
15	26,163	27,064	3.4	1.7	27,523	8
14	26,167	*	*	*	27,523	8

Table 3.14 illustrates the evolution of solution quality as a function of budget limits. Except for the most tightly constrained feasible instance ($\epsilon = 14$), compressed annealing's hill-climbing search procedure produces replacement policies that obtain a lower expected discounted cost than the trade cycles. The policies supplied by compressed annealing offer an average savings ranging from tens of thousands to hundreds of thousands of dollars in net present value. A budget level of $\epsilon = 14$ results

in a relatively small number of feasible policies since all but the most decrepit tractors must be kept to stay within limits on capital expenditures. In these restrictive cases, we suggest an enumerative method to determine an appropriate replacement policy.

In general, policies obtained by compressed annealing have time zero decisions that differ from time zero decisions in a trade cycle. This policy difference is due to the nature of the search performed by compressed annealing. As compressed annealing searches locally for a good replacement policy, it mitigates a budget violation in period t either by changing fleet replacement decisions in period t, or by altering the age-condition distribution of the fleet through changes in fleet replacement decisions for a period $t' < t$. In this manner, the search procedure evaluates replacement policies that consider the future effects of replacement decisions.

3.4.5 Analysis and Conclusions

Our implementation of compressed annealing on the SFRPB exhibits its potential as a viable solution approach. Using a single set of parameters, compressed annealing's average solution is within 4% of a lower bound for the SFRPB instances considering only a single challenger. While compressed annealing generally offers an improvement of less than 1% over a trade cycle approach, it is a more general technique that can be modified to solve more complicated variants of the SFRPB. Such extensions include the consideration of multiple challengers and time-variant budget constraints (for which no simple trade cycle policies exist).

The study of solution quality as a function of budget tightness illustrates a useful aspect of this analysis. This quantitative analysis allows fleet managers to evaluate the effects of capital budgets. As budgets constrict, fleet managers are forced to extend the trade cycle of tractors. To mitigate the increase in maintenance costs due to these extended life-cycles, some trucking companies implement extensive preventive maintenance programs. Improved maintenance has been cited as the most important reason fleets are able to increase vehicle life-cycles [32]. We speculate that an improved replacement policy over a planning horizon in conjunction with comprehensive maintenance program would provide trucking companies a strategic advantage necessary to thrive in an industry characterized by narrowing profit margins.

Another critical factor affecting fleet management is the used-truck market [14]. In our model, we determine salvage value based on the age and physical condition of the asset. In reality, a major component of salvage value is resale potential, which is affected by external market factors. For instance, glutted used-truck markets force fleet values to plummet. Future work should quantitatively measure how market factors affect replacement decisions. Lower used-truck prices may allow the emergence of used-trucks as an economically-advantageous replacement option. If replacement options are limited either by choice or availability, the plummeting values of defenders will alter the timing replacement decisions.

The effect of a trucking company's existing maintenance philosophy on operating and maintenance costs should not be overlooked. The trucking company that donated our sample data implements a thorough preventive maintenance program. This excellent PM program reduces the variance in costs since major repairs are usually avoided. We conjecture that cost data from trucking companies with a less intensive preventive maintenance program would exhibit more cost variance, further increasing the benefit from detailed modeling of stochastic deterioration and the classification of vehicle condition. In any regard, changes in a company's maintenance program will have a rippling effect throughout the fleet and potentially affect fleet replacement policy. This suggests an interesting vein of research concerned with interaction between maintenance and deterioration, and its effect on replacement strategy.

In particular, there are often maintenance problems that characterize classes of trucks, i.e., challengers of the same age and/or brand, in conjunction with problems related only to individual trucks. For example, one fleet manager noted that frame cracks were a prevalent problem among 1989 trucks of brand X. Through the design of challenger-dependent transition probabilities, this behavior can be captured and group replacement strategies could be developed. Furthermore, truck manufacturers have expressed interest in using the SFRPB with challenger-dependent probability structures as a marketing tool to demonstrate their particular brand's strategic advantage.

CHAPTER IV

Traveling Salesman Problem with Time Windows

As defined by the Council of Logistics Management, logistics is "the process of planning, implementing, and controlling the efficient, effective flow and storage of goods, services, and related information from point of origin to point of consumption for the purpose of conforming to customer requirements." With the production trends of "lean manufacturing" and "just-in-time" operations, an inflated premium is placed on the freight industry to provide timely, efficient service. Potential savings in the form of decreased transportation costs, reduced inventory storage costs, and elimination of penalties due to untimely pick-ups or deliveries may result from improved routing assignments.

In general, time windows play a prominent role in routing problems for business organizations that deal with time-sensitive pick-ups and deliveries. As a gateway problem to more complex vehicle routing issues, we consider the traveling salesman problem with time windows (TSPTW). The TSPTW consists of finding a minimum-cost tour, starting from and returning to the same unique depot, that visits a set of customers exactly once, each of whom must be visited within a specific time window. Practical applications of the TSPTW abound in the industrial and service sectors: bank and postal deliveries, busing logistics, manufacture-and-delivery systems, and automated guided vehicle systems. In addition, the TSPTW is mathematically equivalent to time-sensitive production scheduling problems that are prevalent in manufacturing.

4.1 Model Formulation

To formally define the TSPTW, let $G = (C, A)$ be a finite graph, where $C = \{0, 1, \ldots, n\}$ is the finite set of nodes or customers and $A = C \times C$ is the set of arcs connecting customers. We assume that there exists an arc $(i, j) \in A$ for every $i, j \in C$. A tour is defined by the order in which the n customers are visited and denoted by $\Im = \{p_0, p_1, \ldots, p_n, p_{n+1}\}$, where p_i denotes the index of the customer in the i^{th} position of the tour. Let customer 0 denote the depot, i.e., the beginning and end of the route. To formally express the requirement that the route begin and end at the depot, we append the set C with an additional node, $n + 1$, which also represents the depot. Therefore, $p_0 = 0$ and $p_{n+1} = n + 1$. Each of the remaining n customers occupies one position ranging from p_1 to p_n inclusive.

For $j = 0, \ldots, n$, there is a cost, $c(a_j)$, for traversing the arc $a_j = (p_j, p_{j+1})$. This cost of traversing the arc between the j^{th} customer and the $(j + 1)^{\text{st}}$ customer in the route generally consists of the travel time from customer p_j to customer p_{j+1}, plus any service time at customer p_{j+1}. Associated with each customer i is time window, $[e_i, l_i]$, during which the customer i must be visited. We assume that waiting times are permitted; a customer i can be reached before the beginning of its time window,

e_i, but cannot leave before e_i.

The two primary objective functions considered in the literature are: (1) minimize the sum of the traversal cost along the tour and (2) minimize the time to return to the depot. We deal with the former in order to make comparisons with the results of Calvo [16] and Gendreau et al. [48]. To calculate a tour's cost, we track the arrival time at i^{th} customer, A_{p_i}, and the departure time from i^{th} customer, D_{p_i}. Using these variables, we express our problem as

P minimize $f(\Im) = \sum_{i=0}^{n} c(a_i)$

subject to:

$$A_{p_i} = D_{p_{i-1}} + c(a_{i-1}) \quad \text{for} \quad i = 1, \dots, n+1;$$

$$D_{p_i} = \max\{A_{p_i}, e_{p_i}\} \quad \text{for} \quad i = 1, \dots, n+1;$$

$$D_{p_0} = 0,$$

$$D_{p_i} \leq l_{p_i} \quad \text{for} \quad i = 1, \dots, n+1;$$

$$p_i \in \{1, 2, \dots, n\} \quad \text{for} \quad i = 1, \dots n;$$

$$p_i \neq p_j \quad \text{for} \quad i, j = 1, \dots, n, \; i \neq j;$$

$$p_0 = 0,$$

$$p_{n+1} = n + 1.$$

Note that if we were to consider the minimization of tour completion time, we must track the waiting time of the vehicle at each position of the tour, $W_{p_i} = D_{p_i} - A_{p_i}$, for $i = 0, \dots, n+1$. The term $\sum_{i=0}^{n+1} W_{p_i}$ would then be added to the objective function in **P**.

The TSPTW is composed of two main components, a traveling salesman problem and a scheduling problem. The TSP itself is a NP-hard optimization problem, and the scheduling aspect, with release dates and due dates, presents additional feasibility difficulties. Using a penalty method approach, we partially decompose these two components and apply compressed annealing. We consider infeasible solutions by relaxing the time window constraints $\{D_{p_i} \leq l_{p_i}\}$ into the objective function with an exact penalty function of the form

$$p_\lambda(\Im) \;\; = \;\; \lambda \sum_{i=1}^{n+1} \left[\max(0, D_{p_i} - l_{p_i})\right]^s ,$$

where $s > 0$. Hadj-Alouane [54] proves that penalty functions of this form maintain strong duality between the relaxation and the original formulation.

4.2 Compressed Annealing Approach to the TSPTW

In this section, we tailor an application of compressed annealing for the TSPTW. We discuss implementation details of compressed annealing and compare results from this metaheuristic against the best-known results from the literature.

4.2.1 Parameter Calibration

The need for potentially tedious parameter calibration is often a detractor in metaheuristic approaches. In our implementation, we utilize the statistical design approach of Coy et al. [30] to systematically determine robust parameter settings. While the algorithm performs well over a relatively wide range of parameters, we suggest a robust set in Table 4.1. We terminate the compressed annealing runs when the best tour found has not been updated in the last 75 temperature/pressure changes; we set a minimum of 100 total temperature changes. We perform computational

Table 4.1: Compressed Annealing Parameters for TSPTW

Parameter	Value
cooling coefficient (β)	.95
initial acceptance ratio (χ_0)	.94
compression coefficient (γ)	.06
pressure cap ratio (κ)	.9999
iterations per temperature	30000
minimum number of temperature changes	100

experiments to establish an effective neighborhood structure. We find that a 1-opt neighborhood scheme, in which a single customer and its new insertion position are randomly selected, results in quality solutions.

4.2.2 Computational Comparison

To evaluate the suitability of compressed annealing for the TSPTW, we test the algorithm on benchmark problems from the literature. We implement the compressed annealing algorithm in C++ and run the code on a dual AMD 1.8 gigahertz processor with 1 gigabyte of RAM. For each problem instance, we administer 10 runs from randomly generated starting solutions and report the average solution value.

Table 4.2 displays RC2 instances proposed by Solomon [123]. The RC2 instances contain a mix of randomly-spaced and clustered customers. On average, Compressed annealing outperforms the insertion heuristic of Gendreau et al. [48] on all instances. In comparison to the assignment heuristic of Calvo [16], the average compressed annealing solution matches or improves upon the best known result on 26 of the 30 instances. In addition, compressed annealing obtains feasible solutions in all 30 instances, while Calvo reports feasible solutions in only 28 instances.

Table 4.2: Solomon RC2 Instances

Data Set		Solution Value / CPU Seconds		
Problem	n	Gendreau et al. 1998	Calvo 2000	Compressed Annealing
rc201.1	19	**444.54** / 3.00	**444.54** / 0	**444.54** / 5.5
rc201.2	25	712.91 / 6.98	**711.54** / 0	**711.54** / 6.0
rc201.3	31	795.44 / 14.98	**790.61** / 3	**790.61** / 9.9
rc201.4	25	**793.64** / 6.00	**793.64** / 0	**793.64** / 6.0
rc202.1	32	772.18 / 10.55	772.18 / 8	**772.14** / 11.5
rc202.2	13	**304.14** / 2.35	**304.14** / 0	**304.14** / 5.6
rc202.3	28	839.58 / 6.97	839.58 / 0	**837.72** / 7.5
rc202.4	27	**793.03** / 11.55	**793.03** / 2	**793.03** / 9.1
rc203.1	18	**453.48** / 4.03	**453.48** / 0	**453.48** / 7.7
rc203.2	32	**784.16** / 15.67	**784.16** / 4	**784.16** / 11.5
rc203.3	36	842.25 / 16.02	819.42 / 14	**817.80** / 12.8
rc203.4	14	**314.29** / 2.98	**314.29** / 0	**314.29** / 5.9
rc204.1	44	897.09 / 26.43	**868.76** / 35	880.39 / 14.3
rc204.2	32	679.26 / 15.90	**665.96** / 8	666.88 / 10.9
rc204.3	33	460.24 / 11.18	**455.03** / 4	459.38 / 9.0
rc205.1	13	**343.21** / 1.13	**343.21** / 0	**343.21** / 4.2
rc205.2	26	**755.93** / 7.33	**755.93** / 0	**755.93** / 7.1
rc205.3	34	**825.06** / 42.90	**825.06** / 21	**825.06** / 10.7

Table 4.2: *continued*

Problem	n	Gendreau et al. 1998	Calvo 2000	Compressed Annealing
rc205.4	27	762.41 / 6.58	- /	**760.47** / 7.6
rc206.1	3	**117.85** / 0.01	**117.85** / 0	**117.85** / 3.1
rc206.2	36	842.17 / 33.47	853.31 / 10	**828.16** / 11.3
rc206.3	24	591.2 / 6.75	**574.42** / 0	**574.42** / 9.1
rc206.4	37	845.04 / 31.48	837.54 / 8	**832.54** / 11.5
rc207.1	33	741.53 / 14.76	733.22 / 4	**732.68** / 11.0
rc207.2	30	718.09 / 16.28	- /	**701.25** / 10.5
rc207.3	32	684.4 / 17.25	687.28 / 10	**682.40** / 11.4
rc207.4	5	**119.64** / 0.01	**119.64** / 0	**119.64** / 3.0
rc208.1	37	799.19 / 26.58	**789.25** / 10	793.79 / 16
rc208.2	28	543.41 / 20.53	537.33 / 2	**534.68** / 10.0
rc208.3	35	660.15 / 25.63	649.11 / 8	**640.49** / 11.9

In Table 4.3, we present the TSPTW instances of Langevin [87]. For each customer-time window combination, we list the average solution value over the ten different instances of the problem (unless specified otherwise). We compare compressed annealing to known optimal solutions and the solutions obtained by the heuristic procedure in Calvo [16]. Compressed annealing exhibits promising behavior; it achieves the optimal solution on all the instances with 20 and 40 customers. Direct comparison to optimal solutions is not possible for the problems with 60 customers since the optimal solution is only known in seven of the ten instances with 20-minute time windows, eight of the ten instances with 30-minute time windows, and seven of the ten instances with 40-minute time windows. In these 60-customer cases, compressed annealing matches the solutions obtained in Calvo [16].

Table 4.3: Langevin Instances

Data Set		Solution Value / CPU Seconds		
n	Window Width	Optimal	Calvo 2000	Compressed Annealing
20	30	**724.7** / 0.4	**724.7** / 0.0	**724.7** / 5.2
	40	**721.5** / 0.7	**721.5** / 0.0	**721.5** / 5.1
40	20	**982.7** / 1.7	**982.7** / 0.3	**982.7** / 7.0
	40	**951.8** / 7.3	**951.8** / 0.6	**951.8** / 7.1
60	20	1196.4 (7) / 10	**1215.7** / 5.0	**1215.7** / 9.2
	30	1180.6 (8) / 32	**1183.2** / 5.0	**1183.2** / 12.0
	40	1153.9 (7) / 43	**1160.8** / 10.9	**1160.7** / 14.3

We also consider the TSPTW instances of Dumas [38] in Table 4.4. For each customer-time window combination, we list the average solution value over the five different instances of the problem (unless specified otherwise). We compare compressed annealing to known optimal solutions and the results of Calvo [16]. For these data sets, the average compressed annealing solution matches the optimal solution in 9 of the 28 sets of five instances. In the other cases, compressed annealing's average solution is within 1% of the average optimal for every set of five instances.

For the case with 80 customers with time windows of 60 minutes, Calvo's assignment heuristic only finds feasible solutions for four of the five instances; compressed annealing obtains a feasible solution in all five instances.

Table 4.4: Dumas Instances

Data Set		Solution Value / CPU Seconds		
n	Window Width	Optimal	Calvo 2000	Compressed Annealing
20	20	**361.2** / 0.0	**361.2** / 0.0	**361.2** / 5.0
	40	**316.0** / 0.1	**361.2** / 0.0	**361.2** / 5.0
	60	**309.8** / 0.1	**309.8** / 0.0	**309.8** / 5.0
	80	**311.0** / 0.2	**311.0** / 0.0	**311.0** / 5.0
	100	**275.2** / 1.3	**275.2** / 0.0	**275.2** / 6.2
40	20	**486.6** / 0.1	**486.6** / 3.0	**486.6** / 7.1
	40	**461.0** / 0.0	**461.0** / 3.0	**461.0** / 9.7
	60	**416.4** / 4.4	**416.4** / 4.8	**416.4** / 11.5
	80	**399.8** / 7.5	**399.8** / 5.2	400.0 / 11.7
	100	**377.0** / 31.4	**377.0** / 5.6	377.4 / 12.3
60	20	**581.6** / 0.2	**581.6** / 8.4	**581.6** / 13.4
	40	**590.2** / 0.9	**590.4** / 17.2	590.9 / 15.9
	60	**560.0** / 6.8	**560.0** / 20.2	560.2 / 16.2
	80	508.0 / 46.6	509.0 / 18.0	509.2 / 16.5
	100	514.8 / 199.8	516.4 / 26.2	516.5 / 16.4

Table 4.4: *continued*

n	Window Width	Optimal	Calvo 2000	Compressed Annealing
80	20	**676.6** / 0.4	**676.6** / 43.4	676.8 / 19.6
	40	**630.0** / 2.7	630.4 / 69.2	630.1 / 20.7
	60	**606.4** / 55.3	596.5 (4) / 71.6	607.2 / 21.1
	80	**593.8** / 220.3	594.4 / 59.6	595.5 / 21.3
100	20	757.6 / 0.6	757.8 / 102.6	757.8 / 23.9
	40	701.8 / 7.4	703.6 / 128.6	702.2 / 24.9
	60	696.6 / 108.0	696.6 / 148.0	697.8 / 24.9
150	20	**868.4** / 2.4	868.6 / 419.8	869.4 / 35.8
	40	**834.8** / 115.9	837.4 / 529.6	835.7 / 36.4
	60	**805.0** / 463.0	820.4 / 630.0	821.9 / 36.9
200	20	**1009.0** / 6.7	1010.0 / 1456.2	1010.3 / 50.5
	40	**984.2** / 251.4	985.4 / 2105.8	986.7 / 50.4

No known exact method can solve TSPTW instances with wide time windows, making heuristic approaches to these instances particularly useful. In Table 4.5, we consider the TSPTW instances generated in Gendreau et al. [48] by extending the time windows of the instances in Table 4.4. The average compressed annealing solution is the best-known result for 14 of the 27 different sets of instances.

Table 4.5: Gendreau Instances

Data Set		Solution Value / CPU Seconds		
n	Window Width	Gendreau 1998	Calvo 2000	Compressed Annealing
20	120	269.2 / 4	267.2 / 0	**265.6** / 7
	140	263.8 / 4	259.6 / 0	**232.8** / 8
	160	261.2 / 5	260.0 / 0	**218.2** / 8
	180	259.8 / 6	244.6 / 0	**236.6** / 8
	200	245.2 / 6	243.0 / 0	**241.0** / 8
40	120	372.8 / 18	**360.8** / 5	378.5/ 12
	140	356.2 / 19	**348.4** / 9	364.9 / 12
	160	348.0 / 20	337.2 / 10	**326.9** / 12
	180	328.2 / 17	**326.8** / 12	333.9 / 12
	200	326.2 / 23	**315.2** / 16	**314.9** / 12
60	120	492.0 / 52	483.4 / 30	**453.3** / 16
	140	454.8 / 49	454.4 / 28	**454.2** / 16
	160	451.6 / 48	**448.6** / 34	465.3 / 16
	180	439.2 / 52	432.8 / 41	**424.8** / 16
	200	439.6 / 44	**428.0** / 57	430.4 / 16
80	120	581.8 / 121	549.8 / 64	**543.7** / 20
	140	555.2 / 94	525.6 / 75	**512.7** / 20
	160	524.8 / 86	**502.8** / 82	513.3 / 21

Table 4.5: *continued*

n	Window Width	Gendreau 1998	Calvo 2000	Compressed Annealing
	180	511.0 / 99	**489.0** / 116	505.0 / 21
	200	508.6 / 112	**484.0** / 158	485.6 / 20
100	80	675.6 / 118	668.0 / 139	668.9 / 25
	100	671.2 / 130	644.0 / 119	644.9 / 24
	120	624.6 / 204	614.4 / 167	**604.0** / 24
	140	634.6 / 208	591.4 / 201	**581.5** / 24
	160	585.2 / 215	570.4 / 214	589.5 / 24
	180	585.2 / 225	566.0 / 245	568.1 / 24
	200	588.6 / 168	555.6 / 242	562.4 / 24

Comparing computational times is difficult due to differences in computer platforms. Processor comparisons suggest that our algorithm is slower than Calvo's assignment heuristic and Gendreau et al.'s insertion heuristic. However, given today's processor speeds and the general nature of our implementation, our algorithm is certainly capable of solving reasonably large problems in adequate time. In addition, as our results show, run-times for compressed annealing are only minimally affected by increasing numbers of customers and time-window widths. This result is in contrast to the performance of Calvo's and Gendreau's algorithms under the same conditions. Thus, the result suggests that compressed annealing is particularly valuable in circumstances involving large numbers of customers or wide time windows. In addition, our computational experience has shown that, by revising the termination criteria such that the algorithm terminates when the best tour found has not been updated in 25 temperature/pressure changes, run times can be reduced by 20% to 30% for almost all problems. This reduction in computation time also only minimally reduces the quality of the average solution as most average solutions are still within 1% of the optimal or previously best-known heuristic solution.

4.3 Conclusions and Future Considerations

We have presented a solution approach to the TSPTW, a difficult combinatorial problem, utilizing compressed annealing. Using a variable penalty function and stochastic search, we consider solutions infeasible with respect to time windows during our search for good solutions. Computational testing on four series of TSPTW problems demonstrates the potential of the compressed annealing algorithm. Near-optimal solutions can be obtained at a reasonable computational cost in most cases, and feasible solutions are found in every instance.

Future research may include exploring the effect of different penalty functions. In particular, the implementation of the augmented lagrangian penalty function common in continuous mathematical programming could provide an interesting extension to this work. In the current implementation, a single multiplier is used to penalize time window violations for the n customers. Future work may consider the impact of customer-specific penalty multipliers on the search procedure.

We consider the objective of minimizing the sum of the travel time along the TSPTW tour. The performance of compressed annealing with a different objective function, namely the minimization of tour completion time, should be evaluated.

CHAPTER V

Conclusions

This thesis investigates a variant of simulated annealing appropriately parameterized with the penalty multiplier of an integrated penalty function. We research compressed annealing's theoretical behavior and analyze its performance on applications in vehicle replacement and routing. The objectives of this study are to establish conditions for the convergence of compressed annealing and utilize the insight from this theoretical study to enhance the practical performance of the algorithm.

5.1 Summary of Contribution

In Chapter II, we provide a theoretical framework to analyze the integration of an appropriately designed penalty function and the hill-climbing search procedure of simulated annealing. Introducing the novel concept of *levels*, we modify the mechanics in Hajek [56] to cope with the dynamics induced by the variable penalty method. We establish necessary and sufficient conditions for this constrained annealing approach that are the strongest in the literature. We conclude Chapter II with a discussion on the implementation details of the algorithm. In particular, the theoretical results suggest the form of joint cooling and compression schedules.

The parameterized penalty multiplier of compressed annealing allows the algorithm to consider variably-penalized infeasible solutions in its search for optimal or near-optimal solutions. Due to its dynamic search procedure, compressed annealing generally outperforms simulated annealing with a suitable static penalty method. For a traditional simulated annealing approach, setting a static penalty multiplier that allows an adequate search of the solution space often proves to be a difficult chore.

Incorporating maintenance data from trucking companies, we develop a replacement model in Chapter III considering the combined effects of stochastic deterioration, budget constraints, and time-variance induced by constant technological change. Using a penalty function to relax nonlinear constraints on expected expenditure, we implement compressed annealing to determine near-optimal replacement policies over a finite planning horizon. Computational results demonstrate the superiority of the compressed annealing search over a trade cycle approach commonly used in practice.

Chapter IV presents the implementation of compressed annealing on the traveling salesman problem with time windows. Computational experiments on prominent data sets in the literature demonstrate compressed annealing's efficacy. Compressed annealing compares favorably with benchmarks in the literature, obtaining best-known results in numerous instances.

Our implementation of compressed annealing on the SFRPB and TSPTW reveals that nearly identical parameters (see Tables 3.13 and 4.1) result in good solutions. Besides termination criteria, the only significant difference between the two implementations pertained to the number of iterations per temperature. Since the objective function for the SFRPB is computationally expensive relative to the TSPTW objective, fewer iterations are performed at each temperature/pressure setting. Accounting for differences in computational complexity, the similarity between the parameter values for these two different problems vouches for the robustness of the compressed annealing algorithm.

5.2 Future Work

There exist many open research issues in each of the three major components of this study. Throughout this thesis, we noted future research topics which we summarize here.

In our theoretical analysis of compressed annealing, we assume that the values of each penalty multiplier are the same, masking the duality of the problem. Future research is necessary to explore the effect of having distinct values for each penalty multiplier. In the proof of convergence, we focus on the tail behavior of the algorithm. More work is required to scrutinize the search performance during the transient period of compression. One future endeavor is to create a class of problems for which we can derive analytical results on probability of convergence and rate of convergence in order to garner further insight on the algorithm's behavior. A more general objective is to further explore the variable penalty multiplier approach within the framework of other heuristics.

While the computational results in this study reflect favorably on compressed annealing's potential, implementation on other applications will test its robustness. For problems where even finding a feasible solution is difficult, we must focus on techniques to enhance the algorithm's ability to escape a *cup* of solutions composed entirely of infeasible solutions. As the penalty multiplier increases, the solution value of each infeasible solution increases in an amount proportional to its violation. Thus, the algorithm can become trapped in a local network of infeasible solutions, i.e., stuck in the bottom of a flying cup. One possible modification to the compressed annealing algorithm would be to implement an adaptive approach to cooling and compression. Feedback from the search progress would intelligently guide the manipulation of parameters.

For the asset replacement problem, there are a number of interesting extensions to consider in the presence of stochastic deterioration and budget limits. One intriguing possibility is to consider multiple replacement options, i.e., new assets, used assets, rebuilt assets, leased assets, etc. The action space could also be extended to include maintenance actions such as the major rebuild of a defender, creating a more holistic maintenance/replacement model. No simple trade cycle approach exists in the presence of multiple replacement options, so there is potential for marked improvement for such replacement strategies.

Possible refinement of the definition of a vehicle's *state* should be investigated. Recall that a vehicle's maintenance category, i_m, is determined solely by its maintenance cost per mile. As a future research project, we suggest expanding the definition of a vehicle's state to account for possible variance in operating cost as a result in the deterioration of physical condition.

We exclusively study the replacement of linehaul tractors. A holistic approach that simultaneously considers linehaul and P&D fleets would allow fleet managers to incorporate the common conversion of linehaul tractors into pick-up and delivery service. Economies of scale in purchase costs and dis-economies of scale in maintenance costs may also be appropriate concerns for some asset replacement problems; fleet managers comment that early replacement is occasionally motivated by special financing rates and discounts offered by tractor dealers.

Alternate techniques to handle budgets over a planning horizon should also be explored. Morse [98] suggests a multiple knapsack formulation in conjunction with rolling horizon procedure. We treat budget constraints equally across the horizon. However, the budget limits in the first few years of a planning horizon are the most influential. For periods at the end of the planning horizon, estimating budgets becomes difficult and distribution of expected expenditure becomes increasingly uninformative. A time-discounted penalty approach may be appropriate.

Our results on the traveling salesman problem with time windows provide a foundation for extending a compressed annealing approach to other constrained vehicle routing problems. In particular, we plan on exploring the vehicle routing problem with time windows. Other interesting applications include the joint vehicle routing and maintenance problem [126, 53] (prevalent in the aircraft industry) and the vehicle fleet mix problem [102].

APPENDICES

APPENDIX A

Information Defining Maintenance Condition

In the stochastic fleet replacement model with budget constraints, we assume that a vehicle's future maintenance costs are dependent only on its maintenance costs over the past year, i.e., the stochastic process describing the evolution of maintenance category over a vehicle's life possesses the Markov property. Using sample data from the trucking industry, we test the validity of the Markov assumption. We compare various methods of considering a tractor's maintenance history to forecast future maintenance costs. While this comparison does not exhaustively test the countless ways to forecast future maintenance costs, it provides insight on the validity of the Markovian assumption.

Past Year's Maintenance Cost : This approach assumes that a vehicle's future maintenance costs are only dependent on the maintenance costs from the past year. To implement this assumption, we categorize a vehicle-year by the maintenance costs over the year.

Lifetime Average Maintenance Cost : This approach assumes that a vehicle's future maintenance costs are dependent on its entire maintenance history. To test the

Table A.1: Various Methods of Defining Maintenance Condition

Criteria	MAD ($/Mile)
Past Year's Maintenance Condition	0.1873
Lifetime Average Maintenance Cost	0.0207
Weighted Maintenance History	0.1890

effectiveness of this assumption, we categorize a vehicle-year by its average maintenance cost per mile over the vehicle's life up to the present.

Weighted Maintenance History : This approach also assumes that a vehicle's future maintenance costs are dependent on its entire maintenance history. However, in this case, we weight the maintenance history via an exponential smoothing technique. Specifically, we weight the previous year's maintenance cost per mile by 75% and the rest of the maintenance history by 25%.

As shown in Table A.1, the sample data from trucking industry suggests that defining maintenance category by the past year's maintenance category suffices.

In practice, the appraisal of a vehicle's condition may be based on many factors including fleet manager's intuition, maintenance history, driver feedback, and physical appearance. The model's flexibility to allow the fleet manager to determine his own categorization criterion and instill some of his/her own intuition is viewed as a strategic advantage.

APPENDIX B

Aggregation Error

In the stochastic fleet replacement model with budget constraints, we model stochastic deterioration by adopting the aggregation scheme in Derman [33] to define maintenance categories. Aggregation loses some of the information from the continuous range of maintenance costs, and thus induces modeling error. For a survey of aggregation techniques, Rogers et al. [117] discuss aggregate modeling and the level of detail required to model a given problem.

Using data from a trucking company, we measure the aggregation error of the SFRPB. We predict a vehicle's future maintenance cost per mile of operation given the age (a quantitative predictor) and the maintenance condition at the end of the period (a qualitative predictor). This represents a "perfect information" scenario in the sense that we are predicting future maintenance cost given its range, i.e., maintenance category. Analysis with a mixture of quantitative and qualitative predictors requires statistical analysis of covariance [39]. The analysis of covariance reveals that there is no slope interaction between age and maintenance condition, allowing us to fit the data with three separate regression lines (one for each condition level) that

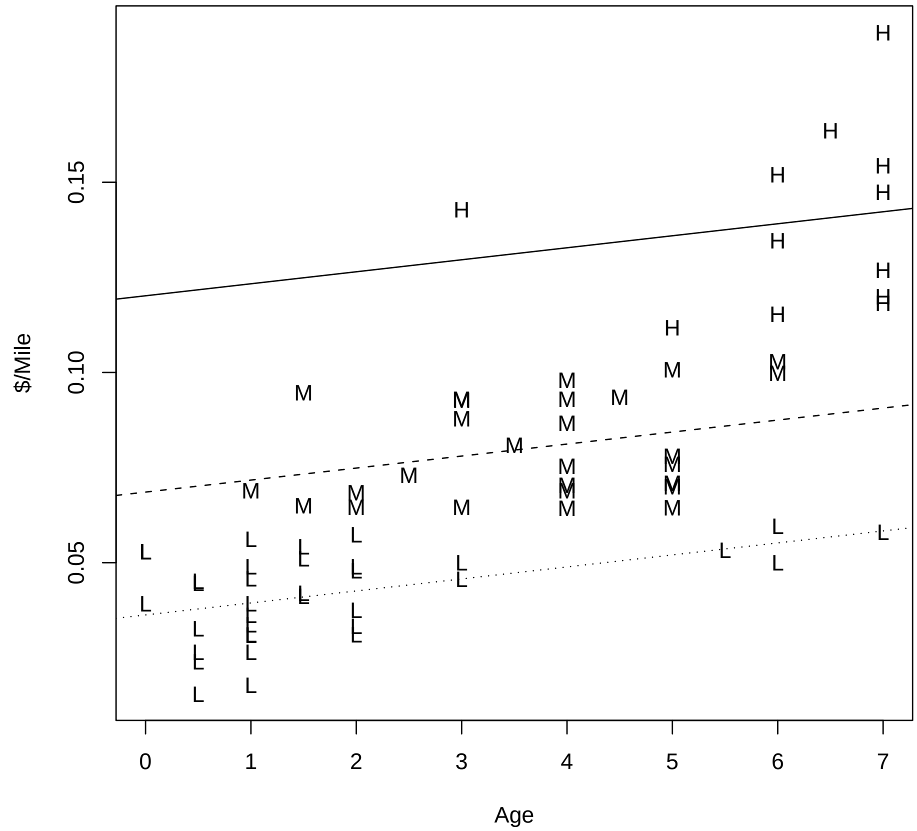

Figure B.1: Regression fit based on age and maintenance condition at end of period, resulting in $R^2 = 0.865$.

share the same slope and differ by their y-intercept. Figure B.1 illustrates the fit and Table B.1 displays the parameters for the respective linear functions. Age and future maintenance condition explain 86.5% of the cost variation. The remaining 13.5% of the variance is the result in the "within-category" variation due to aggregation; even by correctly identifying a vehicle's condition, we have only narrowed the range of values in which the realized cost per mile will occur.

In Table B.2, we benchmark the our stochastic model (implementing the transition probabilities in Table 3.6) with the "perfect information" regression considering age and future maintenance condition. Due to aggregation error, the "perfect in-

Table B.1: Regression Parameters

Maintenance Condition	Y-Intercept	Slope
Low	.036262	0.003155
Medium	.068561	0.003155
High	.120156	0.003155

Table B.2: Accuracy of Probabilistic Model

Model	MAD ($/Mile)
Age + Future Maintenance Condition	0.01106
Expected Value Approach	0.01873

formation" model still results an mean absolute deviation of $0.01106/mile. The probabilistic model results in a mean absolute deviation of $0.01803/mile, implying that the stochastic deterioration increases error by $0.00697/mile, or an average of $697 per tractor per year (100,000 miles per year × $0.00692/mile).

APPENDIX C

Construction of Transition Probabilities

After categorizing annual maintenance cost per mile for each tractor across its lifetime (as in Table 3.5), we develop transition probabilities based on frequency. For example, if a tractor is six years old and in medium condition at the beginning of year 6, we estimate the probability the tractor transitions to low, medium, or high condition by calculating number of times that six-year old tractors beginning in medium condition were in low, medium, and high condition at the end of a period. Figure C.1 displays the empirical probabilities induced by the frequency tabulations. Note that these empirical probabilities fluctuate widely; a larger data sample would be necessary to smooth these curves.

We investigate smoothing curves with a variety of shapes (convex, linear, and concave) that dictate the evolution of the transition probabilities over the lifetime of a tractor. Empirical observation reveals that representing the decay of the probability of *low* → *low*, *medium* → *low*, and *high* → *low* transitions as a convex function of tractor age results in a good fit. Specifically, the functions used to smooth the

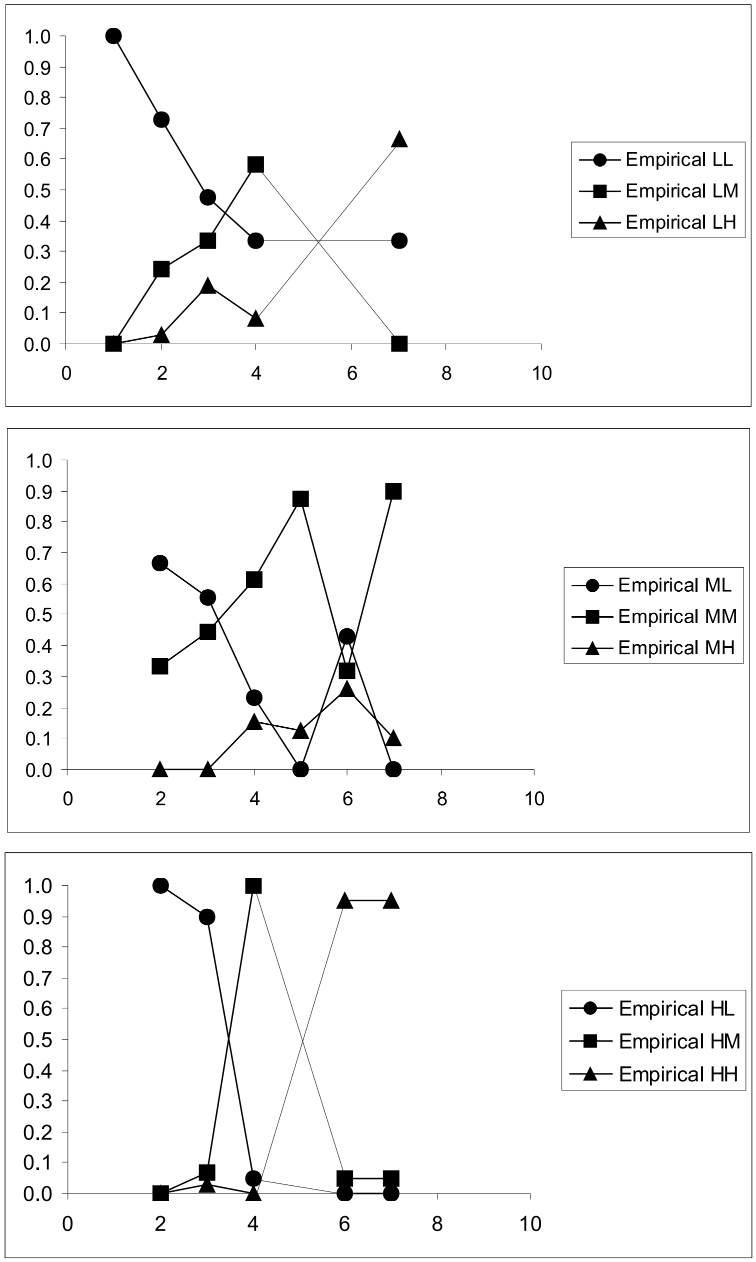

Figure C.1: Empirical unsmoothed transition probabilities.

frequency data are given by

$$
\begin{aligned}
p_{xl}(i_n) &= c_x \left(a_x^{-b_x i_n}\right) \text{ for } x = l, m, h \\
p_{xm}(i_n) &= d_x \left[1 - c_x \left(a_x^{-b_x i_n}\right)\right] \text{ for } x = l, m, h \\
p_{xh}(i_n) &= g_x \left[1 - c_x \left(a_x^{-b_x i_n}\right)\right] \text{ for } x = l, m, h,
\end{aligned}
$$

where $d_x + g_x = 1$ for $x = l, m, h$ to assure the probabilities sum to one. The parameters a_x, b_x, c_x, d_x, and g_x are estimated to fit the data in Figure C.1.

BIBLIOGRAPHY

BIBLIOGRAPHY

[1] E. Aarts and J. Korst. *Simulated Annealing and Boltzmann Machines*. John Wiley & Sons, Chichester, 1989.

[2] E. Aarts and J. Lenstra. *Local Search in Combinatorial Optimization*. John Wiley & Sons, Chichester, 1997.

[3] A. Alchian. *Economic Replacement Policy*. Publication R-224, The RAND Corp., Santa Monica, CA, 1952.

[4] J. Alden and R. Smith. Rolling horizon procedures in nonhomogeneous Markov decision processes. *Operations Research*, 40(Supp. No. 2), 1992.

[5] S. Anily and A. Federgruen. Simulated annealing methods with general acceptance probabilities. *Journal of Applied Probability*, 35:657–667, 1987.

[6] R. Azencott. *Simulated Annealing: Parallelization Techniques*. John Wiley & Sons, New York, 1992.

[7] E. Baker. An exact algorithm for the time constrained traveling salesman problem. *Operations Research*, 31:938–945, 1983.

[8] M. S. Bazaraa, H. D. Sherali, and C. Shetty. *Nonlinear Programming*. John Wiley & Sons, 2nd edition, 1993.

[9] J. Bean, J. Lohmann, and R. L. Smith. A dynamic infinite horizon replacement economy decision model. *The Engineering Economist*, 30:99–120, 1985.

[10] J. Bean, J. Lohmann, and R. L. Smith. Equipment replacement under technological change. *Naval Research Logistics*, 41:117–128, 1994.

[11] J. C. Bean, J. R. Birge, and R. L. Smith. Aggregation in dynamic programming. *Operations Research*, 35:215–220, 1987.

[12] J. C. Bean, W. J. Hopp, and I. Duenyas. A stopping rule for forecast horizons in nonhomogeneous Markov decision processes. *Operations Research*, 40:1188–1199, 1992.

[13] G. Bethuyune. Optimal replacement under variable intensity of utilization and technological progress. *The Engineering Economist*, 43(2):85–105, 1998.

[14] C. Birkland. The used truck market. *Fleet Equipment*, pages 44–48, May 1997.

[15] E. Bonomi and J. Lutton. The N-city traveling salesman problem: Statistical mechanics methods and Metropolis algorithm. *SIAM Review*, 36:551–568, 1984.

[16] R. W. Calvo. A new heuristic for the traveling salesman problem with time windows. *Transportation Science*, 34:113–124, 2000.

[17] W. B. Carlton and J. W. Barnes. Solving the traveling-salesman problem with time windows using tabu search. *IIE Transactions*, 28:617–629, 1996.

[18] Y. Caseau and F. Laburthe. Heuristics for large constrained vehicle routing problems. *Journal of Heuristics*, 5:281–303, 1999.

[19] V. Cerny. A thermodynamical approach to the travelling salesman problem: An efficient simulation algorithm. *Journal of Optimization Theory and Applications*, 45:41–55, 1985.

[20] S. Chand and S. Sethi. Planning horizon procedures for machine replacement models with several possible replacement alternatives. *Naval Research Logistics Quarterly*, 29(3):483–493, 1982.

[21] T. Cheevaprawatdomrong and R. L. Smith. A paradox in equipment replacement under technological improvement. *Operations Research Letters*, 31:77–82, 2003.

[22] Z.-L. Chen. Solution algorithms for the parallel replacement problem under economy of scale. *Naval Research Logistics*, 45:279–295, 1998.

[23] T. Chiang and Y. Chow. On the convergence rate of annealing processes. *SIAM Journal on Control and Optimization*, 26:1455–1470, 1988.

[24] W.-C. Chiang and R. A. Russell. Simulated annealing metaheuristics for the vehicle routing problem with time windows. *Annals of Operations Research*, 63:3–27, 1996.

[25] N. Christofides, A. Mingozzi, and P. Toth. State space relaxation procedures for the computation of bounds to routing problems. *Networks*, 11:145–164, 1981.

[26] N. Christofides, A. Mingozzi, P. Toth, and C. Sandi. *Combinatorial Optimization*. John Wiley and Sons, Chichester, 1979.

[27] D. Coit and A. Smith. Penalty guided genetic search for reliability design optimization. *Computers and Industrial Engineering*, 30:895–904, 1996.

[28] D. Coit, A. Smith, and D. Tate. Adaptive penalty methods for genetic optimization of constrained combinatorial problems. *Journal on Computing*, 8:173–182, 1996.

[29] L. Costa and P. Oliveira. Evolutionary algorithms approach to the solution of mixed integer non-linear programming problems. *Computers & Chemical Engineering*, 25:257–266, 2001.

[30] S. P. Coy, B. L. Golden, G. C. Runger, and E. A. Wasil. Using experimental design to find effective parameter settings for heuristics. *Journal of Heuristics*, 7:77–97, 2000.

[31] K. Deb. An efficient constraint handling method for genetic algorithms. *Computer Methods in Applied Mechanics and Engineering*, 186:311–338, 2000.

[32] B. Deierlein. Life-cycle costing. *Fleet Equipment*, pages 31–35, May 1991.

[33] C. Derman. On optimal replacement rules when changes of state are Markovian. In R. Bellman, editor, *Mathematical Optimization Techniques*. University of California Press, Berkeley, 1963.

[34] M. Desrochers, J. Desrosiers, and M. Solomon. A new optimization algorithm for the vehicle routing problem with time windows. *Operations Research*, 40(2):342–354, 1992.

[35] K. A. Dowsland. Simulated annealing. In C. R. Reeves, editor, *Modern Heuristic Techniques for Combinatorial Problems*, chapter 2. John Wiley and Sons, New York, 1993.

[36] R. Drinkwater and N. Hastings. An economic replacement model. *Operational Research Quarterly*, 18:121–138, 1967.

[37] R. Duff. personal interview, August 1999.

[38] Y. Dumas, J. Desrosiers, E. Gelinas, and M. M. Solomon. An optimal algorithm for the traveling salesman problem with time windows. *Operations Research*, 43:367–371, 1995.

[39] J. Faraway. Practical regression and anova using R. Textbook for Statistics 501, December 1999.

[40] M. L. Fisher, K. Jornsten, and O. B. Madsen. Vehicle routing with time windows–two optimization algorithms. *Operations Research*, 45:488–498, 1997.

[41] F. Focacci, A. Lodi, and M. Milano. A hybrid exact algorithm for the tsptw. *INFORMS Journal on Computing*, 14(4):403–417, 2002.

[42] A. Frigerio and G. Grillo. Simulated annealing with time-dependent energy function. *Mathematische Zeitschrift*, 213:97–116, 1993.

[43] B.-L. Garcia, J.-Y. Potvin, and J.-M. Rousseau. A parallel implementation of the tabu search heuristic for vehicle routing problems with time window constraints. *Computers and Operations Research*, 21:1025–1033, 1994.

[44] M. Garey and D. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W.H. Freeman, San Francisco, 1979.

[45] D. Geman. Random fields and inverse problems in imaging. In P. Hennequin, editor, *Lecture Notes in Mathematics*, pages 117–193. Springer-Verlag, 1990.

[46] S. Geman and D. Geman. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-6:721–741, 1984.

[47] M. Gendreau, A. Hertz, and G. Laporte. New insertion and postoptimization procedures for the traveling salesman problem. *Operations Research*, 40:1086–1094, 1992.

[48] M. Gendreau, A. Hertz, G. Laporte, and M. Stan. A generalized insertion heuristic for the traveling salesman problem with time windows. *Operations Research*, 46:330–335, 1998.

[49] B. Gidas. Nonstationary Markov chains and convergence of the annealing algorithm. *Journal of Statistical Physics*, 39:73–131, 1985.

[50] F. Glover and M. Laguna. *Tabu Search*. Kluwer Academic Publishers, Boston, 1997.

[51] D. E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley, Reading, MA, 1989.

[52] M. Gopalakrishnan, K. Ding, J. Bourjolly, and S. Mohan. A tabu-search heuristic for the capacitated lot-sizing problem with set-up carryover. *Management Science*, 47:851–863, 2001.

[53] R. Gopalan and K. T. Talluri. The aircraft maintenance problem. *Operations Research*, 46(2):260–271, 1998.

[54] A. Hadj-Alouane. *A Dual Genetic Algorithm for Nonlinear Integer Programs*. PhD thesis, University of Michigan, 1994.

[55] A. Hadj-Alouane and J. Bean. A genetic algorithm for the multiple-choice integer program. *Operations Research*, 45, 1997.

[56] B. Hajek. Cooling schedules for optimal annealing. *Mathematics of Operations Research*, 13, 1988.

[57] J. Hartman. A general procedure for incorporating asset utilization decisions into replacement analysis. *The Engineering Economist*, 44:217–238, 1999.

[58] J. Hartman. The parallel replacement problem with demand and capital budgeting constraints. *Naval Research Logistics*, 47:40–56, 2000.

[59] J. Hartman. An economic replacement model with probabilistic asset utilization. *IIE Transactions*, 33:717–727, 2001.

[60] J. Hartman and J. Lohmann. Multiple options in parallel replacement analysis: Buy, lease, or rebuild. *Engineering Economics*, 42:223–248, 1997.

[61] J. C. Hartman. A note on the fleet design problem. *The Engineering Economist*, 41(1):69–75, 1995.

[62] J. C. Hartman. Parallel replacement analysis under variable asset utilization and stochastic demand. submitted to European Journal of Operational Research, March 2003.

[63] N. Hastings. The repair limit replacement method. *Operational Research Quarterly*, 20:337–349, 1969.

[64] J. Hertz, A. Krogh, and R. Palmer. *Introduction to the Theory of Neural Computation*. Addison-Wesley, Redwood City, CA, 1991.

[65] J. Holland. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor, 1975.

[66] R. Holley and D. Stroock. Simulated annealing via Sobolev inequalities. *Communications in Mathematical Physics*, 115:553–569, 1988.

[67] W. J. Hopp. Identifying forecast horizons in nonhomogeneous Markov decision processes. *Operations Research*, 37(2):339–343, 1989.

[68] W. J. Hopp, J. C. Bean, and R. L. Smith. A new optimality criterion for nonhomogeneous Markov decision processes. *Operations Research*, 35:875–883, 1987.

[69] W. J. Hopp, P. C. Jones, and J. Zydiak. A further note on parallel machine replacement. *Naval Research Logistics*, 40:575–579, 1993.

[70] W. J. Hopp and S. K. Nair. Timing replacement decisions under discontinuous technological change. *Naval Research Logistics*, 38:203–220, 1991.

[71] W. J. Hopp and S. K. Nair. Markovian deterioration and technological change. *IIE Transactions*, 26:74–82, 1994.

[72] A. Johnson and S. Jacobson. On the convergence of generalized hill climbing algorithms. *Discrete Applied Mathematics*, 119:37–57, 2002.

[73] D. Johnson, C. Aragon, L. McGeoch, and C. Schevon. Optimization by simulated annealing: An experimental evaluation; part i, graph partitioning. *Operations Research*, 37:865–892, 1989.

[74] P. C. Jones and J. L. Zydiak. The fleet design problem. *The Engineering Economist*, 38(2):83–98, 1993.

[75] P. C. Jones and J. L. Zydiak. A further note on the fleet design problem. *The Engineering Economist*, 44(2):197–201, 1999.

[76] P. C. Jones, J. L. Zydiak, and W. J. Hopp. Parallel machine replacement. *Naval Research Logistics*, 38:351–365, 1991.

[77] P. Kall and S. W. Wallace. *Stochastic Programming*. John Wiley and Sons, Chichester, 1994.

[78] N. Karabakal, J. Lohmann, and J. Bean. Parallel replacement under capital rationing constraints. *Management Science*, 40:305–319, 1994.

[79] N. Karabakal, J. Lohmann, and J. Bean. Solving large replacement problems with budget constraints. *The Engineering Economist*, 45:290–308, 2000.

[80] J. Kemeny and L. Snell. *Finite Markov Chains*. Van Nostrand, Princeton, N.J., 1960.

[81] S. Kirkpatrick, C. Gellat, and M. Vecchi. Optimization by simulated annealing. *Science*, 220:671–680, 1983.

[82] N. Kohl, J. Desrosiers, O. B. Madsen, M. M. Solomon, and F. Soumis. 2-path cuts for the vehicle routing problem with time windows. *Transportation Science*, 33(1), 1999.

[83] N. Kohl and O. B. Madsen. An optimization algorithm for the vehicle routing problem with time windows based on Lagrangean relaxation. *Operations Research*, 45(395-406), 1997.

[84] A. Kolen, A. R. Kan, and H. Trienekens. Vehicle routing with time windows. *Operations Research*, 35(2):1987, 1987.

[85] P. V. Laarhoven and E. Aarts. *Simulated Annealing: Theory and Applications*. Kluwer, Dordrecht, 1988.

[86] M. Laguna. *Scatter Search*. Kluwer Academic Publishers, Boston, 2003.

[87] A. Langevin, M. Desrochers, J. Desrosiers, S. Gélinas, and F. Soumis. A restricted dynamic programming heuristic for the time dependent traveling salesman problem. *European Journal of Operational Research*, 90:45–55, 1996.

[88] E. Lawler. *Combinatorial Optimization: Networks and Matroids*. Holt, Rinehart, and Winston, New York, 1981.

[89] J. R. Lohmann. A stochastic replacement economic decision model. *IIE Transactions*, 18(2):182–194, 1986.

[90] M. Lundy and A. Mees. Convergence of an annealing algorithm. *Math. Prog.*, 34:111–124, 1986.

[91] B. Mahon and R. Bailey. A proposed improved replacement policy for army vehicles. *Operational Research Quarterly*, 2:477–494, 1975.

[92] P. Del Moral and L. Miclo. On the convergence and applications of generalized simulated annealing. *SIAM Journal on Control and Optimization*, 37:1222–1250, 1999.

[93] D. Vander Veen. *Parallel Replacement under Nonstationary Deterministic Demand*. PhD thesis, University of Michigan, 1985.

[94] N. Metropolis, A. Rosenbluth, M. Rosenbluth, A. Teller, and E. Teller. Equation of state calculation by fast computing machines. *Journal of Chemical Physics*, 21:1087–1091, 1953.

[95] Z. Michalewicz. *Genetic Algorithms + Data Structures = Evolution Programs*. Springer-Verlag, New York, third edition, 1996.

[96] D. Mitra, F. Romeo, and A. Sangiovanni-Vincentelli. Convergence and finite-time behavior of simulated annealing. *Advances in Applied Probability*, 18:747–771, 1986.

[97] T. Moore. Trade cycle strategies. *Fleet Owner*, 91:50–53, 1996.

[98] C. Morse. *Stochastic Equipment Replacement with Budget Constraints*. PhD thesis, University of Michigan, 1997.

[99] C. Morse, J. Bean, and N. Karabakal. Nonhomogeneous Markov decision processes linked by side constraints. Department of Industrial and Operations Engineering, 1998.

[100] S. K. Nair and W. J. Hopp. A model for equipment replacement due to technological obsolescence. *European Journal of Operational Research*, 63:207–221, 1992.

[101] R. Oakford, J. R. Lohmann, and A. Salazar. A dynamic replacement economy decision model. *IIE Transactions*, 16:65–72, 1984.

[102] I. H. Osman and S. Salhi. Local search strategies for the vehicle fleet mix problem. In V. Rayward-Smith, I. Osman, C. Reeves, and G. Smith, editors, *Modern Heuristic Search Methods*, chapter 8. John Wiley and Sons, Chichester, 1996.

[103] R. Otten and L. van Ginneken. *The Annealing Algorithm*. Kluwer Academic Publishers, Boston, 1989.

[104] C. Papadimitriou and K. Steiglitz. *Combinatorial Optimization: Algorithms and Complexity*. Prentice-Hall, Englewood Cliffs, NJ, 1982.

[105] P. Pardalos, A. Migdalas, and R. Burkard. *Combinatorial and global optimization*. World Scientific Pub., River Edge, NJ, 2001.

[106] C. S. Park and G. P. Sharp-Bette, editors. *Advanced Engineering Economics*. John Wiley and Sons, New York, 1990.

[107] G. Pesant, M. Gendreau, J.-Y. Potvin, and J.-M. Rousseau. An exact constraint logic programming algorithm for the traveling salesman problem with time windows. *Transportation Science*, 32(1):12–29, 1998.

[108] G. Pesant, M. Gendreau, J.-Y. Potvin, and J.-M. Rousseau. On the flexibility of constraint programming models: From single to multiple time windows for the traveling salesman problem. *European Journal of Operational Research*, 117:253–263, 1999.

[109] J.-Y. Potvin and S. Bengio. The vehicle routing problem with time windows part ii: Genetic search. *INFORMS Journal on Computing*, 8(2):165–172, 1996.

[110] J.-Y. Potvin, T. Kervahut, B.-L. Garcia, and J.-M. Rousseau. The vehicle routing problem with time windows part i: Tabu search. *INFORMS Journal on Computing*, 8(2):158–164, 1996.

[111] S. Rajagopalan. Capacity expansion and equipment replacement: A unified approach. *Operations Research*, 46:846–857, 1998.

[112] S. Rajagopalan, M. R. Singh, and T. E. Morton. Capacity expansion and replacement in growing markets with uncertain technological breakthroughs. *Management Science*, 44(1):12–29, 1998.

[113] S. Rajasekaran. On simulated annealing and nested annealing. *Journal of Global Optimization*, 16:43–56, 2000.

[114] C. R. Reeves. Modern heuristic techniques. In V. Rayward-Smith, I. Osman, C. Reeves, and G. Smith, editors, *Modern Heuristic Search Methods*, chapter 1. John Wiley and Sons, Chichester, 1996.

[115] C. R. Reeves and J. E. Beasley. Introduction. In C. R. Reeves, editor, *Modern Heuristic Techniques for Combinatorial Problems*, chapter 1. John Wiley and Sons, New York, 1993.

[116] M. Robini, Y. Bresler, and I. Magnin. On the convergence of Metropolis-type relaxation and annealing with constraints. *Probability in the Engineering and Informational Sciences*, 16:427–452, 2002.

[117] D. F. Rogers, R. D. Plante, R. T. Wong, and J. R. Evans. Aggregation and disaggregation techniques and the methodology in optimization. *Operations Research*, 39(4):553–582, 1991.

[118] Y. Rossier, M. Troyon, and T. Liebling. Probabilistic exchange algorithms and Euclidean travelling salesman problems. *OR Spektrum*, 8:151–164, 1986.

[119] L.-M. Rousseau, M. Gendreau, and G. Pesant. Using constraint-based operators to solve the vehicle routing problem with time windows. *Journal of Heuristics*, 8:43–58, 2002.

[120] R. A. Russell. Hybrid heuristics for the vehicle routing problem with time windows. *Transportation Science*, 29:156–166, 1995.

[121] M. W. P. Savelsbergh. Local search in routing problems with time windows. *Annals Operations Research*, 4:285–305, 1985.

[122] S. Sethi and S. Chand. Planning horizon procedures for machine replacement models. *Management Science*, 25(2):140–151, 1979.

[123] M. M. Solomon. Algorithms for the vehicle routing and scheduling problems with time windows. *Operations Research*, 35:254–265, 1987.

[124] J. Sussams. *Vehicle Replacement: A Management Guide*. Glower, Glower, 1983.

[125] E. Taillard, P. Badeau, M. Gendreau, F. Guertin, and J.-Y. Potvin. A tabu search heuristic for the vehicle routing problem with soft time windows. *Transportation Science*, 31:170–186, 1997.

[126] K. T. Talluri. The four-day aircraft maintenance routing problem. *Transportation Science*, 32(1):43–53, 1998.

[127] K. Tan, L. Lee, Q. Zhu, and K. Ou. Heuristic methods for vehicle routing problem with time windows. *Artificial Intelligence in Engineering*, 15:281–295, 2001.

[128] J. Tang and K. Tang. A note on parallel machine replacement. *Naval Research Logistics*, 40:569–573, 1993.

[129] G. Terborgh. *Dynamic Equipment Policy*. McGraw-Hill, New York, 1949.

[130] S. R. Thangiah, I. H. Osman, R. Vinayagamoorthy, and T. Sun. Algorithms for the vehicle routing problems with time deadlines. *American Journal of Mathematical and Management Sciences*, 13:323–355, 1993.

[131] V. Theodoracatos and J. Grimsley. The optimal packing of arbitrarily-shaped polygons using simulated annealing and polynomial-time cooling schedules. *Computer Methods in Applied Mechanics and Engineering*, 125:53–70, 1995.

[132] B. W. Thomas. *Anticipatory Route Selection Problems*. PhD thesis, University of Michigan, 2002.

[133] A. Vakharia and Y. Chang. A simulated annealing approach to scheduling a manufacturing cell. *Naval Research Logistics*, 37:559–577, 1990.

[134] P. J. M. van Laarhoven and E. H. L. Aarts. *Simulated Annealing*. P. Reidel Publishing Co., Dordecht, Netherlands, 1987.

[135] H. M. Wagner. *Principles of Operations Research*. Prentice-Hall Inc., New York, 2nd edition, 1975.

[136] L. Wiseman. In search of the perfect vehicle replacement program. *Transmission & Distribution*, pages 28–33, May 1993.

[137] J. Yao. On constrained simulation and optimization by Metropolis chains. *Statistics and Probability Letters*, 46:187–193, 2000.

ABSTRACT

THEORY AND APPLICATIONS OF COMPRESSED ANNEALING

by

Jeffrey W. Ohlmann

Co-chairs: James C. Bean & Shane G. Henderson

Operations managers are often faced with large-scale decision-making problems that are difficult to solve to optimality. We analyze compressed annealing, a heuristic approach for solving large-scale combinatorial optimization problems. Compressed annealing is a variant of simulated annealing that integrates a variable penalty method with heuristic search to address optimization problems with relaxed constraints. The concept of pressure is introduced to parameterize the value of the penalty multiplier.

We present a theoretical framework to study the behavior of compressed annealing. We provide necessary and sufficient conditions that ensure the metaheuristic's convergence in probability to the set of global optima. Guided by theoretical insight, we develop practical joint cooling and compression schedules.

We employ compressed annealing on an asset replacement problem considering

the issues of stochastic deterioration, budget limits, and time-variant costs due to technological change. We perform computational experiments on data sets constructed from information provided by trucking companies. Empirical results illustrate the effectiveness of compressed annealing; replacement plans obtained via compressed annealing outperform a trade cycle approach commonly implemented in the trucking industry.

To test compressed annealing's robustness, we apply the algorithm to the traveling salesman problem with time windows (TSPTW). The variable penalty approach of compressed annealing allows a search considering tours infeasible with respect to the time windows. Compressed annealing obtains best-known results on numerous data sets from the literature.