PLANNING ROBUST FREIGHT TRANSPORTATION OPERATIONS

A Thesis
Presented to
The Academic Faculty

by

Juan C. Morales

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
School of Industrial and Systems Engineering

Georgia Institute of Technology
December 2006

PLANNING ROBUST FREIGHT TRANSPORTATION OPERATIONS

Approved by:

Professor Alan L. Erera, Advisor
School of Industrial and Systems Engineering
Georgia Institute of Technology

Professor Martin Savelsbergh, Advisor
School of Industrial and Systems Engineering
Georgia Institute of Technology

Professor Shabbir Ahmed
School of Industrial and Systems Engineering
Georgia Institute of Technology

Professor Mark Ferguson
College of Management
Georgia Institute of Technology

Professor Mark Goetschalckx
School of Industrial and Systems Engineering
Georgia Institute of Technology

Date Approved: 13 November 2006

To Daniel and Olga L.

ACKNOWLEDGEMENTS

I would like to express my sincere gratitude to my advisors, Professor Alan L. Erera and Professor Martin Savelsbergh, for their unconditional support, guidance and understanding, which helped me overcome many personal and professional challenges, making this thesis possible. Special thanks to Professors Shabbir Ahmed, Mark Ferguson and Mark Goetschalckx, for their valuable comments and suggestions. I would also like to express my appreciation for all faculty, staff and fellow students of the H. Milton Stewart School of Industrial and Systems Engineering, for the wonderful moments and experiences I shared with them during this journey. My deepest gratitude to my wife, Olga L. Ortiz, for her trust, consideration and confidence, which always kept alive the light of my life.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER I

INTRODUCTION

Freight transportation operations are quite often complex dynamic processes. Typically, large freight carriers must manage multiple fleets of heterogeneous resources in order to serve customer demands. However, even relatively small transportation service providers may face challenges managing resources. The focus of this dissertation is problems of *fleet management* in freight transportation systems. Fleet management problems may take a variety of forms, and the "fleet" may represent many types of resources. For example, truckload trucking firms manage fleets of drivers, tractors, and trailers to serve point-to-point full truckload demand of customers. Railroads also serve point-to-point full carload traffic, but face additional complexity since railcars are usually transferred between multiple locomotives and operating crews en route from origin to destination. Container fleet operators, such as those that own tank containers for chemical transport, must manage the global transportation (both loaded and empty) of fleets of various container types to again serve point-to-point customer demands, relying on third parties (such as ocean carriers, railroads, and drayage carriers) to actually move the equipment. Consolidation carriers, such as less-than-truckload (LTL) trucking and parcel express companies, usually manage multiple fleets for different transportation roles. For example, an LTL company will utilize fleets of drivers, tractors, and trailers for long-haul services between terminals, as well as for local pickup and delivery services from terminals.

Fleet management in freight transportation leads to difficult problems of resource allocation and scheduling. There are three primary challenges:

1. *Geography*: Unlike many resource allocation and scheduling problems found in other domains, geography is an essential complicating factor of the problems faced in transportation. Of course, freight transportation systems only exist to overcome the geographical separation of points of production and points of consumption. In today's global economy,

production facilities and consumption markets are often separated by vast distances. International trade agreements, combined with an evolving world economy increasingly fueled by emerging regions with low-cost production capability, has led to many freight transportation systems with global operational scope.

In the context of fleet management, vast system geographies complicate matters since it is not only important *when* resources will be available, but also *where* they will be available, or when they might be made available (at some cost) where they are needed. Thus, one of the primary challenges faced by large fleet managers operating systems that serve large geographic spaces is how to schedule the movement of *empty* resources (*i.e.*, those moving without customer freight).

2. *Dynamics*: System dynamics is another fundamental challenge faced by freight transportation managers, although all scheduling problems by nature are problems with an important time component. From one perspective, freight transportation systems are in a constant state of change. A dynamic flow of events, some exogenous and others driven by fleet management decisions, change the state of the system over time. Fleet managers must monitor these exogenous events and make decisions triggering other events to ensure that the system provides an appropriate level of service to customers at a low operating cost.

Since few, if any, freight transportation systems evolve in a completely predictable way, fleet management almost always involves continual updating of decisions. Importantly, the information that is used to make decisions at any point in time (the *information state*) also evolves dynamically. For example, information regarding the characteristics of the demand request of an individual customer may not be available until a few hours prior to the expected commencement of service. In some local pickup or delivery examples, *all* characteristics of a customer's demand (including its physical quantity or size) may not be available until a resource arrives on site to service the demand. Recent advances in communication technology, such as Global Positioning Systems (GPS) or Radio Frequency Identification (RFID), allow real-time visibility of freight operations, thus offering carriers the opportunity to have better information to update decisions.

3. *Uncertainty*: As mentioned above, the dynamic evolution of the state of a freight transportation system is usually not completely predictable. In some cases, the level of uncertainty is minor and can be ignored when making decisions without compromising service or generating excessive cost. In many systems, however, improper or inadequate consideration of uncertainty during decision-making may lead to substantial adverse impacts.

One key source of uncertainty in many freight transportation systems is customer demand. In most systems, customers provide a notice of intent to ship freight in advance, although the "lead time" of such information may be quite short (for example, a few hours). Even when customers book transportation weeks or months in advance, there is still likely to be some uncertainty about precisely what will require execution and when. Thus, information about customer locations, demand quantities, and demand timing may all be uncertain to some degree in advance of execution. The key to effective fleet management is to make decisions with the best information available at the times when decisions must be made, and to appropriately account for the uncertainty of the parameters that comprise this available information state.

1.1 Freight transportation planning and control

Effective fleet management for freight transportation requires both effective planning and control decisions. In this thesis, *planning decisions* will refer to those made in advance of execution. Effective plans should prepare the system appropriately for the conditions to be faced during future operations. However, since the state of many systems does not evolve predictably, planned decisions often require modification. *Control decisions* refer to the decisions made during execution, and represent an implementation of planned decisions, potentially modified to adapt to the current system state.

Plans are important for freight transportation systems for many different reasons. In some examples, it may be important to develop a plan so that resources can be committed or procured in advance, which often reduces cost or provides some guarantee of service level. For example, in the context of global container management, a fleet manager who books space in advance for the transport of loaded and empty containers may receive a lower price than those who make last-minute requests. The existence of plans usually reduces the complexity of control, since managers can more easily coordinate resources. Furthermore, when plans are communicated to shippers (for example, in the form of schedules), they too can coordinate their activities with those of the carriers from which they receive service.

This thesis will focus on developing methodologies for improving fleet management planning decisions. Better planning decisions lead to more *effective* system control, where control is effective when the system can provide appropriate levels of customer service at low cost. At any given time, planning decisions should be made considering the following two key elements:

1. Current system information state, including known and stochastic parameters; and
2. Future control decisions.

It is important to recognize that if plans are to be effective, they should be decided upon using some knowledge of how control decisions will be made in future operating periods. It is often challenging to develop appropriate means for representing the control decision process within a quantitative model for generating planning decisions.

Clearly, system dynamics and uncertainty complicate decision-making for planning. If the system state were only to evolve dynamically but purely deterministically, the challenge of proper representation of control decisions during planning lies only in the potential complexity of operations, which may require that a solution approach utilize appropriate approximations for tractability. Uncertain dynamics, however, leads to additional complications. The three important questions that need to be considered are:

- What control policies are most appropriate?
- How should control policies be modeled during the planning phase?
- How should uncertainty be incorporated into models for planning?

Given a control policy, fleet planning under uncertainty essentially includes resource allocation and scheduling decisions that generate *flexibility* for the control phase. Flexibility in control implies that additional decision options are available that enable service commitments to be met and with low cost; flexibility is thus a *buffer* or *hedge* against uncertainty. If the control phase is not provided enough options, service or cost will be sacrificed under specific outcomes. Since control flexibility almost always comes at cost, planning decisions must balance the cost of planning flexibility with the subsequent effectiveness of control.

1.2 Fleet management planning for effective control

The primary objective of this thesis is to develop a methodological approach to generate fleet management plans that enable effective control of transportation systems with uncertain dynamics. In practice, plans are often generated using estimates of how the system will evolve in the future; during execution, control decisions need to be made to account for differences between actual realizations and estimates. Although such decisions are necessary, they may come with a high cost. The benefits of plans designed to satisfy customer demands at minimum cost can be negated by performing costly adjustments during the operational phase. Stochastic models must be used to develop plans that explicitly address uncertainty.

Stochastic planning models proposed for fleet management have focused primarily on expectation minimization, which typically results in dynamic programming models that are hard to solve for problems of realistic size. In some stochastic planning models, uncertainty is incorporated using random variables for which probability distribution functions must be assumed. To use such models in practice, it is thus necessary to choose a probability distribution and estimate its corresponding parameters for each random variable in the model. This can be an overwhelming task when planning real-life freight transportation operations, and can often limit the applicability of these traditional models. Some reasons for this are listed below:

- Making plans for freight transportation systems often requires the consideration of long planning horizons. This is especially important for systems that handle long haul movements. Long planning horizons might imply fitting a large number of distribution functions.

- Demand for freight transportation typically has a strong seasonal pattern, and this fact can make the process of fitting probability functions more difficult. Assuming independent and identically distributed random variables can be an oversimplifying assumption.

- Modern transportation systems change rapidly in order to adapt to customer needs, as well as to react to fierce competition. Transportation schedules and business practices are altered periodically, and thus it may become difficult to collect a statistically significant sample of observations in order to fit a probability distribution function and estimate its parameters.

In this research we consider alternative approaches for incorporating stochasticity into planning models, inspired by recent work in robust optimization. This research area has been very active in the past few years, and provides the flexibility to explicitly incorporate uncertainty in a manner that may be easier for practitioners to understand and implement within planning models. In the approaches we develop, uncertainty is often characterized in part by using intervals around nominal forecasted values. Such intervals can be calculated using both point forecasts of quantities and the historical errors of such forecasts, as well as by using the knowledge of experienced fleet managers.

Another feature of the approaches we consider is that risk aversion can be addressed relatively simply. In many decision contexts, it is not adequate to assume that decision makers are risk neutral; clearly, different transportation industries and different decision makers within organizations may not show the same propensity toward taking risks. The approaches applied in this thesis are designed to allow planners to manage solution conservatism with relatively simple controls that restrict both individual and joint realizations of the uncertain parameters considered during solution.

Finally, like most stochastic planning approaches, the methods considered herein may limit the set of allowable control decisions during planning. In cases where the likely cost (or disutility) of operational modifications to plans may be difficult to quantify *a priori*, limitations on the control decision space can be used to build cost-effective plans by forcing the planning phase to generate adequate buffers against uncertainty. Thus, costly modifications to plans can be contained, and random fluctuations can be absorbed by these preplanned buffers.

1.3 Dissertation Outline and Contributions

A planning approach that permits effective fleet control during plan execution is proposed in this dissertation. We develop concepts and methodologies that are applied in two fundamental problem areas: (i) dynamic asset management and (ii) vehicle routing. Three specific and self-contained problems are studied in this dissertation. The first two relate to dynamic asset management, while the third addresses vehicle routing problems with stochastic demands.

Chapter 2 presents the first problem to be addressed, intermodal fleet management in the context of tank containers. Tank containers, also referred to as ISO tanks, are used to transport liquid chemicals. In today's global economy where production facilities and consumption markets are separated over vast distances, a reliable and cost effective way to transport goods is of key importance. Intermodal transportation offers these characteristics, and this explains its phenomenal growth during the past years. The Intermodal Association of North America (IANA) reported that ISO container (20-ft, 40-ft, and 45-ft boxes) international traffic moves increased by 10.7% and 10.8% in 2004 and 2005 respectively.

Managing intermodal fleets is hard because global trades are generally not balanced in both space and time, and transport requires a coordinated set of activities to be performed by multiple carriers. Furthermore, transportation plans need to consider long planning horizons to account for lengthy transit times typically found in worldwide freight transportation. In practice, the fleet management planning and control problem is usually decomposed in two parts; the problem of repositioning empty units to account for flow imbalances, and the problem of allocating units to customer demands and handling the corresponding loaded moves. In Chapter 2, an alternative integrated dynamic model for the tank container management problem is formulated as large-scale multicommodity flow problem on a time-discretized network. The model uses both known customer demands and point forecasts of future demands, and serves joint planning and control roles when implemented in a rolling horizon setting. A computational study shows that with state of the art optimization software, empty and loaded movement decisions can be addressed in an integrated model for fleets of reasonable size. Furthermore, the study provides evidence that by making decisions using an integrated approach, fleet operating costs and fleet sizes may be significantly reduced. However, the results also illustrate clearly how a model that does not consider the inherent uncertainty of forecasted quantities will generate plans without buffers. Such plans are fragile, and may lead to situations in which serving realized customer demands is either impossible or very costly.

Addressing these limitations, Chapter 3 proposes a dynamic and stochastic planning approach for empty repositioning problems. Motivated by the importance of systematic consideration of forecast uncertainty, as illustrated in Chapter 2, this chapter builds the theoretical groundwork of a methodological approach for generating robust repositioning plans that permit effective control. Specifically, the chapter proposes a robust optimization framework for dynamic empty repositioning problems modeled using time-space networks. A repositioning plan is said to be robust with respect to a deviation from forecasted nominal values if there exists a set of feasible control decisions drawn from a restricted control space.

The proposed planning approach incorporates demand and supply uncertainty at every period of the planning horizon using intervals around nominal forecasted parameters. The intervals define the uncertainty space for which buffers need to be built into the plan in order to make it a robust plan. No probability distributions for the parameters are assumed. Risk aversion is also considered; the model user can specify the level of conservatism of the resulting plan using input parameters. More conservative plans are robust with respect to a larger set of joint uncertain parameter realizations than less conservative plans.

There is a key difference between a dynamic model and a model being applied dynamically. As pointed out in Powell et al. [45], “a model is dynamic if it incorporates explicitly the interaction of activities over time”. On the other hand, a model is applied dynamically if the “model is solved repeatedly as new information is received”. In our methodological approach, the model is applied dynamically using a rolling horizon framework. A plan is generated over a finite planning horizon and only the decisions of the first period(s) are implemented. As demand realizations are observed and potentially some dynamic adjustments are made, demand and supply forecasts might need to be reevaluated and the system status changes. The plan must then be updated to incorporate the new information.

The main contributions of this thesis to dynamic asset management can be summarized as follows. This thesis:

- develops a robust optimization framework for integer programming problems with equality constraints and right-hand side uncertainty, a common feature of many logistics planning problems. The framework explicitly incorporates the notion of control decisions to dynamically respond to realizations of the uncertain parameters, transforming initial planned solutions into feasible solutions;
- provides implementations of the proposed robust optimization framework in the context of empty repositioning problems faced by many freight transportation service providers for different sets of allowable recovery actions;
- demonstrates that for the sets of control decisions considered for the empty repositioning problem, the size of the resulting optimization problem does not depend on the size of the uncertain outcome space, and that for the simplest set of recovery actions the resulting optimization problem can be solved in polynomial time; and
- presents a computational study illustrating the value of robust optimization in the context of empty repositioning problems and demonstrating the computational viability of the proposed framework.

In Chapter 4, a fleet management setting with a much smaller geographical scope is considered; a problem faced by fleets performing local delivery or pickup operations. Again, demand uncertainty plays an important role because vehicles may need to restock at the depot if total realized demand exceeds the vehicle capacity. The Vehicle Routing Problem with Stochastic Demands (VRPSD) has received considerable attention from the research community. Our work in this area combines basic ideas and concepts found in the literature on robust optimization and stochastic programming with recourse. This rich framework allows the study of control adjustments and dynamic policies, two key components of this research. Our approach focuses on two specific versions of this problem:

1. The Robust Vehicle Routing Problem with Stochastic Demands (RVRPSD). The traditional approach to address the VRPSD is through expectation minimization of the cost function. Although this approach is useful for building fixed routes with low expected cost, it does not directly consider the maximum potential cost that a vehicle might incur when traversing the tour; that is, the cost associated with the *worst-case demand realization.* A computational study included in this dissertation shows that there can be a significant difference between the maximum and average cost of a solution of the VRPSD. Our approach aims at minimizing the maximum cost.

2. The Vehicle Routing Problem with Stochastic Demands and Duration Constraints (VRPSDDC). In many real-life applications, a fleet of vehicles performs local pickups that are then sorted at a regional terminal for the purpose of performing long-haul transporting of goods between multiple terminals in a network. In this context, duration constraints are very important to ensure that all goods are picked-up and sorted before the long-haul vehicle departs the regional terminal. In this section, we study the traditional VRPSD where duration constraints are enforced on a per tour basis.

Customer demand is assumed to take integer values within closed intervals. We show that the problem of identifying the worst-case demand realization for the tour of a single vehicle can be solved as a longest path problem on an acyclic network. Furthermore, we establish that although in general this problem is solvable in pseudo-polynomial time (as a function of the vehicle capacity), for a family of dynamic policies, referred to as history-independent policies, this problem is polynomially solvable.

A great deal of work is devoted to studying recourse policies. We study a broad spectrum of policies along two dimensions:

1. Customer demand visibility: ranging from observed only when the vehicle arrives at the customer's location, to total demand visibility before the vehicle departs from the depot.

2. Information considered by the dynamic policy: ranging from completely myopic, in which dynamic adjustments are based only on observed demand of the currently visited customer and on-board inventory, to anticipatory policies, in which more information, such as travel times, demand intervals, and even demand estimates, are used to decide when a recourse action needs to take place.

Some relations about the performance of these policies are derived and interesting paradoxes are identified regarding the amount of information considered by a policy and its corresponding cost. We show that the intuitive idea that more information leads to better decisions is not always true in the context of the two problems considered in this section.

Finally, a tabu search heuristic is proposed to solve both the RVRPSD and VRPSDDC. Two computational studies are then conducted using this heuristic; the first aims at understanding the differences between tours generated to minimize expected cost and those generated to minimize maximum cost. The second computational study assesses the effects and implications of including duration constraints in the VRPSD.

The results from the first experiment show that the *mean-max gap*, defined as the percentage difference between the maximum and the expected cost of a tour, can be substantial. In our experiments, the average mean-max gap for tours obtained through the traditional expectation minimization approach was 37%, and cases with gaps of over 100% where identified for a simple myopic recourse policy. Such results point clearly to the potential need for alternative approaches that focus on worst-case performance.

Furthermore, the results in this thesis show that the expected cost of mini-max tours is on average only 3% more expensive than the expected cost of tours obtained using expectation minimization. However, when using a myopic recourse policy, the average mean-max gap of these tours is 30%, a 19% reduction of the gaps generated by expectation minimization tours. While such a reduction is significant, the average mean-max gap of 30% is still quite high. Surprisingly, when we consider anticipatory policies with the robust approach, the tours also show good performance in terms of their expected cost, and the average mean-max gap is reduced to less than 15%. Such results point not only to the value of our approach, but also to the importance of considering non-myopic recourse policies for the

VRPSD.

The main contributions of the research on the vehicle routing problem with stochastic demands can be summarized as follows. This thesis:

- develops a robust optimization approach for the VRPSD, and shows via computational results that the approach may lead to substantial reductions in the mean-max gap. The approach is applied to the VRPSD with duration constraints, a problem variant particularly relevant in practice, which, to our knowledge, has not been previously addressed;

- develops analytical results for non-myopic recourse policies and demonstrates their value over myopic policies, which are the basis of almost all published research on the VRPSD; and

- provides implementations of the proposed approach and presents computational studies that illustrate its value, and also provides insights on how some factors (*e.g.*, number of customers, vehicle capacity and demand variability) affect solution cost.

1.4 Literature Review

In this dissertation, we apply some recently developed ideas on robust optimization to the problem of planning and controlling freight transportation operations. Robust optimization is emerging as a viable alternative approach for modeling and solving decision optimization problems given uncertainty. Soyster [49] is the first work to consider coefficient uncertainty in linear programming formulations, and shows that such uncertainty can be handled by an equivalent linear programming model. The approach, however, is very conservative since it protects feasibility against all potential uncertain outcomes. More recently, Ben-Tal and Nemirovski [6] develops a general framework for robust optimization over a convex cone; also, Ben-Tal and Nemirovski [7] specifically considers linear programs with coefficient uncertainty. To control the conservatism of the robust approach, these references propose the use of an ellipsoidal uncertain outcome space that protects feasibility less conservatively by ignoring very unlikely joint outcomes; for linear programming, the resultant robust optimization problem requires solution of a convex optimization problem over a second-order cone. To avoid the necessity of a nonlinear optimization solution method, Bertsimas and Sim [11] considers robust linear programming with coefficient uncertainty using a uncertainty set with budgets. In this model, each coefficient is assumed to take a value in a symmetric interval around a nominal value, and a budget parameter for each constraint limits the number of coefficient values that can simultaneously take their worst-case value. For this characterization of uncertainty, the resulting robust optimization is still a linear optimization problem. Extending this work, Bertsimas et al. [9] develops robust versions of linear programming problems with coefficient uncertainty sets described by an arbitrary norm.

Robust versions of discrete optimization problems have also received attention. Kouvelis and Yu [37] develops robust versions of many traditional discrete optimization problems with two different objective functions, one which minimizes the maximum absolute cost under any potential outcome, and another which minimizes the maximum cost difference (or regret) between the robust solution and a reoptimized solution for each outcome. The reference shows that robust versions of many polynomial discrete optimization problems become NP-hard. Bertsimas and Sim [10] considers robust discrete optimization and network flow problems where parameter uncertainty occurs only in the objective function, and shows that many polynomially-solvable (or polynomially-approximatable) problems remain so in this case.

Closely related to our work, Atamturk and Zhang [3] considers network flow and design problems with right-hand-side uncertainty using a robust approach. Similar to our approach, this paper develops and analyzes a two-stage approach similar in spirit to the adjustable robust counterpart approach for linear programming developed in Ben-Tal et al. [5]. Also related to our research, Bertsimas and Thiele [13] and Bertsimas and Thiele [14] use a robust optimization framework to address traditional inventory control problems for single point and tree-network supply chain networks that are typically addressed with dynamic programming techniques.

Fleet management is a field of study that in the past years has received considerable attention from the research community. Although first papers in this are date back to the 70s, most of them dealt with static and deterministic models. Recent advances on information technology has created a platform that allows companies to handle real-time information about fleet status. At a touch of a button, carriers and shippers can identify with a high level of accuracy the location of their equipment. Other attributes, such as whether it is loaded or empty, clean or under repair are also available. This high level of visibility, together with a higher degree of flexibility to make and communicate decisions, has raised a new challenge for researchers and practitioners: How can this information be used to make better decisions during the planning and execution phases?

Research on fleet management has been very active in the past few years. For excellent surveys on models for freight transportation operations see Crainic [22] and Powell [43]. First fleet management models were static and deterministic. One of the earliest contributions to this area is Misra [40], which proposes a fleet management model in the contexts of railcar distribution. This paper proposes the application of a well known linear programming problem, the transportation problem, to match railcars to customer demands. Although the model captures important aspects of the problem, the approach fails to consider the dynamic and stochastic nature of the problem. It completely ignores the time dimension of fleet management; it just takes a snapshot of the current availability of railcars and of the set of current requests, and matches them in order to minimize the total cost of moving cars in the railroad network. Further, the model assumes deterministic demands. Surprisingly, this type of technology is still used by some railroad companies in North America (see Powell and Topaloglu [46]).

White [54] proposes one of the first dynamic, still deterministic, models in fleet management in the context of repositioning empty containers. This model considers a fixed planning horizon which allows developing empty repositioning plans over several days. The model is a minimum cost flow problem on a time-space network. Recent extensions of this approach also in the context of container management are Abrache et al. [1] and Erera et al. [28], interest is focused on handling large fleets of containers.

Jordan and Turnquist [36] opened a line of research that guided most of the subsequent work on fleet management. They propose a stochastic and dynamic approach to model the distribution of empty freight cars in the railroad industry. This is one of the earliest contributions in the area of stochastic dynamic optimization in freight transportation. In this model the objective is to maximize the expected profit over a finite planning horizon. The objective function considers the cost of not being able to satisfy demands that randomly show-up in various part of the railroad network, the cost of holding empty cars, as well as the cost of repositioning empty units across the network. The result is a dynamic programming model that is hard to solve for realistic instances. A heuristic solution is proposed based on Frank-Wolfe algorithm, which iteratively makes linear approximations to the the nonlinear function solving linear programming subproblems.

An important contribution in the area of stochastic and dynamic models for fleet managements is Powell [41]. A time-space network, where candidate loaded and empty movements are considered, is used to track and make decisions on equipment movements over a finite planing horizon. Features of this model are motivated by the truck allocation problem for truckload carriers. Equipment requests for different OD pairs in the network is uncertain, and each unit is assumed to be able to handle several loads during the planning horizon under consideration. The model is solved using Frank-Wolfe algorithm because the linearized objective function allows the solution of simple subproblems at each node in the network. Computational results suggest an acceptable computational efficiency. Further, these results illustrate the benefit of the stochastic approach over a deterministic approach.

Powell [42] is the first reference that aside from considering the stochastic and dynamic nature of transportation operations, further explicitly considers demand forecasts in the model. This model combines an assignment model that assigns specific drivers to specific loads in the first time periods of the planning horizon, and a dynamic network that accounts for forecasted demands. The resulting model, called hybrid model, is computationally intractable; therefore, techniques to approximate the recourse function, like scenario aggregation, stochastic gradient methods, successive linear and convex approximation procedures among other are considered. Numerical experiments show that the hybrid model outperforms standard myopic models.

In general, work with stochastic models has focused primarily on expectation minimization. Modelling approaches for stochastic problems, focusing specifically on empty container management, are provided in Crainic et al. [25]. Most computational work in this area has focused on suboptimal solution procedures, developing results for improving value function approximations in iterative dynamic programming solution methods (see, *e.g.*, Frantzeskakis and Powell [29], Cheung and Powell [17], and Powell and Carvalho [44]). More recently, adaptive approaches to approximating nonlinear value functions have been successfully applied to both single commodity and multicommodity problems (see, *e.g.*, Godfrey et al. [33], Godfrey and Powell [34], and Topaloglu and Powell [53]). Topaloglu and Powell [52] provides theoretical justification for this approach, proving the optimality of a particular variant of the general sampling method developed first in Godfrey et al. [33].

The vehicle routing problem with stochastic demands (VRPSD) is a well studied problem. For excellent surveys on this topic see Dror et al. [26] and Gendreau et al. [31]. Most efforts have addressed this problem using three different methodological approaches:

1. Chance constrained models.

2. Stochastic programming with recourse models.

3. Markov decision models.

One of the first chance constrained models is proposed in Stewart and Golden [50], which in its objective function contains the cost of the tours (no expected cost is considered). This model aims at identifying minimum cost tours subject to a threshold constraint on the probability of a route failure. A similar approach is found in Laporte et al. [39], which proposes a model that uses fewer variables but requiters a homogeneous fleet of vehicles. In both cases, the model can be transformed into a deterministic vehicle routing problem. The major shortfall in these models is that, although they control the probability of a failure, the location of the failure is ignored and hence does not take the corresponding costs into account. Tours with the same cost and same failure probability can have quite different expected costs, depending on the possible failure locations.

Stochastic programming with recourse models minimize the cost of the tours plus the expected cost of recourse actions that must be taken when failures occur. Most of the literature on VRPSD follows this approach. Dror and Trudeau [27] propose a model that takes into consideration the location of a failure, where recourse actions take the form of back-and-forth trips to the depot. This type of model has a more meaningful objective function than their chance constrained counterparts, but also becomes harder to solve.

Significant efforts have been devoted to the development of heuristic algorithms for solving recourse models. Stewart and Golden [50] as well as Dror and Trudeau [27] propose algorithms inspired on the ideas of the savings heuristic of Clarke and Wright [19] for deterministic VRPs. An efficient local search heuristic approach for solving this type of problems is presented in Savelsbergh and Goetschalckx [47], a numerical study shows that this approach compares favorably with respect to the other two algorithms. An important contribution is Gendreau et al. [32], which proposes a local search algorithm embedded in tabu search approach. The benefit of this approach is assessed by comparing results obtained with the heuristic algorithm to optimal solutions. The heuristic provided the optimal solution in 89.45% of the cases. Further, only an average deviation from optimal cost of 0.38% was found.

Solving the VRPSD to optimality is very hard. First exact approach that we are aware of is Gendreau et al. [30], where an integer L-shaped method capable of solving small instances is proposed. Laporte et al. [38] presents an improved integer L-shape approach capable of handling larger instances. An important element of their approach is the use of lower bounds at the root node which helps speedup solution times. These bounds are calculated under the assumption that the expected value of demand on any tour is less than or equal to the vehicle capacity.

It is worthwhile to mention that most of the research on stochastic programming with recourse around the VRPSD focuses on simple myopic recourse policies (see for instance Bertsimas [8] and Bertsimas and Simchi-Levi [12]). Markov decision models offer the potential for making optimal decisions every time new information is revealed; the drawback is that these models involve a very large number of states, making it intractable even for modest size instances. Dror et al. [26] considers a single vehicle model where a decision epoch corresponds to the moment the vehicle arrives at a customer location and its demand is realized. At that point two possible decisions can be made prior to the start of the service: (i) not to serve the customer and then move to another location, or (ii) serve the customer and then move to another location. No solution approach or computational study is presented in this paper.

An interesting approach, based on Markov decision models, is presented in Yang et al. [55]. Instead of considering the myopic recourse action of returning to the depot whenever the vehicle runs out of inventory, this paper aims at developing an optimal restocking policy in conjunction with the routing decisions. Under this policy the vehicle might restock at the depot before a stockout actually occurs. For a given tour, they show that the optimal restocking policy has a simple form: after serving a customer if on-board inventory drops below a threshold then take a recourse action (*i.e.*, restock at the depot), else continue to the next customer in the tour. The threshold might be different for each customer. Although the optimal policy is quite simple, solving a model that further considers routing decisions is difficult, so heuristics approaches are considered. Two heuristic algorithms are proposed, one falls under the category of route-first-cluster-second, while the other is a cluster-first-route-second algorithm. A computational study is used to benchmark the results of these two algorithms, and to show the superiority of their approach over deterministic methods with simple recourse actions. Unfortunately, the paper fails to compare the results of their approach to those obtained using algorithms that assume stochastic demands with myopic recourse actions.

CHAPTER II

GLOBAL INTERMODAL TANK CONTAINER MANAGEMENT

The chemical industry is growing steadily, especially in China, where chemical consumption is growing at a rate of 8% annually (China is projected to become the third largest consumer by 2010). The value of global chemical production exceeded US$1.7 trillion in 2003. World trade in chemicals continues to surge as well. In 2002, chemicals led all product groups in global trade growth at over 10%, with total world export value reaching US$660 billion. As a consequence, transport of chemicals represents a significant portion of worldwide transport of goods.

Long-distance, international transportation of liquid chemicals is conducted using one of five modes: pipeline, bulk tankers, parcel tankers, tank containers, or drums. Pipeline and bulk tankers are used primarily in the petrochemical industry for the transport of large quantities of a single product. Parcel tankers are smaller vessels with up to 42 tank compartments and are used to simultaneously transport multiple cargoes. Tank containers, also referred to as ISO tanks, intermodal tanks, or IMO portable tanks, are designed for intermodal transportation by road, rail, and ship. Tank containers have many advantages for the international transport of liquid chemicals:

- They are environment-friendly, since they are less prone to spillage during filling and unloading, as well as leakage during transportation.
- They permit a higher payload when compared to drums stowed in dry containers (43% more volume).
- They can be handled mechanically, which results in cost savings, but also ensures safety when handling hazardous commodities.
- They provide secure door-to-door multi-modal transportation (by road, rail, sea or inland waterways), and do not require specialized port-side infrastructure.

- They are safe and durable, with a design life of 20-30 years.
- They can be cleaned and placed into alternate commodity service with minimum downtime.
- They can be used as temporary storage for customers with limited space or high-cost permanent storage.

A tank container operator manages a fleet of tanks to transport liquid cargo for a variety of customers between essentially any two points in the world. Typically, 60% to 70% of the fleet is owned by the operator; the remaining tanks are leased, usually for periods of 5 to 10 years. To serve a standard customer order, a tank container operator would provide a tank (or multiple tanks) at the customer's origin plant and arrange transportation for the tank across multiple modes to the destination plant. Transportation will usually include a truck leg at origin and destination and a steamship leg between a port near the origin to another port close to the destination. It may also include rail or barge legs at each end. Operators use depots for temporary storage, cleaning, and repair of empty containers.

In this chapter, we consider the management problems faced by tank container operators. Specifically, we are interested in the difficult task of cost-effectively managing a fleet of tank containers, given imbalanced global trade flows. Given the high cost of tanks, high loaded container utilization is very important in this industry.

2.1 Tank container management

Tank container operators do not typically own or manage any of the underlying transportation services used to move a container from origin to destination. Instead, they enter into contracts for transportation service with a number of providers. Tank container operators maintain contracts with trucking companies, railroads, and port drayage companies for inland transportation, and with container steamship lines and nonvessel operating common carriers (NVOCCs) for port-to-port ocean transportation. Each transport service contract specifies transport legs that are available to the operator, and their costs. Operators combine these legs into itineraries to provide origin-to-destination service for customers.

The typical service offering provided by a tank operator to a shipper customer is a one-way trip. To obtain service, customers first place a request for a price quote for a given origin-destination pair, and then subsequently make a booking or multiple bookings under the quote. We will call these steps the *quotation* and *booking* processes respectively. With the exception of certain ancillary charges, the tank container operator charges the customer a fixed price for transportation, and pays the transportation service providers directly out of this fee. Therefore, it is in the operator's interest to minimize the transportation costs for most shipments. However, there must be a level of "reasonableness" in transit times, and some customers may be willing to pay a premium for faster service.

To develop price quotes for customers, container operators currently rely on a port-to-port methodology for developing and pricing itineraries. In this process, the operator associates both the origin and destination customer locations with an appropriate export port and import port. Then, using a database of available ocean carrier contracts and the scheduled sailings for each carrier, at most two or three potential ocean carrier services between these ports are selected. Each such service forms the basis for an itinerary. The price of each itinerary is then determined by adding an inland transportation cost (if necessary), a profit component, and possibly an overhead cost allocation (for example, to account for asset repositioning costs) to the ocean service cost. The itinerary transit time is computed by adding inland transport times if necessary and schedule delays to the transit time of the scheduled ocean service. These identified itineraries now likely provide different combinations of price and transit time. Typically, the low-price itinerary is first presented to the customer, and if the transit time is unacceptable a higher-price shorter-duration option is presented. Once the customer selects an itinerary, the quote is formalized.

Price quotes to customers are usually valid anywhere from 30 to 90 days (and sometimes longer), and one or many bookings may be made over the duration of the quote. When booking, the customer specifies the number and type of tank containers needed, and service time windows at the origin and destination. Minimally, the time window information will include the earliest time containers may be loaded at the origin, and the latest time containers should be delivered to the destination. Given these requirements, the tank operator must verify that the quoted itinerary is feasible. Since ocean sailings are scheduled and service is not provided each day, the sailing used to generate the quote may or may not allow a feasible routing satisfying the time windows. The operator must also verify that the quoted sailings have available space. If all components of the quoted service offering are not available, the operator must determine an alternative best itinerary that meets customer requirements.

Empty repositioning is a critical component of tank container management. Since fleet operators provide global service and loaded flow patterns are not balanced geographically, some regions tend to be net sources of empty tanks and others net sinks. Additionally, loaded flow demands exhibit seasonal patterns. Operators correct geographic and temporal imbalances in container supply and demand by repositioning empty containers between depots.

Some tank container operators have recently begun using decision support tools based on mathematical programming for dynamic operational planning of reposition moves. One method that we are aware of determines weekly repositioning moves using a deterministic multi-commodity network flow model to minimize empty move cost given forecasts of loaded arrivals and departures in each port area. Such models typically use a planning horizon of several months discretized into weeks, and are solved each week with only the first week's decisions implemented in the standard *rolling-horizon* approach. Empty container repositioning has received a fair amount of research attention since it is an integral part of many freight transportation problems (see, *e.g.*, [23, 48, 16, 18]). Repositioning decisions have also been treated directly in large-scale tactical planning models. Recent work by Bourbeau et al. [15], for example, develops parallel solution techniques for large-scale static container network design problems that explicitly consider repositioning decisions.

Many opportunities exist to provide improved operational decision support technology to tank container operators. However, we believe that the most significant opportunity lies in integrating container booking and routing decisions with repositioning decisions. When a tank container is booked, an appropriate empty is assigned to the load, moved from its depot to the customer, loaded and transported via multiple modes to the destination meeting customer requirements, moved back to a depot for cleaning, and then stored or repositioned for future use. An operator making these decisions centrally for a global system via an integrated management model may indeed be able to reduce costs and improve equipment utilization. Such an integrated approach would differ substantially from current practice. Although optimization models are used for repositioning, the inputs to current models are forecasts of weekly loaded flow imbalances at depots; thus, container routings that may improve flow balance without repositioning are ignored. This approach to container management, *i.e.*, breaking up the problem into an empty container allocation phase and a container booking and routing phase, resembles the one proposed about a decade ago by Crainic et al. [23]. They argue that "one would like to develop a single mathematical model to optimize short-term land operations of the company to fully account for the interaction between the various decisions to be made, but given the intrinsic complexity of the problem at-hand and the current state-of-the-art in OR, this is not feasible." We will show that even though such integrated models will indeed be substantially larger than existing repositioning models, recent advances in linear and integer programming technology and in computer hardware technology are able to absorb the increased computational requirements.

It is interesting to note that in a recent article discussing the successful implementation of an optimization-based decision support system for the operational management of Danzas Euronet, the typical decomposition approach, in which repositioning decisions are considered separately from customer order planning, is followed (Jansen et al. [35]). One reason for choosing a decomposition approach may be due to the fact that Danzas manages 2,000-6,000 customer container orders daily, and would like to handle a planning horizon of up to three days for order planning decisions.

2.2 An integrated container management model

This section describes an integrated fleet planning and control model for tank container operators providing decision support for routing booked containers and repositioning empties. For simplicity of exposition, we assume that the operator manages a homogeneous fleet of containers. The problem is formulated as a deterministic multicommodity network flow over a time-expanded network. This type of formulation is often used to model freight transportation problems (see Crainic and Laporte [24] and Crainic [21]).

2.2.1 Problem Definition

Consider an operator managing a global tank fleet. Let D represent the set of container depots used to store and clean containers, and let P be the set of seaports through which the operator maintains ocean transportation service contracts.

Suppose that the operator is making decisions at time $t = 0$ by planning movements for a fixed horizon $[0, t_{\max}]$. In a rolling horizon approach, container movement decisions in some subinterval $[0, t_R]$, $t_R \leq t_{\max}$ are implemented. Let T be the ordered set of time periods to be considered, $T = \{0, ..., t_{\max}\}$.

At time $t = 0$, a set of containers are empty, clean, and available (ECA) for assignment at depot locations. Other containers are in transit and will become ECA at specific depots at future time periods. Additionally, new purchased or leased containers may come into the system at depots at future times, and existing containers may leave the system in the future due to retires or lease returns. If we assume homogeneity, container availabilities can be summarized by net ECA container inflows at depots at specific points in time. For each depot $d \in D$ and time $t \in T$, let b_{dt} be the net ECA container inflow, where a negative number indicates a net outflow. Of course, this modeling convention does not track specific individual container numbers which may be important if modeling lease returns.

During T, the operator must must serve a set of origin-destination customer demands Δ which we assume to be known with certainty. Let L be the set of all customer locations (origin and destination). Each demand $\delta \in \Delta$ may represent booked containers or forecasted future container bookings, and minimally has the following characteristics:

$o_\delta \in L$	origin customer location
$d_\delta \in L$	destination customer location
n_δ	number of containers to be transported
$e_\delta \in T$	earliest time empty containers can be delivered to origin
$l_\delta \in T$	latest time loaded containers can begin unloading at destination
τ_δ^L	time required to load all containers at origin
τ_δ^U	time required to unload all containers at destination

This representation simplifies the customer time windows. If more detailed information such as the latest time loaded containers can depart the origin customer or the earliest time loaded containers can be delivered to the destination were available, it could be used to restrict the problem further.

Operators contract with providers for two types of transportation service. *Scheduled service* is provided by ocean carriers between seaport pairs in $P \times P$. Let Φ be the set of scheduled service contracts, where $\phi \in \Phi$ minimally has the following attributes:

$o_\phi \in P$	origin port of the service
$d_\phi \in P$	destination port
τ_ϕ	total sailing time
$T_\phi^D \subseteq T$	time periods at which service departs origin port
$T_\phi^A \subseteq T$	time periods at which service arrives destination port
c_ϕ	cost per container

Sets T_ϕ^D and T_ϕ^A describe the schedule of any service over the planning horizon. For now, we ignore service commitments or capacities.

Unscheduled services are available every time period. Unscheduled service may represent local or long-haul trucking, rail service, and barge feeder service. Let Θ be the set of unscheduled service contracts, where $\theta \in \Theta$ minimally has the following attributes:

$o_\theta \in D \cup L \cup P$	origin depot, customer, or port
$d_\theta \in D \cup L \cup P$	destination depot, customer, or port
τ_θ	total transport time
c_θ	cost per container

Containers may be repositioned empty between all depot pairs using both scheduled services Φ and unscheduled services Θ. Suppose that all depot-to-depot repositioning options are specified by set Γ. Each $\gamma \in \Gamma$ specifies:

$o_\gamma \in D$	origin depot
$d_\gamma \in D$	destination depot
$T_\gamma^D \subseteq T$	time periods at which option departs origin depot
$T_\gamma^A \subseteq T$	time periods at which option arrives destination depot
τ_γ	total travel time
c_γ	total cost per container

We assume that each $\gamma \in \Gamma$ is generated by transportation service contracts. When repositioning for a depot-depot pair can be conducted directly via an unscheduled service $\theta \in \Theta$, we create a repositioning option which departs the origin depot at all time periods. When repositioning between depots i and j uses a scheduled service $\phi \in \Phi$ (*i.e.*, needs to be routed through a seaport), we create an option combining the costs and times of an unscheduled service to connect i to o_ϕ, the scheduled service connecting o_ϕ to d_ϕ, and then an unscheduled service to connect d_ϕ to j. Such composite options depart the origin depot at time periods appropriate for the container to meet the scheduled service sailings without delay.

Finally, we consider important container processing steps at depots and customers. It is assumed that customers return containers to the depot that is closest in distance to the customer location, and that this transportation requires only unscheduled transportation service. After a container is returned to a depot, it must undergo cleaning before they can be stored or reused. Let τ_d^C represent the duration of the cleaning process at depot $d \in D$, which may vary due to available equipment; in fact, some depots must send containers out for cleaning. In this initial problem, however, we assume that depot cleaning times are independent of the previous commodity transported. Let c_d represent the per container cleaning cost at d. Furthermore, suppose there are no capacity restrictions for cleaning or storage at the depots.

At customer locations, arriving empty containers must immediately begin loading and must depart once loaded. On the other end of the move, suppose that loaded containers arriving at a destination customer must immediately begin unloading, and once unloaded must be immediately dispatched to a depot.

Given this description, the goal of the tank container operator is to route loaded containers and reposition empty containers to serve all customer demands according to time window requirements, while minimizing total transportation and depot costs incurred.

2.2.2 Multi-commodity network flow model

The optimization problem outlined in the previous section can be formulated as a deterministic multi-commodity network flow on a time-expanded network. For simplicity of exposition, suppose that T is a uniform discretization, where each time period represents a day: $T = \{0, 1, 2, ..., t_{\max}\}$.

Decision variables Each variable represents integer flows of containers, both loaded and empty, through different stages of routing.

Container flow variables

$x^{\delta}_{\phi}(t) =$ containers of demand δ transported by scheduled option ϕ at time t

$y^{\delta}_{\theta}(t) =$ containers transported by unscheduled option θ at time t associated with demand δ, either empty or loaded

$u^{\delta}(t) =$ containers of demand δ beginning unloading at time t

$v^{\delta}(t) =$ containers of demand δ beginning loading at time t

$z_{\gamma}(t) =$ containers transported with repositioning option γ at time t

$s_d(t) =$ inventory of ECA containers carried at depot d from time t to $t + 1$

$w_d(t) =$ containers beginning cleaning at depot d at time t

Of course, each of these flow variables need not be present at each time period of the model. For example, we can exclude loaded flows x^{δ}_{ϕ} and y^{δ}_{θ} for time periods $t < e_{\delta} + \tau^{L}_{\delta}$ and $t > l_{\delta}$.

Instance-specific preprocessing may be able to further tighten the time windows for which certain flow variables require definition.

The model description to follow does not prevent a customer demand of more than one container to be split as it is routed through the network, as long as each container in the demand reaches the destination by the customer deadline. If demands are unsplittable in the sense that they should be loaded and unloaded simultaneously and use the same transportation options, an alternative formulation is required.

Constraints The decision variables must satisfy constraints of three general types: (a) flow balance, (b) time-window demand satisfaction, and (c) integrality and non-negativity.

Port nodes

In the model, each port $p \in P$ only explicitly handles loaded container flows, strictly assigned to some demand δ. (Empty repositioning flows between depots that require the use of scheduled services through ports are handled independently using set Γ, as explained in the previous section.) We require flow balance constraints for both containers to be loaded onto scheduled services (export), and containers to be unloaded from scheduled services (import):

Port export flow balance

$$\sum_{\phi \in \Phi | o_\phi = p,\ t \in T_\phi^D} x_\phi^\delta(t) - \sum_{\theta \in \Theta | o_\theta = o_\delta,\ d_\theta = p} y_\theta^\delta(t - \tau_\theta) = 0 \qquad \forall\, p \in P,\ t,\ \delta \tag{1}$$

Port import flow balance

$$\sum_{\theta \in \Theta | o_\theta = p,\ d_\theta = d_\delta} y_\theta^\delta(t) - \sum_{\phi \in \Phi | d_\phi = p,\ t \in T_\phi^A} x_\phi^\delta(t - \tau_\phi) = 0 \qquad \forall\, p \in P,\ t,\ \delta \tag{2}$$

Depot nodes

Each depot $d \in D$ is treated as two nodes at each point in time. One node allows for the consolidation of arrivals of dirty empty containers from destination customer locations, while the other models flow balance of ECA containers. Thus, we need two sets of flow balance constraints for each depot at each point in time.

Depot dirty flow balance

$$w_d(t) \quad - \sum_{\theta\in\Theta|d_\theta=d} \quad \sum_{\delta\in\Delta|d_\delta=o_\theta} y_\theta^\delta(t-\tau_\theta) = 0 \qquad \forall\, d \in D,\ t \tag{3}$$

Depot ECA flow balance

$$\sum_{\theta\in\Theta|o_\theta=d} \quad \sum_{\delta\in\Delta|o_\delta=d_\theta} y_\theta^\delta(t) \quad + \sum_{\gamma\in\Gamma|o_\gamma=d,\ t\in T_\gamma^D} z_\gamma(t) + s_d(t) - \tag{4}$$

$$\sum_{\gamma\in\Gamma|d_\gamma=d,\ t\in T_\gamma^A} z_\gamma(t-\tau_\gamma) \quad - \quad w_d(t-\tau_d^C) - s_d(t-1) = b_{dt} \qquad \forall\, d \in D,\ t \tag{5}$$

Customer locations

Customer locations $\ell \in L$ during the planning horizon may be origins of freight, destinations of freight, or both. In the absence of capacity restrictions on loading and unloading demands, constraints at customers can be separated by demand. Constraints (6) and (7) model loading operations, where arriving empty containers begin loading and upon loading completion are immediately dispatched:

Demand loading flow balance

$$v^\delta(t) - \sum_{\theta\in\Theta|d_\theta=o_\delta} y_\theta^\delta(t-\tau_\theta) = 0 \quad \forall\, \delta \in \Delta, \quad e_\delta \le t \le l_\delta \tag{6}$$

Demand outbound flow balance

$$\sum_{\theta\in\Theta|o_\theta=o_\delta} y_\theta^\delta(t) - v^\delta(t-\tau_\delta^L) = 0 \quad \forall\, \delta \in \Delta,\ e_\delta + \tau_\delta^L \le t \le l_\delta \tag{7}$$

Similarly, constraints (8) and (9) model the arrival of loaded containers to their destination. These tanks are immediately unloaded, and then dispatched dirty to a container depot:

Demand inbound flow balance

$$u^\delta(t) - \sum_{\theta\in\Theta|d_\theta=d_\delta} y_\theta^\delta(t-\tau_\theta) = 0 \quad \forall\, \delta \in \Delta,\ e_\delta + \tau_\delta^L \le t \le l_\delta \tag{8}$$

Demand unloading flow balance

$$\sum_{\theta\in\Theta|o_\theta=d_\delta} y_\theta^\delta(t) - u^\delta(t-\tau_\delta^U) = 0 \qquad \forall\, \delta \in \Delta,\ e_\delta + \tau_\delta^U \le t \le l_\delta + \tau_\delta^U \tag{9}$$

In addition to flow balance, we must also ensure that all containers in each demand are loaded and unloaded during the appropriate customer time window. The following constraints are used:

Loading demand time windows

$$\sum_{e_\delta \le t \le l_\delta} v^\delta(t) = n_\delta \qquad \forall\, \delta \in \Delta \tag{10}$$

Unloading demand time windows

$$\sum_{e_\delta \le t \le l_\delta} u^\delta(t) = n_\delta \qquad \forall\, \delta \in \Delta \tag{11}$$

We note that constraints (10) and (11) introduce difficulty into this problem model, since they bundle and constrain flow for multiple demands.

Integrality and non-negativity

$$x_\phi^\delta(t), y_\theta^\delta(t), u^\delta(t), v^\delta(t), z_\gamma(t), s_d(t), w_d(t) \ge 0 \quad \text{and integer} \tag{12}$$

Objective Function The objective is to minimize the total cost of all empty and loaded transportation, and depot costs:

$$\min EC + LC + DC \tag{13}$$

where the empty cost is given by

$$EC = \sum_t \sum_{\delta \in \Delta} \sum_{\theta \in \Theta | (d_\theta = o_\delta) || (o_\theta = d_\delta)} c_\theta \; y_\theta^\delta(t) + \sum_t \sum_{\gamma \in \Gamma} c_\gamma \; z_\gamma(t),$$

the loaded cost by

$$LC = \sum_t \sum_{\delta \in \Delta} \sum_{\phi \in \Phi} c_\phi \; x_\phi^\delta(t) + \sum_t \sum_{\delta \in \Delta} \sum_{\theta \in \Theta | (o_\theta = o_\delta) || (d_\theta = d_\delta)} c_\theta \; y_\theta^\delta(t)$$

and the depot cost by

$$DC = \sum_t \sum_{d \in D} c_d \; w_d(t).$$

Each of the sums above is assumed to only include time periods $t \in T$ during which the corresponding variable is defined.

2.3 A Computational Study

As mentioned earlier, we believe that substantial benefits can be derived from integrating container booking and routing decisions with repositioning decisions. We have conducted a small computational study to (1) demonstrate that it is now computationally feasible to solve realistic instances of an integrated model using commercially available integer programming software on a high-end personal computer, and (2) assess the magnitude of the benefits of an integrated model in terms of reduced repositioning costs and increased asset utilization.

2.3.1 Instance Generation

We consider container management problems with a planning horizon of six months, discretized in days. To create representative global problems, we build instances around 10 ports of global importance: Singapore, Hong Kong, Shanghai, Kobe, Hamburg, Rotterdam, Southampton, Seattle, Los Angeles, and Savannah. Each port is paired with a nearby container depot. The ports are grouped together to form six regions as described in Table 1.

Table 1: Region characteristics for the computational study.

Region	Ports	# customer locations
Singapore	Singapore	5
Asia	Hong Kong, Shanghai	30
Japan	Kobe	5
UK	Southampton	5
Continental Europe	Hamburg, Rotterdam	25
America	Seattle, Los Angeles, Savannah	30

Fixed customer locations were generated within each region according to a uniform distribution on a rectangular geographical zone representing the region (see Figure 1 for a depiction of the instance geography).

Transportation between locations in the same region is assumed to be land-based and therefore available whenever desired. Transportation between locations in different regions involves at least one sea leg and therefore depends on available sailings between ports. To determine a representative set of sailings between the ports, we use a published schedule

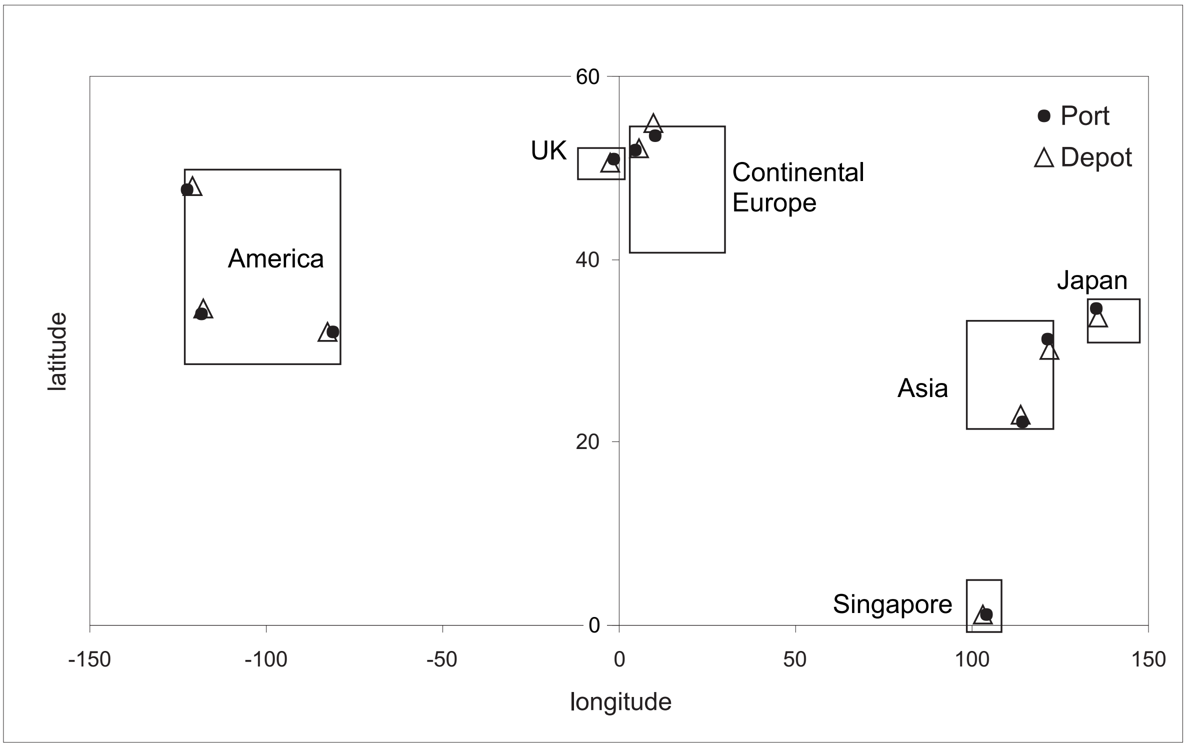

Figure 1: Geographic regions for the computational instances.

(available online) for a large ocean carrier. For the six-month planning period, over 8,000 port-to-port sailings are considered.

A set of 900 customer booking demands is generated for the six-month horizon, with each booking requesting between 1 and 5 containers. To reflect geographic trade imbalances, the origin region and destination region for each booking are selected according to the probabilities given in Table 2.

Table 2: Distributions of origin-destination customer locations for computational study.

Region	Probability of demand originating at region	Conditional probability destination Region 1	Conditional probability destination Region 2	Conditional probability destination Region 3	Conditional probability destination Region 4	Conditional probability destination Region 5	Conditional probability destination Region 6
1	0.15	0	0.10	0.05	0.10	0.25	0.50
2	0.40	0.10	0	0.05	0.10	0.25	0.50
3	0.10	0.05	0.10	0	0.10	0.25	0.50
4	0.05	0.15	0.30	0.15	0	0.30	0.10
5	0.15	0.15	0.30	0.15	0.30	0	0.10
6	0.15	0.15	0.40	0.10	0.10	0.25	0

Within the origin and destination regions, the actual origin and destination customer locations are chosen from the set of available locations randomly with equal probability. For this study, no demands are generated with origin and destination locations in the same region.

For each booking demand δ, the earliest available day e_δ is generated using a discrete uniform distribution over the planning horizon days. To capture the possibility that different bookings may have varying travel time requirements, we determine the latest delivery day l_δ for booking δ as follows. Let $\overline{\tau}_\delta$ be the average transportation time (in days) from the origin customer to the destination customer, where the average is computed over all itineraries in the planning horizon. Then,

$$l_\delta = e_\delta + \lceil \overline{\tau}_\delta \rceil + DiscUniform(\lfloor -0.2 \cdot \overline{\tau}_\delta \rfloor, \lceil 0.2 \cdot \overline{\tau}_\delta \rceil).$$

We further ensure that at least one feasible transport itinerary connects the origin customer to the destination customer by simply discarding any booking demand generated during the construction of the instance for which no feasible transport itinerary exist.

The final component of an instance is the number of available tank containers $\hat{b}$. For simplicity, we assume that the fleet is initially distributed over the depots according to the demand destination probabilities implied by Table 2. We consider three fleet scenarios in the computational experiments that follow. The first scenario models an excess capacity environment with $\hat{b} = 1000$; the system contains many more containers than necessary to handle all demand. The second scenario models an adequate capacity environment with $\hat{b} = 800$. Finally, the third scenario represents a tight capacity environment with $\hat{b} = 600$. In the tight environment, there are certain times during the planning horizon when nearly all containers are in use (either traveling loaded with a demand, repositioning, or traveling empty between a customer and a depot).

2.3.2 Repositioning Strategies

We model several repositioning strategies in this study, including a base strategy designed to emulate current state-of-the-practice in tank container management, and three strategies representing different options for deploying an integrated model.

In the base strategy, we consider an environment in which a centralized decision maker determines repositioning decisions independently from routing decisions. We assume as in practice that these decisions are made once per week using estimates of weekly inflows and outflows of containers at depots. We call this the *2-phase* strategy. To emulate this strategy with our model for a given instance, we consider each booking demand δ and first label the closest container depot to the origin customer location as the origin depot, and the closest depot to the destination customer as the destination depot. Next, we consider all transportation itineraries (combinations of unscheduled and scheduled services) that can be used to feasibly cover δ within the time window defined by e_δ and l_δ, and for each itinerary, determine the day that containers would depart the origin depot and be returned to the destination depot. We then assume that containers will depart the origin depot on its modal day, and be returned to the destination depot on its modal day. Aggregating across all bookings, we compute for each depot the expected weekly inflow and outflow of empty containers. These counts are converted to weekly external container inflows and outflows at depots by assuming that the inflow occurs on the final day of the week, and that the outflow occurs on the first day of the week. Next, the integrated model is solved without any demands to determine the weekly repositioning decisions. Finally, the external container inflows and outflows at depots are removed, the repositioning decisions are fixed, and the integrated model is solved with the booking demands to determine the routing decisions.

The three alternative strategies for deploying the integrated container management model also assume a centralized decision-maker, but determine routing and repositioning decisions simultaneously. The three strategies impose different restrictions on the availability of repositioning options, as follows:

- *Weekly Repositioning (WR)*: Depot-to-depot repositioning is allowed only on the first day of each week. Further, for a given pair of depots, only the minimum cost repositioning option γ is included in the model.

- *Bounded Daily Repositioning (BDR)*: Depot-to-depot repositioning is allowed every day, and all repositioning options are included in the model. However, if a repositioning move is initiated on a given day, a lower bound on the number of containers repositioned is imposed to avoid repeatedly sending small numbers of containers.

- *Unbounded Daily Repositioning (UDR)*: Depot-to-depot repositioning is allowed every day with all repositioning options. No lower bounds are imposed.

2.3.3 Computational Results

The integrated container management model was implemented using ILOG's OPL Studio 3.6.1 and instances were solved using ILOG's CPLEX 8.1. All computational experiments were conducted on a PC with a 1.6 GHz processor and 1 Gb of memory. Given a repositioning strategy and a container fleet size, we solved many instances representing different realizations of customer request demand over the six-month planning horizon. Since the variations between the results was relatively small for problems with the same characteristics, we have chosen to present results for a representative six-month demand realization generated with a particular seed rather than averaging results over many instances.

First, we compare the repositioning costs determined under the 2-phase strategy to those determined when using the integrated model with the unbounded daily repositioning (UDR) strategy. Table 3 presents the results

Table 3: Comparison of repositioning costs between 2-phase approach and integrated UDR approach.

Container Fleet Size	Repositioning Costs: 2-phase	Repositioning Costs: Integrated UDR	Percent Improvement	CPU Time: Integrated UDR (seconds)
1000	125,969	113,286	10.0	81.46
800	152,174	131,489	13.5	68.42
600	-	165,546	-	174.53

Note that all instances were solved within 3 minutes of CPU time and that no feasible solution was found by the 2-phase strategy in the tight capacity scenario.

The results in Table 3 demonstrate the value of integrated container management since empty repositioning costs are substantially reduced (by 10.0% in the overcapacity scenario and by 13.5% in the adequate capacity scenario). However, perhaps even more interesting is the fact that the integrated model is able to produce a feasible schedule in an environment with far fewer containers in the system (the 2-phase strategy was also unable to find a feasible solution with 700 containers in the system). The integrated container management model is able to fully exploit any routing flexibility offered by the service time windows for each booking. This indicates that it may be very beneficial for container operators to try to collaborate with their customers to obtain timely and accurate information about bookings, since it may allow the container operator to free up capital that is otherwise tied up in expensive tank containers.

In the next experiment, we investigate the relative performance of the three alternative deployment strategies for the integrated model. We first compare the costs determined by the UDR, BDR, and WR strategies in Table 4. Note that for the BDR strategy, we assume that a lower bound of 4 containers must be transported whenever a depot-to-depot repositioning move is initiated.

Table 4: Cost comparison for different integrated repositioning strategies.

Container Fleet Size	Repositioning Strategy	Minimum repositioning quantity	Repositioning costs	Land-based transportation costs	Ocean-based transportation costs	Total costs
1000	UDR	1	113,286	298,908	503,800	915,994
1000	WR	1	129,372	298,716	504,374	932,462
1000	BDR	4	113,668	298,668	503,920	916,256
600	UDR	1	165,546	294,372	524,589	984,507
600	WR	1	173,611	294,564	524,268	992,443
600	BDR	4	169,727	294,624	526,709	991,060

The results indicate that it is worthwhile to make repositioning decisions daily as opposed to weekly, especially in environments with overcapacity (a difference of about 12% in repositioning costs, and of about 2% in total cost). It also appears that in the overcapacity environment, deciding the proper timing of repositioning is more important than deciding on the number of containers to reposition because imposing a lower bound on the repositioning quantity has relatively little impact on cost. When capacity is tight, however, imposing such a lower bound does lead to a slight cost increase.

The container utilization statistics associated with the schedules produced for these strategies are presented in Table 5. The final four columns in the table present the percentage of time containers were in inventory at a depot, the percentage of time containers were being moved empty, the percentage of time containers were being moved loaded, and the percentage of time containers were being moved.

Table 5: Utilization comparison for different integrated repositioning strategies.

Container Fleet Size	Repositioning Strategy	Minimum repositioning quantity	Inventory time	Empty transport time	Loaded transport time	Total transport time
1000	UDR	1	45.28%	18.28%	36.43%	54.72%
1000	WR	1	49.24%	14.43%	36.33%	50.76%
1000	BDR	4	44.95%	18.37%	36.68%	55.05%
600	UDR	1	18.66%	26.37%	54.97%	81.34%
600	WR	1	21.17%	23.60%	55.24%	78.83%
600	BDR	4	18.91%	26.15%	54.94%	81.09%

The results in Table 5 demonstrate the potential value of integrated container management, since high levels of utilization can be realized, *i.e.*, containers being transported over 80% of the time (when capacity is tight). In reality the utilization is even higher since the results presented in the table include some start and end effects. In the four months in the middle of the planning horizon, *i.e.*, months two through five, total transport time reaches 95% when capacity is tight. High asset utilization is of key importance to container operators due to relatively high asset capital costs. By using integrated container management models, operators may be able to increase their revenue (satisfying more demands) without having to increase the number of containers under management and thus without requiring additional capital investments.

An important limitation on the approach presented in this chapter is that it does not consider demand uncertainty. Although dynamic, the model discussed in this section is not stochastic; therefore, it fails to capture a very important component of the problem. It is crucial to identify approaches to handle uncertainty in the input data, especially uncertainty related to the forecasted demand. If the freight plan that results from solving the model fails to consider discrepancies between forecasted and actual demand realizations, the tank container operator may suffer substantial degradation of customer service levels, excessive control costs, or both.

Consider for instance Figure 2 where inventory levels for three depots, Los Angeles, Shanghai and Hong Kong, are shown for a solution obtained by solving our integrated model under the UDR strategy. It is clear that given the probability distributions in Table 2, empty containers need to be repositioned to Shanghai and Hong Kong in order to account for trade imbalances. Observe that since demand uncertainty is not addressed, a solution to the deterministic model repositions only the number of containers strictly necessary to satisfy forecasted demands. This leads to inventory accumulating in Los Angeles, and dangerously low inventory levels in Shanghai and Hong Kong. Since no buffers are included in the plan, some depots are vulnerable to stockouts during execution. Furthermore, cost-effective control is difficult; if either of the two depots in the Asia region requires additional containers, they would most likely need to be repositioned from another region, representing a potentially costly control decision. This example illustrates that it is imperative to identify methods that generate cost-effective freight plans that hedge against demand uncertainty.

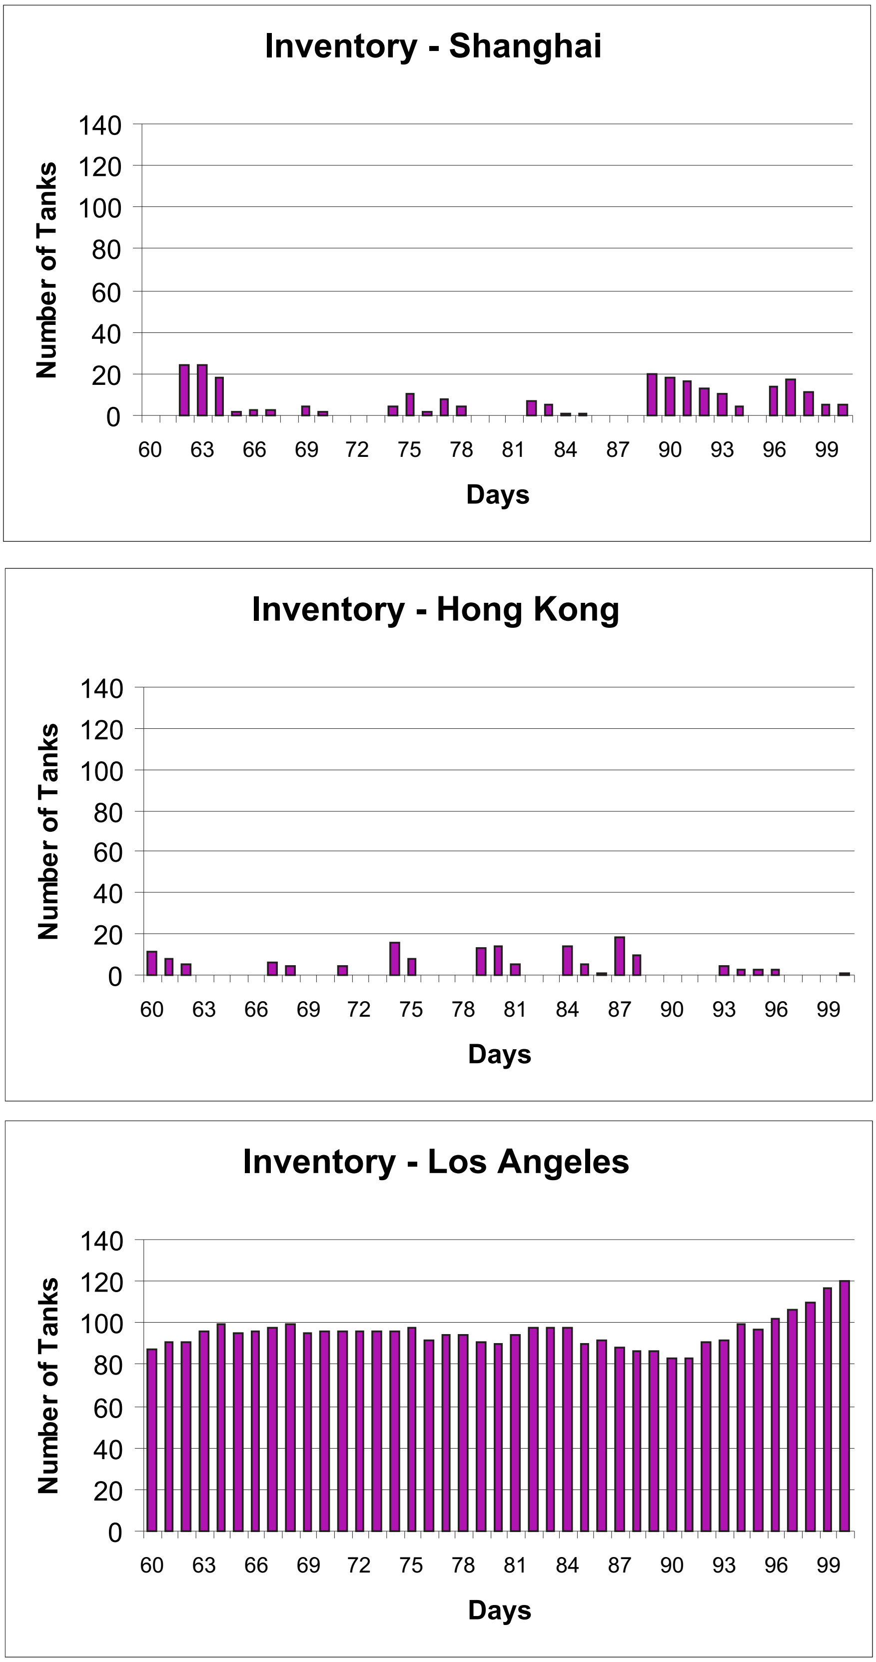

Figure 2: Inventory levels in three depots for a solution of one instance of the problem.

CHAPTER III

ROBUST OPTIMIZATION FOR EMPTY REPOSITIONING PROBLEMS

In the previous chapter, a fleet management planning and control problem for global tank-container operators was addressed. Problems of this type arise for many large freight transportation service providers, and are typically decomposed in two parts: the empty repositioning problem and the problem of allocating units to customer demands. In this chapter we consider the former problem– the management of empty resources over time– and propose a planning model that explicitly considers demand uncertainty.

Almost all freight transporters serve a set of load requests that is imbalanced in both time and space. Thus, when a resource such as a container or truck driver arrives at the destination location of a loaded move, there may not be an opportunity to match that resource in a timely way with a new loaded move outbound from that location. To correct imbalance, transporters move resources empty between locations; planning and executing empty moves that enable future customer demands to be served at low cost is a primary challenge.

Currently, the more sophisticated transporters address this planning problem using deterministic flow optimization models over time-space networks. Network nodes are defined at relevant decision points, and connect forward in time with other nodes via arcs that represent management decisions (and their costs) such as holding inventory of empty resources, or repositioning such resources between locations. Next, point forecasts are developed for the net expected supply for resources at some or all of the time-space network nodes, and initial and final resource states are specified. A feasible flow on such a network represents a set of feasible empty management decisions, and network optimization algorithms can be used to find an optimal flow. For problems that can be decomposed by resource, the resulting problems are often single-commodity minimum cost network flow problems, which can be solved very efficiently. In practice, these models are used in a rolling horizon implementation where a solution is obtained for a long planning horizon, but only the decisions in an initial set of time periods are implemented.

One major deficiency of this traditional approach is that there may be significant uncertainty in the forecasts of resource net supply at each time-space node, especially towards the end of the planning horizon. When realized demands differ from forecasts, the implemented empty allocation plan may be far from optimal. Stochastic models for these problems typically replace point forecasts of expected net supply with distribution forecasts, and attempt to find solutions that minimize total expected cost over the planning horizon; in this case, difficult-to-solve stochastic dynamic programming or stochastic integer programming problems result.

In contrast to existing stochastic approaches that focus on expected cost minimization, our research develops a robust optimization approach for repositioning problems. Our approach borrows ideas from both Ben-Tal et al. [5] and Bertsimas and Sim [10]. Since during the execution of the plan, dynamic adjustments might need to be taken, we model the repositioning problem using a two-stage framework similar to the adjustable robust counterpart (ARC) for linear programming problems (see Ben-Tal et al. [5]). To avoid requiring estimates of probability distributions, we model forecast uncertainty using intervals about a nominal expected net supply at each time-space node and limit system-wide deviations from nominal values through an uncertainty budget (see Bertsimas and Sim [10]). Thus, our robust approach seeks to find minimum cost solutions that satisfy special feasibility conditions for any demand realization in which (1) each time-space node demand lies within its forecast interval, and (2) no more than k demands can simultaneously take their worst-case value. Parameter k can be seen as a planner's view on the accuracy of the nominal expected net supply values: when $k = 0$, the planner has complete confidence in the nominal expected net supply values and the problem reduces to the deterministic problem, and when $k = \infty$, the planner has no confidence in the nominal expected net supply values and solutions must satisfy the special feasibility conditions for all realizations. In other words, the planner's risk-aversion can be captured with the value of k.

The special feasibility conditions require that a robust repositioning plan (1) satisfies flow balance equalities and flow bounds with respect to the nominal expected net supply values, and (2) is *recoverable*, that is, it can be transformed into a plan that satisfies flow balance equalities and flow bounds for every demand realization using dynamic adjustments, also referred to as recovery actions. Note that any solution feasible for a specific realization of net supplies is necessarily infeasible for every other realization. Dynamic adjustments are similar to recourse actions in two-stage stochastic programming models and in the adjustable robust counterpart approach. In the simplest planning problem considered in this chapter, dynamic policies allows recovery flow changes only on inventory arcs between the same space point in consecutive time periods; this scenario, therefore, corresponds to the case where each spatial location must hedge independently against uncertain future outcomes. We also consider dynamic policies that allow limited reactive repositioning between locations.

The robust repositioning problem that allows only inventory recovery actions is shown to be polynomially-solvable. For the robust repositioning problems that allow dynamic adjustments that use reactive repositioning, we develop sets of feasibility conditions whose sizes, while not polynomial, do not grow with the size of the uncertain outcome space. We illustrate the application of the robust repositioning framework developed with a set of computational experiments involving a global container repositioning problem of realistic size. The results provide insight into how different levels of confidence in the nominal expected net supply values, measured by parameter k, and different degrees of flexibility for performing recovery actions affect the cost of a robust repositioning plan, and therefore the price of robustness.

3.1 Robust Feasibility-Recovery Optimization Framework for Problems with Right-hand-side Uncertainty

Before focusing specifically on empty repositioning flow problems, we first outline a general approach for optimization problems with right-hand-side uncertainty that considers feasibility recovery actions. The approach applies the adjustable robust counterpart approach developed for uncertain linear programs in Ben-Tal et al. [5] to mixed-integer programming problems in a specific way that is useful for rolling-horizon implementations of dynamic optimization models where ensuring the existence of future feasible solutions is important.

Consider first a nominal optimization problem:

$$\mathbf{NP} \quad \min_x \left\{ c^T x : Ax = b,\ x \in X \right\} \tag{14}$$

where c is an n-vector, A is an m by n matrix, b is a deterministic m-vector referred to as the *vector of nominal right-hand-side values*, and x is an n-vector of decision variables. Set X is used to describe any additional constraints on the components of x, such as lower and upper bounds as well as integrality.

Now suppose that the right-hand side vector b may be uncertain, and further that each potential realization is a realization of some random vector $\widetilde{b}$. Given an assumed distribution for $\widetilde{b}$, stochastic approaches for extending **NP** include both chance-constrained programming as well as two-stage stochastic programming with recourse. An alternative is to use a robust optimization approach, planning for worst-case outcomes of $\widetilde{b}$. Since it will usually be overly-conservative to seek solutions that are robust with respect to all potential realizations within the finite support of an appropriate distribution for $\widetilde{b}$, robust approaches typically instead consider only outcomes that are members of some smaller, user-defined *uncertainty set* $\mathcal{Z} \subset \mathbb{R}^m$. Note that we assume that the nominal vector b is a member of $\mathcal{Z}$.

In typical robust optimization problems, one searches for a solution x that remains feasible given any uncertain outcome in $\mathcal{Z}$. However, it is clear that any feasible solution x to **NP** is infeasible given any non-zero perturbation $\delta \in \mathbb{R}^m$ such that $b + \delta \in \mathcal{Z}$. An alternative approach is to search for a solution x that can be made feasible using *adjustable* variables w, following the approach developed in [5] for linear programs.

Consider then the following adjustable robust optimization problem:

$$\mathbf{ARP} \quad \min_x \left\{ c^T x : Ax = b,\ x \in X,\ \forall\, b + \delta \in \mathcal{Z}\ \ \exists\, w \in W\ :\ Ax + Bw = b + \delta \right\} \tag{15}$$

where B is an m by p matrix and w is a p-vector of adjustable decision variables that may differ for each realization $b + \delta \in \mathcal{Z}$. Set W does not depend on δ, and represents a set of *feasibility recovery constraints* that may place additional restrictions on the adjustable decisions and should be developed in the context of the application problem. In the context of freight transportation operations, feasibility recovery constraints can be used to model dynamic policies.

Formulation **ARP** is essentially a special case of the adjustable recourse counterpart modeling framework developed in Ben-Tal et al. [5], but generalized to allow integer variable restrictions through X. It is a special case since **ARP** considers only uncertainty in the right-hand-side vector, and additionally limits x to solutions that are feasible with respect to the nominal vector b. Note that **ARP** does not capture in its objective the potential costs of adjustable decisions w, but rather only ensures that a feasible adjustment of x exists for all realizations. However, by carefully choosing which decisions to make fixed versus adjustable and by judiciously restricting the adjustable decisions using the constraints W, it is possible in application to appropriately capture the dominant system costs in the objective function $c^T x$ while ignoring the smaller costs of adjustable decisions w.

In this spirit, it may be useful in many dynamic applications of such a framework to link an adjustable variable w_a to each fixed variable x_a, and to allow no others. If we treat $x+w$ as a transformation of the initial decisions x, we can set $B = A$ and define a *transformable* robust optimization problem as

$$\textbf{TRP} \quad \min_x \left\{ c^T x : Ax = b,\ x \in X,\ \forall\, b + \delta \in \mathcal{Z}\ \exists\, w \in W\ :\ Aw = \delta,\ x + w \in X \right\}. \quad (16)$$

Note that in **TRP**, we explicitly provide an additional linkage between the initial decisions and the adjustable decisions by forcing $x + w \in X$, which may be used to model bounds that apply to both initial and transformed decisions. Again, while it may seem that such a transformable formulation allows the model to change any fixed decision x_a, the user can use the bounding mechanism provided by W to ensure that only "low-cost" adjustments to the initial plan are allowed. While it is natural, for example, to force $w_a = 0$ for any decision x_a that is truly fixed when planning, it may also be useful to do so for any "high-cost" decisions.

To simplify notation for the remainder of this chapter, if we define set $H(W, \delta)$ as

$$H(W, \delta) = \{x \mid\ \exists\, w \in W\ :\ Aw = \delta,\ x + w \in X\}, \quad (17)$$

then we can write **TRP** as

$$\mathbf{TRP}(W,\varphi) \quad \min_x \left\{ cx : Ax = b,\ x \in X,\ x \in \bigcap_{\delta\in\varphi} H(W,\delta) \right\}, \tag{18}$$

where $\varphi = \{\delta :\ b + \delta \in \mathcal{Z}\}$. We can write (18) in an extended form as

$$\begin{array}{lll} \mathbf{TRP}(W,\varphi) & \text{minimize} & c\,x \\ & \text{s.t.} & Ax = b \\ & & x \in X \\ & & Aw_\delta = \delta \quad \forall\,\delta \in \varphi \\ & & x + w_\delta \in X \quad \forall\,\delta \in \varphi \\ & & w_\delta \in W \quad \forall\,\delta \in \varphi \end{array}$$

where w_δ is a transformation vector applied to x given uncertain outcome δ. Observe that the number of decision variables and the number of constraints in the extended formulation of (18) may dramatically increase from (14) when the size of the outcome space φ is large.

3.2 Robust Empty Repositioning

We now consider the $TRP(W,\varphi)$ problem framework in the context of developing empty repositioning plans. Consider a transport operator managing a homogeneous fleet of reusable resources using a centralized control. Such resources may represent for instance containers, tank-containers, railroad cars or trucks. Further suppose that to manage these resources, the decision-maker need only track two state attributes: location and empty/loaded status. Other potential state attributes (*e.g.*, those required for maintenance) are ignored.

To manage this system, the operator maintains a network of storage depots from which all empty resources are sourced and to which all empty resources are returned. The empty repositioning problem of the operator, then, is to determine a plan for repositioning empty resources between these depots to satisfy loaded move requirements.

In practice this problem arises naturally for tank container fleet operators. Such operators move loads globally using ocean transportation, and face significant loaded flow imbalance. One large operator in this industry manages its empty repositioning flows using a deterministic time-space network flow formulation, considering a six-month planning horizon discretized into weeks [51]. The operator develops point forecasts of the expected net supply of containers at each depot during each week of the horizon; negative net supply corresponds to demand for containers. Using costs for repositioning containers between depots, the operator determines an empty repositioning plan for the horizon and then implements the decisions for the current week. The process is repeated weekly.

3.2.1 The Nominal Repositioning Problem

Suppose that b, the vector of nominal right-hand side values, is the vector of time-space net supply forecasts in the container repositioning problem described above. The nominal empty repositioning problem can then be modeled using a time-expanded network $G = (\mathcal{N}, \mathcal{A})$. Assume for simplicity that a planning horizon has been discretized into $\rho + 1$ periods, $\{0, 1, 2, ..., \rho\}$. Let D be the set of depots in the system, and let $V^d_\tau = \{v^d_0, v^d_1, ..., v^d_\tau\}$ for each $d \in D$ be the ordered set of nodes v^d_t representing depot d at each time period t up to time τ. Let $V^d = V^d_\rho$, the complete node set for depot d, and $\mathcal{V} = \cup_{d \in D} V^d$. Let $b(v)$ represent the component of b corresponding to $v \in \mathcal{V}$.

Containers can be held in inventory at a depot from one time period to the next. Therefore, an *inventory arc* (v^d_t, v^d_{t+1}) exists for each $d \in D$ and $0 \leq t < \rho$. Let I be the set of all inventory arcs. A *repositioning arc* (v^i_t, v^j_{t+h}) is defined between depots i and j at time t when available, where $h \geq 1$ is the travel time in periods along this arc. Let R be the set of all repositioning arcs. To complete the network specification, we add to G a sink node s with net supply $b(s) = -\sum_{v^d_t \in V} b(v^d_t)$ and an arc connecting v^d_ρ to s for all $d \in D$. For consistency we add these arcs to set I; for simplicity, node s can be labeled as $v^d_{\rho+1}$ for any $d \in D$. Let $\mathcal{N} = \mathcal{V} \cup \{s\}$ and $\mathcal{A} = I \cup R$.

For each $a \in \mathcal{A}$, let $c(a)$ be the unit cost of flow on arc a. For $a \in R$, $c(a)$ represents the repositioning costs per unit transported, while for $a \in I$, $c(a)$ represents per period holding costs per unit. In virtually all freight transportation settings, the (actual) cost of holding a resource in inventory is much smaller than the cost of moving a resource. Furthermore, since differences in per unit holding costs at different depots are also minor, in many applications holding costs are assumed to be negligible.

The nominal repositioning problem may now be written as:

$$\textbf{NP} \quad \min_x \left\{ cx : Ax = b, x \in \mathbb{Z}_+^{|\mathcal{A}|} \right\} \tag{19}$$

where the decision vector x corresponds to the empty container flow on each arc and A is the node-arc incidence matrix implied by G defining the typical network flow-balance constraints $Ax = b$. Let $x(a)$ represent the flow on arc $a \in A$. Note that this formulation is equivalent to (14) where $X = \mathbb{Z}_+^{|\mathcal{A}|}$.

It is well-known that problem (19) can be solved to optimality in polynomial time with standard minimum cost network flow algorithms, or via linear programming. We note that a feasible solution may not exist for this problem as posed; well-known techniques can address this problem, but for clarity and simplicity we assume that a feasible solution to (19) exists.

3.2.2 Three Robust Repositioning Problems

Since the nominal estimate of the net supply $b(v)$ at each node $v \in \mathcal{V}$ may be uncertain, we now apply the robust framework $TRP(W, \varphi)$ to this repositioning problem. Let $\widetilde{b}(v) \in \mathbb{Z}$ be the random variable representing net supply at $v \in \mathcal{V}$. Now suppose that the decision-maker is also able to estimate an interval around $b(v)$ which contains all of the potential outcomes of $\widetilde{b}(v)$ for which future feasibility should be protected. Assuming the interval is symmetric around $b(v)$, we represent it as

$$\widetilde{b}(v) \in [b(v) - \hat{b}(v), b(v) + \hat{b}(v)] \quad \forall \ v \in \mathcal{V}, \tag{20}$$

where $\hat{b} \geq 0$. We assume that the decision-maker always knows with certainty the net supplies in the initial period, and therefore $\hat{b}(v_0^i) = 0$ for each $i \in D$.

To allow further control of the conservatism of our robust repositioning models, we adopt the approach proposed in Bertsimas and Sim [10] to restrict the allowable *joint* realizations of $\widetilde{b}$ using an uncertainty budget. To do so, we define a limited perturbation set φ_k as a function of a budget parameter k as follows:

$$\varphi_k = \left\{ \delta \in \mathbb{Z}^{|\mathcal{N}|} \ : \ \delta(v) = \hat{b}(v) z(v), \sum_{v \in \mathcal{V}} |z(v)| \leq k, \ |z(v)| \leq 1 \ \forall \ v \in \mathcal{V}, \ \ \delta(s) = -\sum_{v \in \mathcal{V}} \delta(v) \right\}, \tag{21}$$

and we assume that every realization of $\widetilde{b}$ can be represented by $b + \delta$ for some $\delta \in \varphi_k$. Note that the constraint on $\delta(s)$ is a technical condition used only to preserve balance (for this reason, we will ignore $\delta(s)$ for the majority of the discussion to follow). Parameter k specifies the maximum number of net supplies that may simultaneously take on an extreme value in a realization, and thus can be used to control the conservatism of the uncertainty set used by the model. When $k = 0$, the decision-maker is most aggressive assuming that every realization will conform to nominal. When $k \geq |N|$, the decision-maker protects against all potential outcomes that fall within the intervals specified by (20).

It is important to observe that (21) ignores the fact that a reduction in net supply at a given time in a specific depot is most likely coupled with an increase in net supply at another time epoch or depot. For instance, a reduction in net supply might be the result of a customer requesting additional (more than forecasted) empty containers at a demand origin depot, which therefore implies a net supply increase some time later at the corresponding destination depot. Ignoring net supply correlations may result in a worst-case scenario that includes joint realizations that are highly unlikely, and perhaps not possible; this may lead to plans that are overly conservative. Handling supply correlations is a topic for future research.

Since the robust framework presented proposes to determine a feasible solution to the nominal repositioning problem (19) that also is recoverable for all potential outcomes, a concept that will be important for the analysis to follow is the *vulnerability* of a set of nodes to demand perturbations.

Definition 1 (Node Set Vulnerability) *For a set of nodes $V \subseteq \mathcal{V}$, its vulnerability $\vartheta(V, k)$ is defined as*

$$\vartheta(V, k) \ = \ \max_z \left\{ \sum_{v \in V} \hat{b}(v)\, z(v) \ : \ \sum_{v \in V} |z(v)| \leq k \ , \ |z(v)| \leq 1 \ \forall \ v \in V \right\}. \tag{22}$$

Observe that $\vartheta(V,k)$ corresponds to the maximum aggregate deviation from the nominal values for the nodes in V over all demand realizations in φ_k. The vulnerability of a node set V with m members can be determined easily in polynomial time. Suppose $V = \{v_1, v_2, ..., v_m\}$ where $\hat{b}(v_1) \geq \hat{b}(v_2) \geq ... \geq \hat{b}(v_m)$. Then, $\vartheta(V,k) = \sum_{i=1}^{\min(m,k)} \hat{b}(v_i)$. Further, we denote by $\eta(V,k) = \cup_{i=1}^{\min(m,k)} \{v_i\}$ the set of nodes that realize their worst-case value in the determination of $\vartheta(V,k)$.

We now specify several different robust repositioning problems, each defined by a different feasibility recovery set W. For each problem, we develop necessary and sufficient conditions for feasible solutions.

3.2.2.1 The Inventory Robust Repositioning Problem

Suppose that a decision-maker would like to develop an empty repositioning plan in which each depot hedges independently against uncertainty using its own inventory. To model this case, let W_1 be the set of feasibility recovery transformation vectors w that allow integer flow changes only on inventory arcs:

$$W_1 = \{w \in \mathbb{Z}^{|\mathcal{A}|} \mid w(a) = 0 \quad \forall \; a \in R\}.$$

Flow changes on inventory arcs can be interpreted as using containers in inventory to satisfy a larger-than-expected demand (a negative flow change), or adding extra containers to inventory in the event of a larger-than-expected supply (a positive flow change).

The recoverable set $H(W_1, \delta)$ for a given perturbation δ from nominal is given by:

$$H(W_1, \delta) = \left\{ x \mid \; \exists \; w \in W_1 \; : \; Aw = \delta, \; x + w \in \mathbb{Z}_+^{|\mathcal{A}|} \right\}.$$

We can now define the inventory robust optimization problem using our earlier notation:

$$\mathbf{TRP1} = \mathbf{TRP}(W_1, \varphi_k).$$

Any repositioning plan x satisfying the feasibility conditions of **TRP1** is called *k-robust inventory feasible*. **TRP1** seeks the minimum cost k-robust inventory feasible solution, and can be written in extended form as

$$
\begin{array}{llll}
\mathbf{TRP1}(W_1, \varphi_k) & \text{minimize} & c\,x & \\
 & \text{s.t.} & Ax = b & (23) \\
 & & x \in \mathbb{Z}_+^{|\mathcal{A}|} & (24) \\
 & & Aw_\delta = \delta \quad \forall\, \delta \in \varphi_k & (25) \\
 & & x + w_\delta \in \mathbb{Z}_+^{|\mathcal{A}|} \quad \forall\, \delta \in \varphi_k & (26) \\
 & & w_\delta(a) = 0 \quad \forall\, a \in R,\ \ \delta \in \varphi_k & (27) \\
 & & w_\delta \in \mathbb{Z}^{|\mathcal{A}|} \quad \forall\, \delta \in \varphi_k & (28)
\end{array}
$$

While correct, the above formulation requires for each potential uncertain outcome δ a vector w_δ of decision variables representing the feasibility recovery transformation and an associated set of flow balance constraints. Clearly, such a formulation becomes intractable as the size of the outcome space grows. We now show that **TRP1** alternatively can be solved as a minimum cost network flow problem with flow lower bound constraints using only the original flow variables x.

To do so, consider any inventory arc $a = (v_t^d, v_{t+1}^d) \in I$ and the corresponding node set V_t^d. Given a specific uncertain outcome $\delta \in \varphi_k$, let $\sigma(a)$ be the cumulative deviation from nominal net supply at depot d by time t:

$$\sigma(a) = \sum_{v \in V_t^d} \delta(v).$$

Since the vulnerability $\vartheta(V_t^d, k)$ is the maximum cumulative deviation from the nominal net supply values for the nodes of depot d up to the tail node of arc a in any realization, it is clear then that $|\sigma(a)| \le \vartheta(V_t^d, k)$ for all $\delta \in \varphi_k$.

The relationship between the flow on an inventory arc $x(a)$ and the vulnerability of V_t^d will determine whether or not a solution is k-robust inventory feasible. This motivates the following definition.

Definition 2 (Weak Arc) *For a given repositioning plan x, an inventory arc $a = (v_t^d, v_{t+1}^d) \in I$ is a weak arc if*

$$x(a) < \vartheta(V_t^d, k).$$

Observe that if a is a weak arc, then the inventory at time t at depot d is not sufficient to protect against every potential uncertain realization in φ_k.

The following theorem now characterizes the set of feasible solutions for **TRP1**:

Theorem 1 *A feasible solution x for the nominal problem* (19) *is also a feasible solution for **TRP1** if and only if*

$$x(a) \geq \vartheta(V_t^d, k) \;\; \forall \;\; a = (v_t^d, v_{t+1}^d) \in I. \tag{29}$$

Proof. Given the definition of W_1, for any $\delta \in \varphi_k$ the only transformation vector w that can feasibly satisfy constraints (25), (27), and (28) in **TRP1** is given by

$$w(a) = \sigma(a) \quad \forall \;\; a \in I. \tag{30}$$

Thus, we focus attention on constraints (26).

To show necessity by contradiction, suppose that there exists a feasible solution x for **TRP1** such that $x(a) < \vartheta(V_t^d, k)$ for some arc $a = (v_t^d, v_{t+1}^d)$. Now consider the uncertain outcome $\delta \in \varphi_k$ such that $\delta(v) = -\hat{b}(v)$ for all $v \in \eta(V_t^d, k)$ and $\delta(v) = 0$ for all other $v \in V$. Thus, from (30) note that $w(a) = -\vartheta(V_t^d, k)$ and thus the transformed flow on arc a is $x(a) - \vartheta(V_t^d, k) < 0$ which violates constraint (26). Therefore, x cannot be a feasible solution for **TRP1**.

Sufficiency can also be shown by contradiction. Let x be a feasible solution of (19) satisfying (29). Now, consider uncertain outcome $\delta \in \varphi_k$ such that after applying transformation (30) to x there exists an arc $a = (v_t^d, v_{t+1}^d) \in I$ such that $x(a) + w(a) < 0$. This implies $\vartheta(V_t^d, k) \leq x(a) < -\sigma(a)$, which then implies $\delta \notin \varphi_k$. □

Theorem 1 shows that **TRP1** can be solved for any given φ_k by adding a flow lower bound constraint for each inventory arc $a \in I$ to (19). Thus, **TRP1** is polynomially-solvable using standard minimum cost network flow algorithms. Also, observe that the lower bound constraints for a specific depot j are independent of the vulnerability of the arcs for any other depot in the system.

It is also important to note that in order to develop the necessary and sufficient conditions in Theorem 1, we need only consider perturbations $\delta \leq 0$. Any positive component in δ implies an unexpected addition of containers into some depot of the system, and such an event does not act against the interests of a decision-maker attempting to determine a feasible container allocation. In the remainder of this chapter, we only consider perturbation vectors $\delta \leq 0$.

3.2.2.2 The Inventory-Pooling Robust Repositioning Problem

Suppose that container depots can hedge against uncertainty not only using their own inventory but also using inventory at other depots in the system. We use the term *reactive repositioning* to refer to dynamic adjustments that correspond to depot-to-depot container repositioning conducted in response to a perturbation from expected net supplies. For a given value of k, a group of depots may be able to jointly hedge against uncertainty with fewer total container resources.

This idea is illustrated for a simple two-depot system in Figure 3. The number inside each node corresponds to the nominal net supply value, the number above each arc corresponds to its flow, and the interval above a node determines the range in which δ can take values. Observe that the conditions of Theorem 1 are satisfied for the inventory arcs of depot A, but not for those of depot B. However, if we could reactively reposition a unit of inventory at time 1 from depot A to depot B, then we could recover feasibility for this problem given any realization in φ_1. Since depots A and B share a single container resource in inventory to collectively hedge against uncertainty, such a solution is called an *inventory pooling 1-robust* solution.

We now formally define a feasibility recovery set W_2 for the inventory-pooling scenario:

$$W_2 = \{w \in \mathbb{Z}^{|\mathcal{A}|} \mid w(a) \geq 0 \quad \forall \ a \in R, \quad w(a) = 0 \quad \forall \ a = (v_0^j, u) \in R\}.$$

The set W_2 allows any integer flow change on each inventory arc, and non-negative integer flow changes on repositioning arcs. By enforcing non-negativity, we assume that only minor changes are allowed to the initial repositioning plan. Further, we do not allow any reactive flow changes on repositioning arcs that begin in the initial time epoch, since such decisions

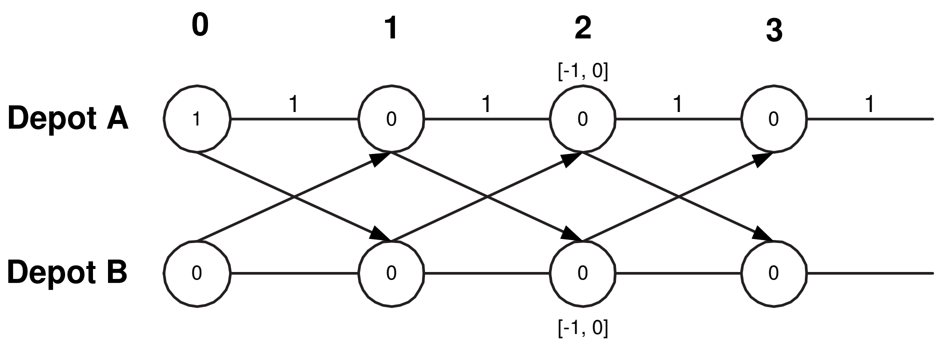

Figure 3: The solution is robust for $k = 1$ if containers in Depot A can be repositioned reactively to Depot B.

are assumed to be fixed. If the decision-maker intends to fix decisions for multiple initial time epochs, additional constraints could be added to W_2.

In this case, the recoverable set $H(W_2, \delta)$ for a given perturbation δ from nominal is given by

$$H(W_2, \delta) = \left\{ x \mid \exists\, w \in W_2 \ : \ Aw = \delta,\ x + w \in \mathbb{Z}_+^{|\mathcal{A}|} \right\},$$

and the general inventory-pooling robust repositioning problem is then:

$$\mathbf{TRP2} = \mathbf{TRP}(W_2, \varphi_k).$$

We can formulate Problem **TRP2** in extended form:

$$\begin{aligned}
\mathbf{TRP2}(W_2, \varphi_k) \ \min \quad & c\,x \\
\text{s.t.} \quad & Ax = b && (31)\\
& x \in \mathbb{Z}_+^{|\mathcal{A}|} && (32)\\
& Aw_\delta = \delta \quad \forall\, \delta \in \varphi_k && (33)\\
& x + w_\delta \in \mathbb{Z}_+^{|\mathcal{A}|} \quad \forall\, \delta \in \varphi_k && (34)\\
& w_\delta(a) = 0 \quad \forall\, a = (v_0^j, u) \in R,\ \forall \delta \in \varphi_k && (35)\\
& w_\delta(a) \geq 0 \quad \forall\, a \in R,\ \forall \delta \in \varphi_k && (36)\\
& w_\delta \in \mathbb{Z}^{|\mathcal{A}|} \quad \forall\, \delta \in \varphi_k && (37)
\end{aligned}$$

This direct integer programming formulation for **TRP2** may become difficult to solve as the space of feasible uncertain outcomes defined by φ_k grows large. Therefore, following the approach for the inventory robust repositioning problem, we seek methods for determining an optimal inventory-pooling robust solution that do not rely on enumerating the outcome space.

Creating a tight set of necessary and sufficient constraints for a nominal solution x to be robust with respect to any set of allowable recovery actions W requires ensuring that x is in the recoverable set $H(W, \delta)$ for every $\delta \in \varphi_k$. A useful methodology for testing this condition for the recovery actions considered in this chapter is to use the existence conditions for a feasible flow on a properly defined *recovery network*.

Let $G_W = (N_W, A_W)$ refer to the recovery network corresponding to allowable recovery action set W. The node set N_W is the same as the node set $\mathcal{N}$ of G. The arc set A_W contains all inventory arcs in I, and all repositioning arcs in R on which recovery flow is permitted to be nonzero by W. In the case of recovery set W_2, G_{W_2} contains only inventory arcs and each repositioning arc departing a depot at time $t > 0$.

To determine whether a repositioning plan x is recoverable using the action set W_2 for a specific realization $\delta \in \varphi_k$, we add appropriate net supplies to the nodes N_{W_2} and search for a feasible flow on G_{W_2}. To do so, let $\mathcal{I}^{x,\delta}(v)$ at each time-space depot node represent the marginal net inventory of containers available (or needed) in the recovery problem given x and δ:

$$
\begin{aligned}
\mathcal{I}^{x,\delta}(v_0^d) &= 0 \quad \text{for all } d \in D \\
\mathcal{I}^{x,\delta}(v_1^d) &= x(v_1^d, v_2^d) + \delta(v_1^d) \quad \text{for all } d \in D \\
\mathcal{I}^{x,\delta}(v_t^d) &= x(v_t^d, v_{t+1}^d) - x(v_{t-1}^d, v_t^d) + \delta(v_t^d) \quad \text{for all } d \in D,\ 1 < t \leq \rho
\end{aligned}
$$

where we note that $\delta(v_0^d) = 0$ for all $d \in D$. To understand this definition, suppose initially that $\delta = \mathbf{0}$. Given repositioning plan x, $\mathcal{I}^{x,0}(v_1^d)$ is the initial inventory at depot d that could be repositioned to serve recovery needs elsewhere. For $t > 1$, a positive value of $\mathcal{I}^{x,0}(v_t^d)$ indicates an increase in the number of units in inventory at depot d at time t, and therefore an additional number of containers that may be repositioned if warranted, or held in inventory to serve container needs at future times. On the other hand, a negative value of $\mathcal{I}^{x,0}(v_t^d)$ indicates a reduction in container inventory at time t. Such a reduction represents a demand for containers at that time. A container shortage will occur if the reduction is not satisfied by inbound containers either from inventory or via reactive repositioning. Observe that $\sum_{s=1}^{t} \mathcal{I}^{x,0}(v_s^d)$ corresponds to the actual inventory at time t at depot d. Since x is feasible for the nominal problem, this inventory is nonnegative for all values of t.

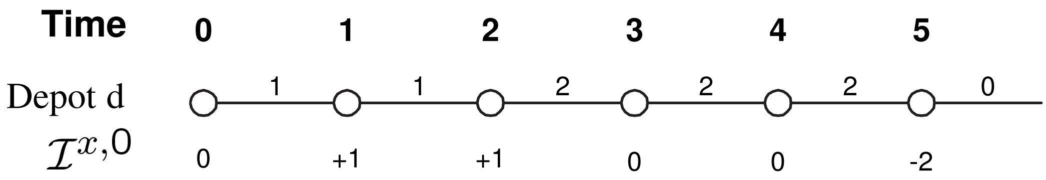

Figure 4: Given the flow on the inventory arcs of depot d, the value of $\mathcal{I}^{x,0}$ indicates that one unit at time 1 is available for reactive repositioning, and that an additional unit becomes available at time 2. The negative value of $\mathcal{I}^{x,0}$ indicates that at least 2 units must be in inventory by time 5. If any units are repositioned out of d, the same number must be repositioned back no later than time 5.

Given a nonzero realization vector δ, the net inventory at node v_t^d is changed by $\delta(v_t^d)$. Hence, the definition of $\mathcal{I}^{x,\delta}(v)$ models the net inventory availability (or requirement) at each node after the realization. Observe that $\sum_{s=1}^{t} \mathcal{I}^{x,\delta}(v_s^d)$ is not necessarily greater than or equal to 0 for all values of t. A negative value for this expression implies the necessity to reactively reposition units into the depot by time t in order to avoid a container shortage.

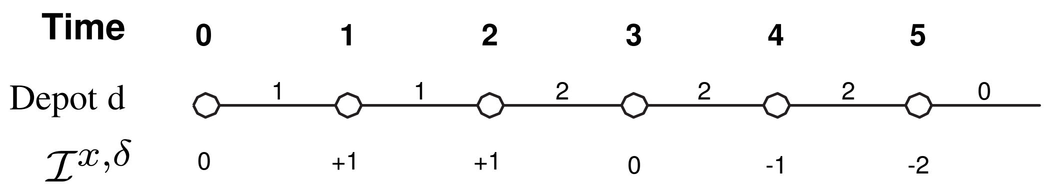

Figure 5: Value of $\mathcal{I}^{x,\delta}$ after perturbation $\delta(v_4^d) = -1$, $\delta(v_1^d) = \delta(v_2^d) = \delta(v_3^d) = \delta(v_5^d) = 0$. At least one unit must be in inventory by time 4 and at least 2 additional units by time 5 to recover feasibility.

Using $\mathcal{I}^{x,\delta}$, we can complete the definition of the recovery network G_{W_2}. Each arc $a \in A_{W_2}$ is given a flow lower bound $\ell(a) = 0$ and upper bound $u(a) = +\infty$. The net supplies d at each node are given by

$$
\begin{aligned}
d(v_t^d) &= \mathcal{I}^{x,\delta}(v_t^d) \quad \forall\, d \in D,\ t = 0, 1, 2, \cdots, \rho \\
d(s) &= -\sum_{v \in V} d(v)
\end{aligned}
$$

Let $G_{W_2}(x, \delta)$ refer to the recovery network with net supply vector d defined as above.

Proposition 1 *A feasible solution x for the nominal problem* (19) *belongs to the recoverable set $H(W_2, \delta)$ for a given δ if and only if there exists a feasible flow in $G_{W_2}(x, \delta)$.*

Proof. By construction of the network and its associated net supplies d, a feasible flow in $G_{W_2}(x, \delta)$ defines a set of feasible reactive repositioning decisions and inventory flow changes w restoring the feasibility of x given δ: for arcs $a \in R$, $w(a)$ is simply the flow on the corresponding arc in A_{W_2}, and for arcs $a = (v_t^d, v_{t+1}^d) \in I$ for $t > 1$, $w(a)$ is the flow on the arc in A_{W_2} minus $x(a)$. It is also not difficult to see that a w corresponding to an $x \in H(W_2, \delta)$ can be used to construct a feasible flow in $G_{W_2}(x, \delta)$: for $a \in R$, the flow on the associated arc in A_{W_2} is $w(a)$ and for $a \in I$, the flow is $x(a) + w(a)$. $\square$

We now derive necessary and sufficient conditions for the existence of a feasible flow in general recovery networks $G_W(x, \delta)$, where W allows general flow changes on all inventory arcs (v_t^d, v_{t+1}^d) for $t \geq 1$ and non-negative or zero flow changes on repositioning arcs. To do so, we first introduce two definitions.

Definition 3 (Competing Arc Set) *A set of arcs $S \subseteq I$ is competing if every directed path P in G_W has $|P \cap S| \leq 1$.*

Such arcs essentially compete for inventory to protect against uncertainty, since containers moved to satisfy a need of one arc cannot be later used to satisfy the need of any other in the set since no path for flow exists.

Definition 4 (Inbound-closed Node Set) *A set of nodes $C \subseteq N_W$ is inbound-closed if there exists no directed path P in G_W from any node $i \in N_W \setminus C$ to any node $j \in C$.*

Since no reactive flow paths exist into an inbound-closed node set, no additional containers can be brought into these nodes to satisfy net demand generated in excess of nominal. Thus, each such set must contain enough pooled inventory to hedge against a worst-case outcome.

Using these definitions, we can define a set of cuts in G_W that characterize feasible flows in $G_W(x,\delta)$. Let $\Delta^{\text{out}}(U) = \{(u,v) \in A_W \mid u \in U,\ v \in N_W \setminus U\}$. Define

$$\mathcal{U}_W = \{U \subseteq N \mid U \text{ is inbound-closed, } \ \Delta^{\text{out}}(U) \cap I \text{ is competing}\}.$$

Proposition 2 *There exists a feasible flow in $G_W(x,\delta)$ if and only if for every set of nodes $U \in \mathcal{U}_W$*

$$\sum_{v \in U} d(v) \geq 0.$$

Proof. It is known (see *e.g.*, Cook et al. [20]) that there exists a feasible flow in $G_W(x,\delta)$ if and only if

$$\sum_{a \in \Delta^{\text{out}}(U)} \ell(a) \leq \sum_{v \in U} d(v) + \sum_{a \in \Delta^{\text{out}}(N \setminus U)} u(a) \qquad \text{for all } U \subseteq N. \tag{38}$$

Consider then any node set $U \subseteq N$ such that $U \notin \mathcal{U}_W$. By definition of $\mathcal{U}_W$, there must exist a $v_1 \in N \setminus U$ and $v_2 \in U$ where $(v_1, v_2) \in A_W$. Since $u((v_1, v_2)) = +\infty$, (38) is always satisfied for such U.

Now consider any node set $U \in \mathcal{U}_W$, and note that because U is inbound-closed, there are no arcs into set U. Since $\ell(a) = 0$ for all $a \in A_W$, (38) reduces to

$$0 \leq \sum_{v \in U} d(v) \quad \text{for all } U \in \mathcal{U}_W.$$

□

The necessary and sufficient conditions in Proposition 2 can be enforced through a set of constraints on the nominal flow variables, as proposed by the following theorem:

Theorem 2 *A feasible solution x of the nominal problem* (19) *is also feasible for **TRP2** if and only if for every set of nodes $U \in \mathcal{U}_{W_2}$*

$$\sum_{a \in \Delta^{out}(U) \cap I} x(a) \ \geq \ \vartheta(U,k)$$

Proof. Let x be a feasible solution of the nominal problem (19) and $\delta \in \varphi_k$. By definition of $\mathcal{I}^{x,\delta}$ it is clear that

$$\sum_{s=1}^{t} \mathcal{I}^{x,\delta}(v_s^d) = x(v_t^d, v_{t+1}^d) + \sigma((v_t^d, v_{t+1}^d)) \qquad \forall\, d \in D$$

and therefore that

$$\sum_{v \in U} \mathcal{I}^{x,\delta}(v) = \sum_{a \in \Delta^{\text{out}}(U) \cap I} x(a) + \sum_{d \in D} \sigma(a^d)$$

for every $U \in \mathcal{U}_W$, where $a^d \in \Delta^{\text{out}}(U) \cap I$ is the inventory arc for depot d in cut $\Delta^{\text{out}}(U)$. Thus, by Propositions 1 and 2, solution x is feasible for **ROP2** if and only if

$$\sum_{v \in U} \mathcal{I}^{x,\delta}(v) = \sum_{a \in \Delta^{\text{out}}(U) \cap I} x(a) + \sum_{d \in D} \sigma(a^d) \geq 0 \qquad \forall\, U \in \mathcal{U}_{W_2} \tag{39}$$

holds for each δ. But since $\delta \in \varphi_k$, $-\sum_{d \in D} \sigma(a^d)$ can be bounded:

$$-\sum_{d \in D} \sigma(a^d) \leq \vartheta(U,k) \quad \forall\, d \in D.$$

We note that this bound is tight for at least one $\delta \in \varphi_k$ (namely, $\delta(v) = -\hat{b}(v)$ for all $v \in \eta(U,k)$). Thus, condition (39) simplifies to

$$\sum_{a \in \Delta^{\text{out}}(U) \cap I} x(a) \geq \vartheta(U,k) \quad \text{for all } U \in \mathcal{U}_{W_2}.$$

□

As an aside, we note that the conditions in Theorem 2 (and the techniques used to develop the theorem) could also be used for the inventory robust problem **TRP1** given an appropriately-defined recovery network G_{W_1} and set $\mathcal{U}_{W_1}$.

Using Theorem 2, an alternative integer programming formulation for **TRP2** is:

$$\begin{aligned}
\min \quad & c\,x \\
\text{s.t.} \quad & Ax = b \\
& \sum_{a \in \Delta^{\text{out}}(U) \cap I} x(a) \geq \vartheta(U,k) \quad \text{for all } U \in \mathcal{U}_{W_2} \\
& x \in \mathbb{Z}_+^{|\mathcal{A}|}
\end{aligned}$$

Observe that the size of the constraint set specifying necessary and sufficient conditions for a feasible solution to **TRP2** is independent of the size of the uncertain outcome space characterized by φ_k. Note also that the number of variables required for the robust formulation is equal to the number required for the nominal problem. However, the formulation requires a separate constraint for each element of $\mathcal{U}_{W_2}$, which is the set of all inbound-closed subsets of the nodes N_W that satisfy the competing inventory arc condition. It should be clear that enumerating this constraint set may be computationally expensive, and that the resultant integer program may be difficult to solve if the constraint set is large. However, for many problems arising in practice, this will not be the case; we will elaborate on these ideas later in this chapter.

3.2.2.3 A Restricted Inventory-Pooling Robust Repositioning Problem

Given a set of depots that pool inventory to hedge against uncertainty, a decision-maker may still wish include dynamic policies to limit further the allowable options for reactive repositioning during the planning phase. One approach that is appealing in practice designates *a priori* those depots that serve only as providers of reactive repositioning, and those that serve only as recipients. Suppose that D then is partitioned into two subsets: the depots in D_s can reposition containers reactively to other depots, but do not receive such support, while the depots in D_r may receive reactive repositioning containers but do not provide them. For example, in a geographic region a container operator may have a large hub depot, and many smaller depots. The operator then might wish to include the hub in D_s, and the smaller depots in D_r.

The feasibility recovery set W_3 can be specified for this scenario by a simple modification of W_2 where we restrict outbound reactive repositioning from depots in D_r and inbound reactive repositioning to depots in D_s:

$$
\begin{aligned}
W_3 = \{w \in \mathbb{Z}^{|\mathcal{A}|} \mid w(a) &\geq 0 \quad \forall \quad a \in R, \\
w(a) &= 0 \quad \forall \quad a \in \{(v_0^j, u) \in R\} \;\cup\; \{(v_t^a, v_{t+h}^b) \in R \mid a \in D_r \text{ or } b \in D_s\}\}
\end{aligned}
$$

The recoverable set $H(W_3, \delta)$ for a given realization $\delta \in \varphi_k$ is then

$$H(W_3, \delta) = \left\{ x \mid \; \exists\, w \in W_3 \; : \; Aw = \delta, \; x + w \in \mathbb{Z}_+^{|\mathcal{A}|} \right\},$$

and a restricted inventory-pooling robust repositioning problem is

$$\mathbf{TRP3} = \mathbf{TRP}(W_3, \varphi_k).$$

While it is not too difficult to extend the analysis developed in Section 3.2.2.2 to determine a tight set of necessary and sufficient conditions defining feasible solutions x for **TRP3**, it turns out that these problems tend to require larger sets of constraints; the basic intuition behind this result is that fewer available reactive arcs lead to fewer reactive flow paths and therefore more competing arc sets. Since effective solution procedures may therefore require techniques to generate such constraints dynamically as needed rather than *a priori*, we develop in this section some additional concepts that should prove useful for such techniques.

In the **TRP3** setting, it is clear that each depot in D_s must not dispatch inventory for reactive repositioning that it will need later to cover its own needs. Recall that given a nominal solution x and a realization $\delta \in \varphi_k$, the inventory available at time t at depot $d \in D_s$ is given by

$$\sum_{s=1}^{t} \mathcal{I}^{x,\delta}(v_s^d) = x(v_t^d, v_{t+1}^d) + \sigma((v_t^d, v_{t+1}^d)) \qquad \forall\, d \in D.$$

Furthermore, any containers that are reactively repositioned from depot d at or before time t will reduce this available inventory. Since this adjusted inventory cannot fall below zero, we can define the available container *support* $\mathcal{S}^{x,\delta}(v_t^d)$ at d at time t as the maximum number of containers that can be reactively repositioned from depot d by time t given nominal flow x and uncertain outcome δ, such that no container shortage occurs at d after time t. Mathematically,

$$\mathcal{S}^{x,\delta}(v_t^d) = \begin{cases} 0 & \text{if } t = 0 \\ \min_j\{x(v_j^d, v_{j+1}^d) + \sigma((v_j^d, v_{j+1}^d)) \mid t \leq j \leq \rho\} & \text{otherwise} \end{cases}$$

Note that for a fixed d, $\mathcal{S}^{x,\delta}(v_t^d)$ is a *non-decreasing* function of t. Support at time 0 is defined again to indicate that no reactive repositioning is allowed at that time epoch.

Following the approach for **TRP2**, we first define a recovery network G_{W_3} that will be used to develop conditions for the existence of a feasible recovery flow given a nominal problem solution. The network G_{W_3} is specified using the procedure in Section 3.2.2.2; the arc set A_{W_3} thus contains no repositioning arcs outbound from depots in D_r and none inbound to depots in D_s.

The supply vector d in this case can be specified using the definitions of $\mathcal{I}^{x,\delta}$ and $\mathcal{S}^{x,\delta}$. For nodes associated with depots in D_r, the definition is unchanged. However, for nodes associated with depots in D_s, the net supply is equal to the incremental support available at time t:

$$\begin{aligned} d(v_0^d) &= 0 \quad \forall\, d \in D \\ d(v_t^d) &= \mathcal{I}^{x,\delta}(v_t^d) \quad \forall\, d \in D_r,\ t = 1, \cdots, \rho \\ d(v_t^d) &= \mathcal{S}^{x,\delta}(v_t^d) - \mathcal{S}^{x,\delta}(v_{t-1}^d) \quad \forall\, d \in D_s,\ t = 1, \cdots, \rho \\ d(s) &= -\sum_{v \in N \setminus \{s\}} d(v) \end{aligned}$$

A negative net supply at some node v_t^d where $d \in D_r$ indicates a demand for containers that must be served from inventory, or via reactive repositioning. The net supplies for nodes v_t^d where $d \in D_s$ specify the maximum number of additional units that can be repositioned out of depot d at time t so that no container shortage occurs later in time. Note that by the definition of support, a negative net supply can only occur at such a node at $t = 1$; clearly in this case, there exists no feasible recovery flow.

Let $G_{W_3}(x, \delta)$ represent the recovery network G_{W_3} along with the associated net supply vector d. Again, a feasible flow in this recovery network has a one-to-one correspondence with a valid feasibility recovery vector for a given realization δ.

Proposition 3 *A feasible solution x of the nominal problem (19) is a member of the recoverable set $H(W_3, \delta)$ for a given δ if and only if there exists a feasible flow in $G_{W_3}(x, \delta)$.*

Proof. Parallel to proof of Proposition 1. □

Although it is true that valid necessary and sufficient conditions for the existence of a feasible flow in $G_{W_3}(x, \delta)$ are given by Proposition 2 using set $\mathcal{U}_{W_3}$ since W_3 is a recovery set of the form required by the Proposition, we now develop an explicit representation of these conditions using the special structure of the recovery set W_3. To do so, we first introduce some additional notation. Given a set of inventory arcs $\alpha \subseteq I$, let $T(\alpha)$ be the set of tail nodes of arcs in α. Let $C(\alpha)$ be the set of nodes from which the tail nodes of arcs in α can be reached, *i.e.*, the set of nodes from which there exists a directed path to the tail node of an arc in α. Note that $T(\alpha) \subseteq C(\alpha)$, and that $C(\alpha)$ is an inbound-closed set. Finally, let $D(\alpha)$ be the set of depots corresponding to α.

In this case, instead of considering all subsets of nodes $U \in \mathcal{U}_{W_3}$, we can determine whether or not a feasible flow in $G_{W_3}(x, \delta)$ exists by considering only sets of nodes $C(\alpha)$ defined by sets of competing inventory arcs α at depots in D_r, where each arc $a \in \alpha$ has a shortage with respect to δ: $x(a) + \sigma(a) < 0$.

Proposition 4 *There exists a feasible flow in $G_{W_3}(x, \delta)$ if and only if*

$$d(v_1^d) = \mathcal{S}^{x,\delta}(v_1^d) \geq 0 \quad \forall\, d \in D_s \tag{40}$$

and

$$\sum_{v \in C(\alpha)} d(v) \geq 0 \tag{41}$$

for all $\alpha \subseteq I$ where α is competing, $D(\alpha) \subseteq D_r$, and $x(a) + \sigma(a) < 0$ for each $a \in \alpha$.

Proof. From Proposition 2, necessary and sufficient conditions for the existence of a feasible flow in $G_{W_3}(x, \delta)$ are

$$\sum_{v \in U} d(v) \geq 0 \quad \forall\ U \in \mathcal{U}_{W_3} \tag{42}$$

It is now shown that the conditions in the proposition are equivalent to the conditions given by (42).

The necessity of (40) and (41) is clear, since each constraint in the two sets corresponds directly to some set $U \in \mathcal{U}_{W_3}$ for which the expression in (42) must hold. For each $d \in D_s$, (40) corresponds to $U = \{v_0^d, v_1^d\}$ which is clearly inbound-closed by definition of D_s. Each $C(\alpha)$ generating a constraint of type (41) is also inbound-closed. Further, the arc set $\Delta^{\text{out}}(C(\alpha)) \cap I$ can be shown to be competing. This set is comprised of α and $\overline{\alpha}$, where each $a \in \overline{\alpha}$ is associated with a different depot $d \in D_s$. Therefore, since α is a competing arc set, and no path exists containing arcs both in α and arcs in $\overline{\alpha}$ by definition of $C(\alpha)$, and no path exists containing more than one arc $\overline{\alpha}$ by definition of D_s, $\alpha \cup \overline{\alpha}$ is competing. Thus, constraint of type (41) corresponding to $C(\alpha)$ also has a corresponding constraint in (42).

We now show the sufficiency of the conditions by showing that if they hold, conditions (42) hold for all $U \in \mathcal{U}_{W_3}$. Consider any $U \in \mathcal{U}_{W_3}$, and let $D_u \subseteq D_r$ be the set of depots d where there exists an arc $a^d \in \Delta^{\text{out}}(U) \cap I$ satisfying $x(a^d) + \sigma(a^d) \geq 0$. Note then that for each $d \in D_u$,

$$\sum_{v \in V^d \cap U} d(v) = x(a^d) + \sigma(a^d) \geq 0. \tag{43}$$

We claim first then that the conditions for U in (42) are redundant with those for $\widetilde{U} = U \setminus \bigcup_{d \in D_u}(V^d \cap U)$ in this case. Clearly this is true if $\widetilde{U} = \emptyset$. If $\widetilde{U} \neq \emptyset$, note that $\widetilde{U} \in \mathcal{U}_{W_3}$ since the subset $\Delta^{\text{out}}(\widetilde{U})$ of competing arcs $\Delta^{\text{out}}(U) \cap I$ is competing, and since no outbound reactive repositioning arcs exist in A_{W_3} from any depot $d \in D_u$ the set $\widetilde{U}$ remains inbound-closed. Suppose first that $\widetilde{U}$ contains only nodes associated with depots in D_s. In this case, conditions (40) guarantee that $\sum_{s=1}^{t} d(v_s^d) \geq 0$ for any $0 \leq t \leq \rho$ by definition of $\mathcal{S}^{x,\delta}$, and that therefore for any such $\widetilde{U}$, (42) holds. Along with (43), this in turn implies that (42) is satisfied for U.

Finally, suppose instead that $\widetilde{U}$ contains nodes for each depot in some set $\widetilde{D}_r \subseteq D_r$, in addition perhaps to nodes for some depots in D_s. Let $\alpha \subset \Delta^{\text{out}}(\widetilde{U}) \cap I$ where $D(\alpha) = \widetilde{D}_r$. Clearly, α is competing, and $x(a) + \sigma(a) < 0$ for each $a \in \alpha$ by the definition of $\widetilde{U}$. Further, $C(\alpha) \subseteq \widetilde{U}$, where any additional nodes in $\widetilde{U}$ must be associated with depots in D_s. Since $d(v_t^d) \geq 0$ for $d \in D_s$ by (40) and the definition of $\mathcal{S}^{x,\delta}$, if (41) holds for $C(\alpha)$ then (42) holds for $\widetilde{U}$, and finally (43) implies further that (42) is also satisfied for U. □

Proposition 4 can now be used to specify necessary and sufficient conditions for a feasible nominal repositioning plan x to be a feasible solution to **TRP3**. To do so, we must introduce several additional definitions. First, we identify sets of arcs, denoted *vulnerable* sets, that require a constraint of type (41) to protect against a joint container shortage that will arise for at least one uncertain realization $\delta \in \varphi_k$:

Definition 5 *Given a feasible solution x to the nominal problem* (19)*, a set of inventory arcs $\alpha \subset I$ that are competing in G_{W_3} and where $D(\alpha) \subseteq D_r$ is* vulnerable *if there exists a $\delta \in \varphi_k$ such that $x(a) + \sigma(a) < 0$ for all $a \in \alpha$.*

Let $\Psi(x) = \{\alpha \subset I \mid \alpha \text{ is vulnerable}\}$. If α is vulnerable then there is $\delta \in \varphi_k$ such that $x(a) + \sigma(a) \leq -1$ for each $a \in \alpha$. Therefore, if α is vulnerable then

1. Each arc $a \in \alpha$ is weak, and

2. $\sum_{a\in\alpha} x(a) + |\alpha| \ \leq \ \vartheta\left(\bigcup_{(v_t^d, v_{t+1}^d)\in\alpha} V_t^d, k\right)$.

For a given $\delta \in \varphi_k$, a constraint (41) is required by Proposition 4 for arc set α. Note that a vulnerable set corresponds to a set of arcs α for which we know there exists at least one $\delta \in \varphi_k$ that would create a container shortage at all depots $D(\alpha)$ unless containers are reactively repositioned. Further, note that an arc set α which is not vulnerable will not require a constraint (41) for any $\delta \in \varphi_k$.

For a vulnerable set $\alpha \in \Psi(x)$, let $C'(\alpha) = C(\alpha) \bigcap \{\bigcup_{d\in(D\setminus D(\alpha))} V^d\}$ be the set of nodes associated with depots in $D \setminus D(\alpha)$ from which an arc in the vulnerable set can be reached. Let $D'(\alpha) \subseteq D \setminus D(\alpha)$ be the set of depots with nodes from which an arc in the vulnerable set can be reached, *i.e.*, the set of depots $d \in D \setminus D(\alpha)$ for which $C'(\alpha) \cap V^d \neq \emptyset$. Let $t_\alpha^d = \text{argmax}_t\{v_t^d \in C'(\alpha)\}$ for all $d \in D'(\alpha)$. Finally, let

$$I^d(\alpha) = \left\{(v_t^d, v_{t+1}^d) \in I \ \mid \ t \geq t_\alpha^d\right\} \quad \forall\, d \in D'(\alpha).$$

Definition 6 *A layer of a vulnerable set $\alpha \in \Psi(x)$, denoted $\theta(\alpha) \subset \bigcup_{d\in D'(\alpha)} I^d(\alpha)$, is a set of inventory arcs where*

$$|\theta(\alpha) \cap I^d(\alpha)| = 1 \quad \forall\, d \in D'(\alpha).$$

Each layer of α, therefore, contains one inventory arc from each depot in $D'(\alpha)$ leaving d at some time greater than or equal to t_α^d. Let $\Theta(\alpha) = \{\theta(\alpha) \subset I \mid \theta(\alpha) \text{ is a layer of } \alpha\}$.

Consider a vulnerable set α and a layer $\theta(\alpha)$, and let

$$\widetilde{U}(\alpha, \theta(\alpha)) \ = \ \left(\bigcup_{(v_t^d, v_{t+1}^d)\in\alpha} V_t^d\right) \bigcup \left(\bigcup_{(v_t^d, v_{t+1}^d)\in\theta(\alpha)} V_t^d\right)$$

Definition 7 *Given a feasible solution x to the nominal problem* (19)*, a vulnerable set $\alpha \in \Psi(x)$, and a layer $\theta(\alpha) \in \Theta(\alpha)$, let*

$$\sum_{a\in\alpha} x(a) + \sum_{a\in\theta(\alpha)} x(a) \geq \vartheta\left(\widetilde{U}(\alpha, \theta(\alpha)), k\right)$$

be the layer $\theta(\alpha)$ inequality.

We are now ready for the main theorem in this subsection:

Theorem 3 *A feasible solution x for the nominal problem* (19) *is also a feasible solution for **TRP3** if and only if*

$$x(a) \geq \vartheta(V_t^d, k) \quad \forall \; a = (v_t^d, v_{t+1}^d) \in I \; \textit{where} \; d \in D_s$$

and for every vulnerable set $\alpha \in \Psi(x)$, the layer $\theta(\alpha)$ inequality is satisfied for each $\theta(\alpha) \in \Theta(\alpha)$.

Proof. Let x be a feasible solution of the nominal problem (19). By Propositions 3 and 4, x is also a feasible solution for **TRP3** if and only if conditions (40) and (41) are satisfied for all $\delta \in \varphi_k$.

First we consider conditions (40). By the definition of $\mathcal{S}^{x,\delta}$, these conditions become

$$\min_{1 \leq j \leq \rho} \left\{ x(v_j^d, v_{j+1}^d) + \sigma((v_j^d, v_{j+1}^d)) \right\} \geq 0 \quad \forall \, d \in D_s \quad \forall \, \delta \in \varphi_k,$$

which are equivalent to the following sets of constraints for each $d \in D_s$:

$$x(v_t^d, v_{t+1}^d) \geq -\sigma((v_t^d, v_{t+1}^d)) \quad \forall \, t \in \{1, \cdots, \rho\} \quad \forall \, \delta \in \varphi_k. \tag{44}$$

For any $\delta \in \varphi_k$,

$$-\sigma((v_t^d, v_{t+1}^d)) \leq \vartheta(V_t^d, k)$$

where we note that this bound is tight for at least one $\delta \in \varphi_k$. Therefore (44) simplifies to

$$x(v_t^d, v_{t+1}^d) \geq \vartheta(V_t^d, k) \quad \forall \, t \in \{1, \cdots, \rho\},$$

which must hold for all $d \in D_s$. Thus,

$$x(a) \geq \vartheta(V_t^d, k) \quad \forall \, a = (v_t^d, v_{t+1}^d) \in I \text{ where } d \in D_s.$$

Now consider conditions (41). Using the definition of vulnerable sets, these conditions become

$$\sum_{v \in C(\alpha)} d(v) \geq 0 \quad \forall \, \alpha \in \Psi(x), \tag{45}$$

which again by Proposition 3 must hold for all $\delta \in \varphi_k$. The sum can be rewritten by partitioning the depots into those associated with the vulnerable arcs and those providing reactive repositioning support to the vulnerable arcs, and simplified as follows:

$$\begin{aligned}
\sum_{v \in C(\alpha)} d(v) &= \sum_{d \in D(\alpha)} \sum_{v \in C(\alpha) \cap V^d} d(v) + \sum_{d \in D'(\alpha)} \sum_{v \in C'(\alpha) \cap V^d} d(v) \\
&= \sum_{d \in D(\alpha)} \sum_{v \in C(\alpha) \cap V^d} \mathcal{I}^{x,\delta}(v) + \sum_{d \in D'(\alpha)} \sum_{s=1}^{t^d_\alpha} \left(\mathcal{S}^{x,\delta}(v^d_s) - \mathcal{S}^{x,\delta}(v^d_{s-1}) \right) \\
&= \sum_{a \in \alpha} x(a) + \sum_{d \in D(\alpha)} \sigma(a^d) + \sum_{d \in D'(\alpha)} \mathcal{S}^{x,\delta}(v^d_{t^d_\alpha}) \\
&= \sum_{a \in \alpha} (x(a) + \sigma(a)) + \sum_{d \in D'(\alpha)} \min_{a \in I^d(\alpha)} \{x(a) + \sigma(a)\}.
\end{aligned}$$

where a^d is the inventory arc for depot d in α. Thus, the condition for a specific α in (45) can now be replaced by a set of inequalities, one for each layer $\theta(\alpha) \in \Theta(\alpha)$:

$$\sum_{a \in \alpha} (x(a) + \sigma(a)) + \sum_{a \in \theta(\alpha)} (x(a) + \sigma(a)) \geq 0.$$

Rewriting yields

$$\sum_{a \in \alpha} x(a) + \sum_{a \in \theta(\alpha)} x(a) \geq -\sum_{a \in \alpha} \sigma(a) - \sum_{a \in \theta(\alpha)} \sigma(a). \tag{46}$$

Since $\delta \in \varphi_k$, we have a tight bound of the right-hand side of (46) using $\vartheta(\widetilde{U}(\alpha, \theta(\alpha)), k)$, therefore

$$\sum_{a \in \alpha} x(a) + \sum_{a \in \theta(\alpha)} x(a) \geq \vartheta(\widetilde{U}(\alpha, \theta(\alpha)), k)$$

which is the layer $\theta(\alpha)$ inequality. $\square$

Theorem 3 allows **TRP3** to be formulated as the following integer program:

$$\begin{aligned}
\min \quad & c\,x \\
\text{s.t.} \quad & Ax = b \\
& x(a) \geq \vartheta(V^d_t, k) \quad \forall\, a = (v^d_t, v^d_{t+1}) \in I, \;\; d \in D_s \\
& \sum_{a \in \alpha} x(a) + \sum_{a \in \theta(\alpha)} x(a) \geq \vartheta(\tilde{U}(\alpha, \theta(\alpha)), k) \quad \forall\, \alpha \in \Psi(x), \;\; \forall \theta(\alpha) \in \Theta(\alpha) \\
& x \in \mathbb{Z}_+^{|\mathcal{A}|}
\end{aligned}$$

Note that the layer inequality constraints are specified for each vulnerable set α, and that the vulnerable sets depend on the solution x. If the layer inequality constraints were alternatively required regardless of x for each competing arc set α in G_{W_3} where $D(\alpha) \subseteq D_r$, the formulation would remain valid. However, Theorem 3 shows that a given solution x can be checked for feasibility to **TRP3** using potentially fewer constraints. These ideas motivate computational approaches to this problem as future research.

3.2.3 Alternative Uncertainty Sets

Although the results developed above have assumed perturbation sets of the form specified by (21), it turns out that several different types of "budget" uncertainty sets may be used without changing the fundamental results. We now show that different budget uncertainty sets lead only to different specifications of the vulnerability ϑ of a node set V. In this section, let $\vartheta(V, \varphi)$ represent the vulnerability of node set V given an uncertainty set represented by the perturbation set φ.

1. *Maximum Scaled Deviation Per Depot*: Rather than limiting the maximum scaled deviation from the nominal net supply estimates for all depots jointly, a decision-maker may wish to limit each depot's deviation separately. Assuming for simplicity that each depot has an identical budget k, the alternative perturbation set for this case is given by $\varphi_k^D = \{\delta \in \mathbb{Z}^{|\mathcal{N}|} \;:\; \delta(v) = \hat{b}(v)z(v),\, \sum_{v \in V^d} |z(v)| \leq k \;\forall\; d \in D,\, |z(v)| \leq 1 \;\forall\; v \in \mathcal{V},\;\; \delta(s) = -\sum_{v \in \mathcal{V}} \delta(v)\}$. In this case, the vulnerability of a set $V \subseteq \mathcal{V}$ of nodes can be determined by
$$\vartheta(V, \varphi_k^D) \;=\; \max_z \left\{ \sum_{v \in V} \hat{b}(v)\, z(v) \;:\; \sum_{v \in V \cap V^d} |z(v)| \leq k \;\; \forall\, d \in D \;,\; |z(v)| \leq 1 \;\forall\; v \in V \right\},$$
a simple optimization problem similar to that given by (22) which can be solved by summing for each depot d the k largest values of $\hat{b}(v)$ for nodes $v \in V \cap V^d$, and then adding the sums for all depots.

2. *Telescoping Maximum Scaled Deviation*: One limitation in applying uncertainty set φ_k to problems with time-space nodes is that as the value of k is increased, the method may generate very conservative decisions in the first few planning periods

since a large proportion of the demands may simultaneously take on their worst case values. This limitation can be avoided by using a telescoping uncertainty set. Let $\kappa = \{k_1, k_2, ..., k_\rho\}$ be a vector of budget parameters, $k_t \leq k_{t+1}$, where parameter k_t represents the maximum number of time-space net supplies that may take on their worst-case value by time period t. If we define $V_t = \cup_{d \in D} V_t^d$, the perturbation set for this case is given by $\varphi_k^T = \{\delta \in \mathbb{Z}^{|\mathcal{N}|} \ : \ \delta(v) = \hat{b}(v)z(v), \sum_{v \in V_t} |z(v)| \leq k_t \ \forall \ t \in \{1, ..., \rho\}, \ |z(v)| \leq 1 \ \forall \ v \in \mathcal{V}, \ \ \delta(s) = -\sum_{v \in \mathcal{V}} \delta(v)\}$. Then, the vulnerability of $V \subseteq \mathcal{V}$ can be determined by

$$\vartheta(V, \varphi_k^T) \ = \ \max_z \left\{ \sum_{v \in V} \hat{b}(v)\, z(v) \ : \ \sum_{v \in V \cap V_t} |z(v)| \leq k_t \ \ \forall \, t \in \{1, ..., \rho\} \ , \ |z(v)| \leq 1 \ \forall \ v \in V \right\}.$$

This optimization problem is also easy to solve, given that the nodes in V are sorted by non-decreasing order of $\hat{b}(v)$: set $z(v_t^d) = 1$ for the largest value of $\hat{b}(v_t^d)$, as long as $\sum_{v \in V \cap V_\tau} |z(v)| \leq k_\tau$ for $\tau \geq t$, then proceed to the next largest value of $\hat{b}(v)$ and repeat. Note that it would also be simple to construct such a telescoping maximum scaled deviation uncertainty set where the maximums were applied per depot rather than system-wide.

3. *Maximum or Telescoping Maximum Absolute Deviation*: Finally, while budget uncertainty sets that limit the maximum total scaled deviation from nominal have the benefit of being independent of the vectors b and $\hat{b}$, there may be cases where the decision-maker would rather limit the maximum absolute deviation from nominal for which he would like to plan; in fact, generating an appropriate value for such a statistic may be relatively simple from an analysis of past forecast accuracy. In the non-telescoping case, such a perturbation set could be represented by $\varphi_k^A = \{\delta \in \mathbb{Z}^{|\mathcal{N}|} \ : \ \delta(v) = \hat{b}(v)z(v), \sum_{v \in \mathcal{V}} \hat{b}(v)|z(v)| \leq k \, , \ |z(v)| \leq 1 \ \forall \ v \in \mathcal{V}, \ \ \delta(s) = -\sum_{v \in \mathcal{V}} \delta(v)\}$. In this case, the vulnerability of $V \subseteq \mathcal{V}$ can be determined simply by

$$\vartheta(V, \varphi_k^A) \ = \ \min \left\{ k, \sum_{v \in V} \hat{b}(v) \right\}.$$

A telescoping set and vulnerability could be similarly defined.

3.3 *Using the Robust Repositioning Models in Practice*

Large carriers are typically organizations with regional business units. A region usually encompasses a group of facilities that are relatively close to one another geographically. Since repositioning costs tend to be smaller between facilities within a region than between those in separate regions, it may make sense to limit the allowable *reactive* control decisions considered while generating plans to those that represent intra-regional repositioning moves. Importantly, note that such an assumption does not mean that inter-regional moves are not planned or executed. Instead, the restriction simply forces the planning model to assume that facilities within the same region can share resources reactively, but that they cannot be shared reactively inter-regionally. In addition to eliminating costly reactive moves from consideration, such a restriction may also improve the computational tractability of the approach proposed in this research by limiting the size of the constraint sets required to specify the robust repositioning problems **TRP2** and **TRP3**.

For example, consider an international container manager that provides service from many container depots located across the globe. These depots may be naturally grouped into regions: *e.g.*, Southeast Asia, East Coast North America, Northern Europe, etc. While depots within a region might pool inventory to hedge against uncertainty, it might not be practical to allow reactive sharing across regional boundaries. This scenario can be modelled using a simple modification of **TRP2** or **TRP3** where inter-regional repositioning arcs are included in the nominal problem network G, but excluded from the specification of G_W. It should be clear, then, that the conditions in Theorems 2 and 3 decompose by region, leading to much smaller sets of constraints guaranteeing feasibility. Such decomposition also could allow the manager to specify certain regions where complete reactive pooling of the type specified by **TRP2** is allowed, and other regions where the restricted reactive pooling of **TRP3** is used.

The idea of partitioning a service area into regions is incorporated in the computational study now described. The computational experiment demonstrates that realistic instances of the empty repositioning problem can be solved effectively with the proposed robust optimization framework, and provides insight into how different levels of conservatism (determined by parameter k) and different degrees of flexibility for performing reactive repositioning (determined by the feasibility recovery constraints) affect the cost of a repositioning plan and therefore the price of robustness.

Consider now an example problem setting representative of those found in the tank-container industry. Twenty depots, each located nearby a seaport of global importance, were chosen to be part of the network where the tank-container operator provides transportation services. The depots were partitioned into eight geographic regions as described in Table 6. Transportation times between seaports were determined using a published schedule of port-to-port sailing for a large ocean carrier, in which the largest time corresponded to 5 weeks between ports in the east coast of North America and some ports in Asia. Transportation costs were assumed to be proportional to transportation times.

Table 6: Regions and depots in the computational test network.

Region	Depots
1. South East Asia	Singapore, Port Kelang
2. East Asia	Hong Kong, Shanghai, Busan
3. Japan	Kobe, Tokyo
4. Northern Europe	Southampton, Rotterdam, Hamburg
5. Southern Europe	Algeciras, Gioia Tauro
6. North America West	Los Angeles, San Francisco, Seattle
7. North America East	New York, Norfolk, Savannah
8. South America	Buenos Aires, Rio de Janeiro

In order to generate point forecasts of the net expected supply for each time period for each depot (*i.e.*, the vector of nominal values), a total of $3,000$ customer requests for loaded tanks were randomly generated uniformly over a time period of 57 weeks. Each request requires a number of tanks, which was randomly generated uniformly from the set $\{1, 2, 3, 4, 5\}$. Each request has an associated origin depot and destination depot (the determination of which is described in the following paragraph), where tanks are assumed to be sourced from and returned to, respectively. Aggregating across all requests, the weekly net inflow of empty tanks was calculated for each depot for each week; this value was then used as the corresponding point forecast $b(v)$ at each node v. Uncertain intervals were then constructed assuming fluctuations of up to 10% of nominal value: $\hat{b}(v) = \lceil 0.10 b(v) \rceil$.

Similar to the approach used in the computational study of Chapter 2, to incorporate geographic trade imbalances into the test problem, the origin region and destination region of each request was randomly generated using the probabilities given in Table 7. Within in a region, a specific depot is randomly selected with equal probability. For this distribution data, regions in North America and Europe are on average net sources of empty tanks, while regions in Asia and South America correspond to net sinks.

Table 7: Origin-destination distribution information for loaded demands between regions for computational test.

		Conditional probability of destination region							
Region	Probability of Origin	1	2	3	4	5	6	7	8
1	0.15	0	0.10	0.10	0.20	0.15	0.20	0.20	0.05
2	0.35	0.05	0	0.05	0.15	0.10	0.30	0.30	0.05
3	0.10	0.05	0.05	0	0.15	0.10	0.30	0.30	0.05
4	0.10	0.10	0.20	0.10	0	0.10	0.20	0.25	0.05
5	0.05	0.10	0.25	0.15	0.05	0	0.15	0.25	0.05
6	0.05	0.10	0.30	0.15	0.20	0.15	0	0.00	0.10
7	0.10	0.10	0.30	0.15	0.20	0.15	0.00	0	0.10
8	0.10	0.05	0.05	0.10	0.15	0.15	0.20	0.30	0

To avoid beginning and ending effects created by this approach for generating time-space net supplies, we truncated the problem horizon. The first 9 weeks and the final 8 weeks were eliminated from the initial 57 weeks of data, resulting in an instance with a 40 week planning horizon. The size of the container fleet was set at three different levels: 500, 600 and 700; each level captures different capacity limits for the carrier. Initial inventories of tank-containers at each depot were determined proportional to the probability of a demand originating in its corresponding region.

All instances were solved using both **TRP1** and **TRP2**, where for the latter, reactive repositioning was only allowed between depots in the same region. Solutions were obtained using CPLEX 9.0 with default parameter values on a PC with a 1.6 GHz processor and 1Gb of memory. In all cases, fewer than 4 seconds of computation time were required to instantiate and solve an instance.

Control parameter k was varied from 0 (*i.e.*, solution to the nominal problem) to 9. When $k = 9$, the corresponding solution must be recoverable when any nine net supplies in the region simultaneously take their worst-case forecasted value. Figure 6 summarizes

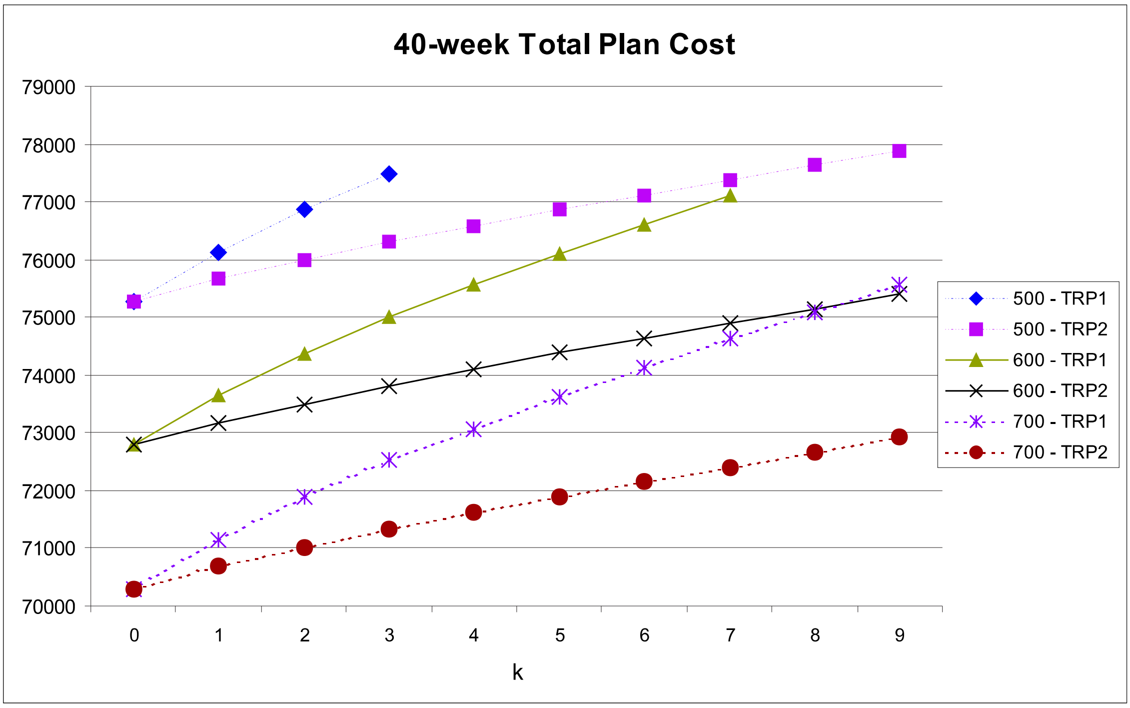

Figure 6: Total plan cost by fleet size and value of control parameter k for **TRP1** and **TRP2**.

cost results. Interestingly, for a fleet size of 600 containers the total 40-week plan cost for **TRP2** increases only 3.6% from $k = 0$ to $k = 9$. This increment corresponds to the price of robustness. A similar increment was obtained for 700 containers, and in the case of 500 containers the corresponding increment was less than 4%. This suggests that even under tight capacity, the price of robustness of the repositioning plans obtained using our approach remains within reasonable limits.

For a fleet size of 600 containers, there are no feasible solutions for **TRP1** for values of k greater than 7. The same is true for 500 containers, but for values of k greater than 3. Interestingly, feasible solutions for **TRP2** exist for all levels of k. Furthermore, observe that for $k = 9$ the cost of the optimal solution obtained for **TRP2** with 600 containers is less than the cost of the optimal solution for **TRP1** with 700 containers. These results show the power of inventory pooling to hedge against uncertainty in a cost-effective way.

Lastly, to examine how decisions change when using a robust repositioning approach, we examine the planned average inventory levels at depots that result for different levels of k for both **TRP1** and **TRP2**. To avoid any bias resulting from our choice of initial inventory locations, a 2-phase approach was employed for this analysis. In Phase I, we solve the full 40-week problem, and assume that decisions for the first 12 weeks of the plan are implemented. Based on these decisions, we update the nominal net supply values for weeks 13 through 17 (recall that the maximum transportation time is 5 weeks). Then, in Phase II, we solve a 28-week instance corresponding to weeks 13 through 40. Our inventory analysis is based on the results of the 28-week repositioning instance solved in Phase II.

Figures 7 and 8 summarize average inventory by region for different levels of k for **TRP1** and **TRP2**, respectively; note that these figures include regions that are net demanders of empty containers. Not unexpectedly, observe that in both cases, average inventory per region increases as the value of k increases. At $k = 0$, which corresponds to the nominal problem, little inventory is kept at these depots; on the other hand, for values of $k > 0$, safety-stock inventory is built by the repositioning plan to hedge against uncertainty.

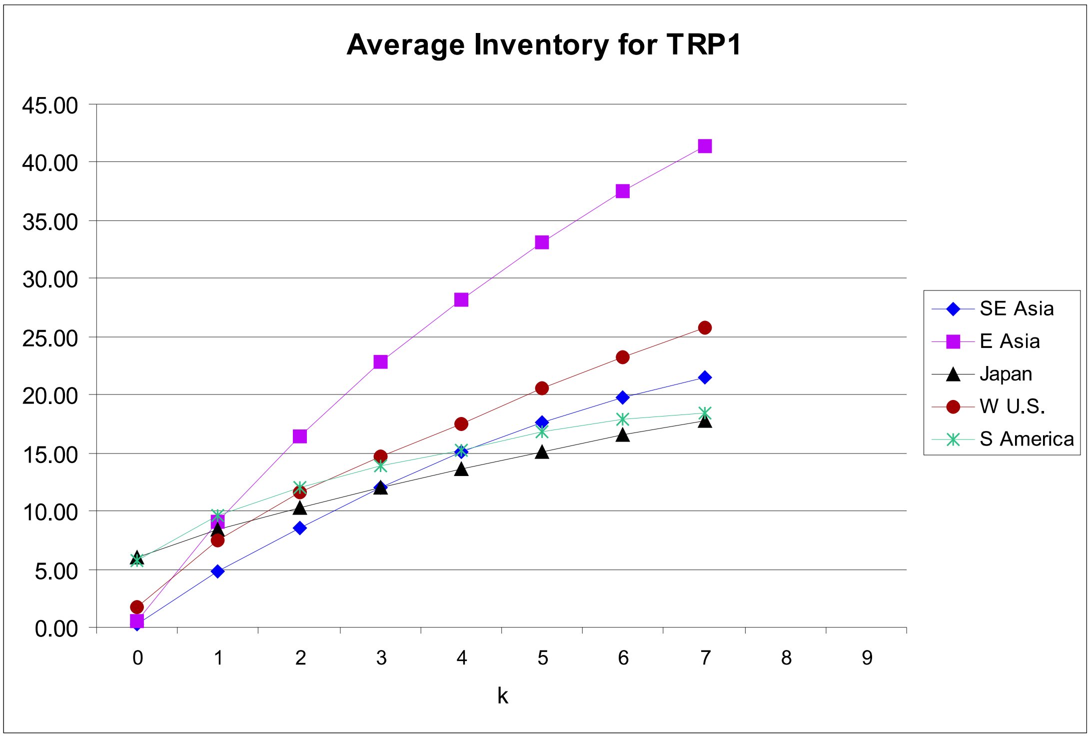

Figure 7: Average inventory per region by value of control parameter k for **TRP1** given a fleet size of 600 tanks.

The effect of inventory pooling can be observed by contrasting Figures 7 and 8. When reactive repositioning is allowed between depots in the same region, the plan is recoverable with respect to the same level of uncertainty, defined by parameter k, with far fewer tank containers of inventory per region.

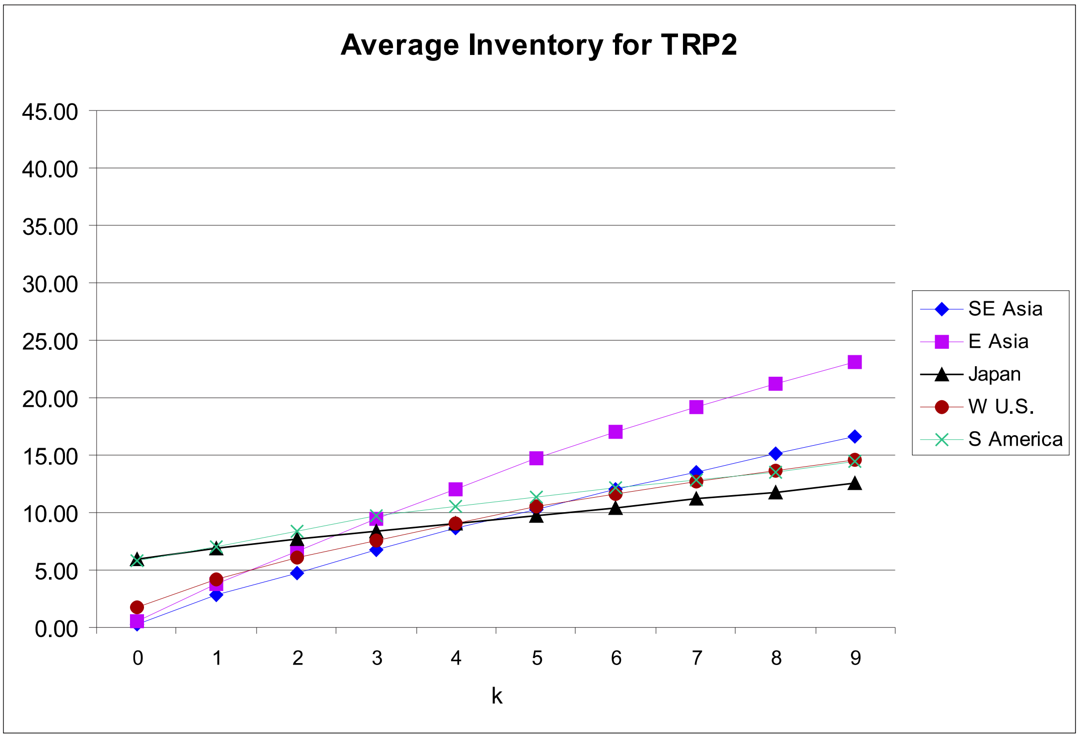

Figure 8: Average inventory per region by value of control parameter k for **TRP2** given a fleet size of 600 tanks.

CHAPTER IV

ROBUST OPTIMIZATION APPROACHES FOR THE VEHICLE ROUTING PROBLEM WITH STOCHASTIC DEMANDS

In this chapter, we study fleet management problems arising in a completely different environment. A homogeneous fleet of vehicles based at a single depot is used to serve the daily demand of a set of customers. The daily demand at a customer is not known with certainty, but falls between known lower and upper bounds. A set of fixed routes is employed to serve the customers. That is, each vehicle visits the same set of customers in the same order every day. Since customer demand is uncertain, it may happen during the execution of a route that a vehicle arrives at a customer and finds that it has insufficient product to satisfy demand. In that case, the vehicle deviates from its fixed route, restocks at the depot and resumes the route. The process of restocking at the depot is referred to as a *recourse action*. *Recourse policies* specify system conditions under which a vehicle takes a recourse action. The environment outlined above is an example of a Vehicle Routing Problem with Stochastic Demands (VRPSD). The VRPSD and its variations have been studied extensively by the research community.

In this research, we study two specific variants of the VRPSD. First, we consider the *Robust Vehicle Routing Problem with Stochastic Demands* (RVRPSD), where the objective is to identify, for a given recourse policy, a set of fixed tours such that the maximum actual travel time for any tour under any possible demand realization is minimized. The travel time of a fixed tour includes two components: the time associated with traversing the fixed tour, and the time associated with performing recourse actions. The objective function has the min-max form commonly found in robust optimization settings. Second, we consider the *Vehicle Routing Problem with Stochastic Demands and Duration Constraints* (VRPSDDC), where the objective is to minimize the total expected travel time of the fixed tours, but subject to a hard duration constraint that restricts the maximum travel time of any tour under any possible realization.

4.1 Problem definitions

The Vehicle Routing Problem with Stochastic Demands (VRPSD) is defined on a directed graph $G = (V_0, A)$ where $V = \{1, \cdots, N\}$ represents the set of customers. We use 0 to refer to the depot, therefore $V_0 = \{0\} \cup V$. $A = \{(i,j) \,|\, i \neq j,\, i,j \in V_0\}$ is the set of arcs defined over V_0. Matrix $(l(i,j))$ is defined on A and coefficient $l(i,j)$ is used to denote some cost metric; in this research we assume it represents travel time between customers i and j. It is assumed that travel times satisfy the triangle inequality: $l(i,j) \leq l(i,k) + l(k,j)$ for all i,j,k in V_0. The vehicle capacity is denoted by Q, and assume for exposition that vehicles carry inventory to be delivered to customers.

Customer demands are integer-valued random variables with known distributions, and are denoted by vector $\tilde{d} \in \mathbb{Z}_+^{|V|}$. It is assumed that vectors $\underline{d}, \overline{d} \in \mathbb{Z}_+^{|V|}$ are known, and that $\underline{d} \leq \tilde{d} \leq \overline{d}$. It is further assumed that $\overline{d}(i) \leq Q$ and that $\overline{d}(i) > 0$ for all $i \in V$. The uncertainty space considered in this problem is then denoted by

$$\mathcal{U} = \left\{ d \in \mathbb{Z}_+^{|V|} \;:\; \underline{d} \leq d \leq \overline{d} \;\right\}.$$

A tour specifies the *a priori* sequence in which some subset of customers is visited by a single vehicle; all tours start and end at the depot. Each customer is assumed to be served by one and only one vehicle. Let $\mathcal{T}_k = \{i_1, i_2, \cdots, i_{n_k}\}$ denote the tour of vehicle k, where $i_j \in V$. Observe that in this notation we do not include the depot (node 0) at the beginning and at the end of the tour.

Each time a vehicle traverses its tour, all customer demands must be satisfied; therefore, the vehicle might need to take one or more recourse actions and restock at the depot without skipping any customer in the tour. A *recourse policy*, denoted by $\mathcal{P}$, defines the system conditions that lead to vehicle recourse actions. It is assumed that for any demand realization, $\mathcal{P}$ determines uniquely when and where recourse actions will be undertaken in each tour. Let $\mathrm{L}(\mathcal{T}_k)$ be the total travel time required by vehicle k to complete its *a priori* tour if no recourse actions were required, and let $\phi(\mathcal{T}_k, \mathcal{P}, d)$ be the total additional travel time due to recourse required given policy $\mathcal{P}$ and demand realization $d \in \mathcal{U}$. Then, the maximum duration of tour k is:

$$\mathcal{L}(\mathcal{T}_k, \mathcal{P}) \;=\; \mathrm{L}(\mathcal{T}_k) \,+\, \max_{d \in \mathcal{U}} \phi(\mathcal{T}_k, \mathcal{P}, d).$$

Since the additional travel time incurred due to recourse actions depends on the demand realization, obtaining the maximum travel time that can be incurred due to recourse actions is an optimization problem. We will refer to it as the *adversarial problem* and its objective is to identify the worst-case demand realization. For convenience, we introduce the following notation:

$$\Phi(\mathcal{T}_k, \mathcal{P}) \;=\; \max_{d \in \mathcal{U}} \phi(\mathcal{T}_k, \mathcal{P}, d).$$

We now formally state the two problems considered in this chapter. Given a fleet of m vehicles, the Robust Vehicle Routing Problem with Stochastic Demand (RVRPSD) develops a set of m vehicle tours with minimum maximum travel time duration:

$$\textbf{RVRPSD} \qquad \min_{\{\mathcal{T}_k\}} \;\max_{k}\; \mathcal{L}(\mathcal{T}_k, \mathcal{P}) \tag{47}$$

Observe that no probability distribution is necessary for demand vector $\tilde{d}$ in this case, since no expectations are computed.

Before stating the second problem, we introduce the function

$$\mathcal{L}_E(\mathcal{T}_k, \mathcal{P}) \;=\; \mathrm{L}(\mathcal{T}_k) \,+\, \mathbf{E}\left[\phi(\mathcal{T}_k, \mathcal{P}, d)\right],$$

where $\mathbf{E}$ denotes the expectation operator with respect to the customer demand uncertainty space $\mathcal{U}$. $\mathbf{E}[\phi(\mathcal{T}_k, \mathcal{P}, d)]$ denotes the expected additional travel time incurred by vehicle k due to recourse actions under recourse policy $\mathcal{P}$. The second problem studied, the Vehicle Routing Problem with Stochastic Demands and Duration Constraints (VRPSDDC), finds the set of tours with minimum total expected duration subject to a hard constraint on individual vehicle duration:

$$\begin{aligned} \textbf{VRPSDDC} \qquad & \min_{\{\mathcal{T}_k\}} && \sum_k \mathcal{L}_E(\mathcal{T}_k, \mathcal{P}) \\ & \text{s.t.} && \mathrm{L}(\mathcal{T}_k) + \Phi(\mathcal{T}_k, \mathcal{P}) \;\leq\; D \quad \forall\, k, \end{aligned} \tag{48}$$

where D denotes the maximum travel time duration allowed for a tour.

4.2 The adversarial problem

Consider now a single vehicle and its tour $\mathcal{T} = \{1, 2, \cdots, n\}$ of some subset of the customers. The adversarial problem corresponds to determining the demand realization $d \in \mathcal{U}$ that maximizes $\phi(\mathcal{T}, \mathcal{P}, d)$. If function ϕ is non-decreasing in $d(i)$ for $\mathcal{P}$ and for all $i \in \mathcal{T}$, it is clear that an optimal solution to this problem is to set d equal to $\overline{d}$. As will be shown later, this is not the case for all recourse policies; that is, unlike most robust optimization applications, the worst scenario in this problem is not always found at an extreme point of the box that defines the uncertainty space. Therefore, we investigate solution approaches for the adversarial problem. Observe that for a given tour $\mathcal{T}$, the size of the uncertainty space is given by

$$\prod_{i \in \mathcal{T}} \left(\overline{d}(i) - \underline{d}(i) + 1 \right),$$

which rules out enumerative approaches.

All recourse policies considered in this research are assumed to have the following characteristics:

1. For any given demand realization, the customers at which recourse actions are triggered are uniquely determined.

2. The number of recourse actions is a non-decreasing function of the demand $d(i)$ for all customers i in the tour.

The process resulting from operating tour $\mathcal{T}$ under policy $\mathcal{P}$ can be described in terms of state variables (r, i, I_i) defined for every customer i in the tour, where $r = k$ if the k^{th} recourse action is triggered by customer i and $r = 0$ if no recourse action is triggered by this customer, and I_i denotes the load of the vehicle when it departs from customer i after completing delivery. Therefore, the state space for each customer $i \in \mathcal{T}$ is $0 \leq I_i \leq Q$ and $i \leq r \leq R$, where R denotes the maximum number of recourse actions that can occur. For the policies we consider, the value of R can be calculated by evaluating the tour using $\mathcal{P}$ for demand realization $\overline{d}$; R is also clearly bounded from above by n, the number of customers in the tour. Note also that by earlier assumption, recourse policies have the property that for a given demand realization, the values of the state variables are uniquely determined.

To solve the adversarial problem for a given tour $\mathcal{T}$, a given recourse policy $\mathcal{P}$, and demand uncertainty set $\mathcal{U}$, we characterize the set of recourse-triggering states S for which there exists a demand realization $d \in \mathcal{U}$ that leads to that state. Consider, for example, tour $\mathcal{T}$ operated using recourse policy $\mathcal{P}$. Because $\underline{d}$ and $\overline{d}$ are known, it is possible to determine the set of customers at which the first recourse action may be triggered, and for each such customer a corresponding demand realization(s) and the resulting vehicle load(s) at the departure from that customer. Given tour $\mathcal{T}$ and recourse policy $\mathcal{P}$, let $C^1(i, I_i)$ for $i \in \mathcal{T}$ denote the set of necessary and sufficient conditions (on $\underline{d}$ and $\overline{d}$) that ensures existence of a demand realization $d \in \mathcal{U}$ for which the *first* recourse action is triggered for customer i and such that the vehicle load at the departure from customer i is I_i. Furthermore, let $C^{r,r+1}(i, I_i, k, I_k)$ for all $i, k \in \mathcal{T}$ such that $i < k$ and for $r = 1, \cdots, R-1$ denote the set of necessary and sufficient conditions that ensures existence of a demand realization $d \in \mathcal{U}$ for which the $(r+1)^{th}$ recourse action is triggered for customer k with vehicle load I_k at departure *given* that the r^{th} recourse action is triggered at customer i with vehicle load I_i at departure. Conditions $C^1(i, I_i)$ and $C^{r,r+1}(i, I_i, k, I_k)$ are referred to as *recourse conditions*.

The recourse conditions are used in the construction of an acyclic digraph $\mathcal{G}(\mathcal{P}) = (\mathcal{N}, \mathcal{A})$. The set of nodes of $\mathcal{G}(\mathcal{P})$ is constructed as follows

$$\mathcal{N} = \{s\} \cup \{t\} \cup \big\{(r, i, I_i) \mid i \in \mathcal{T},\ r \in \{i, \cdots, R\}, I_i \in \{0, 1, \cdots, Q\}\big\}.$$

The set of arcs is constructed as follows

1. $(s, (1, i, I_i)) \in \mathcal{A}$ for $i \in \mathcal{T}$ and I_i that satisfies $C^1(i, I_i)$. The cost of the arc is the increase in travel time incurred when a recourse action is triggered by customer i. Each arc is associated with a demand realization for which the first recourse is triggered by i yielding corresponding load I_i.

2. For $r = 1$ to $R-1$, $((r, i, I_i), (r+1, k, I_k)) \in \mathcal{A}$ for $i, k \in \mathcal{T}$, $i < k$, and I_i, I_k that satisfy $C^{r,r+1}(i, I_i, k, I_k)$. The cost of the arc is the increase in travel time incurred when a recourse action is triggered for customer k. Each arc is associated with a demand

realization for which the r^{th} recourse action is triggered by customer i yielding vehicle load I_i at departure and the $(r+1)^{th}$ recourse action is triggered by customer k yielding vehicle load I_k at departure.

3. For all $(r, i, I_i) \in \mathcal{A}$ such that $indeg(r, i, I_i) > 0$ and the $outdeg(r, i, I_i) = 0$, $((r, i, I_i), t) \in \mathcal{A}$ with cost 0.

Figure 9 shows digraph $\mathcal{G}(\mathcal{P}) = (\mathcal{N}, \mathcal{A})$ for an instance with $Q = 1$ and $n = R = 4$.

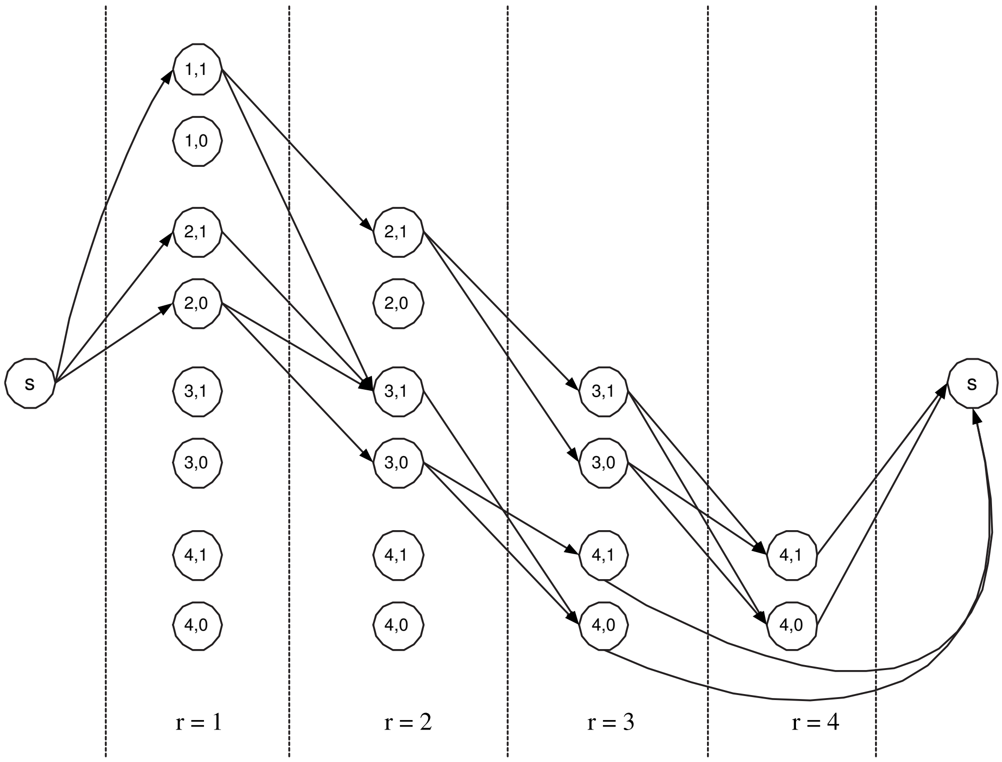

Figure 9: $\mathcal{G}(\mathcal{P})$ for an instance with $Q = 1$ and $n = R = 4$, with (i, I_i) in each node.

Let $L(\mathcal{G}(\mathcal{P}))$ be the longest $s-t$ path in the graph. If there does not exist an $s-t$ path, then by definition $L(\mathcal{G}(\mathcal{P})) = 0$.

Lemma 1 *For a tour $\mathcal{T}$ operated using recourse policy $\mathcal{P}$, $\Phi(\mathcal{T}, \mathcal{P}) = L(\mathcal{G}(\mathcal{P}))$.*

Proof. It is assumed that for any demand realization d there is a unique set of customers on $\mathcal{T}$ where recourse actions take place. From the construction of network $\mathcal{G}(\mathcal{P})$, it is clear that every $d \in \mathcal{U}$ is associated with one and only one $s-t$ path in $\mathcal{G}(\mathcal{P})$, and also, every $s-t$ path in the network is associated with one and only one demand realization. In order to evaluate $\Phi(\mathcal{T}, \mathcal{P})$, the longest path on the network needs to be identified because this is associated to a demand realization with maximum additional travel time; hence, $\Phi(\mathcal{T}, \mathcal{P}) = L(\mathcal{G}(\mathcal{P}))$. $\square$

Lemma 2 *$L(\mathcal{G}(\mathcal{P}))$ can be calculated in $O(n^3Q^2)$.*

Proof. Evaluating $L(\mathcal{G}(\mathcal{P}))$ can be done in $O(|\mathcal{A}|)$ because $\mathcal{G}(\mathcal{P})$ is an acyclic network. Observe that instances such that $\underline{d}(i) = 0$ and $\overline{d}(i) = Q$ for all $i \in V$ have the largest value for $|\mathcal{A}|$ because this gives the highest flexibility to the adversary. For this case, the number of nodes (i, r_i, I_i) with $r_i = r$ is clearly bounded by $(n - r + 1)(Q + 1)$. Take a node (i, r, I_i) in $\mathcal{N}$, observe that it can be connected directly with any other node $(k, r + 1, I_k)$ such that $i < k$; hence, the number of such arcs with tail (i, r, I_i) is $O((n - i)Q)$, and for a fixed i, there can be $Q + 1$ of such nodes in the network, therefore the number of arcs with tail (i, r, I_i) for $I_i \in \{0, \cdots, Q\}$ is $O((n - i)Q^2)$. For a given r, i is bounded between r and R, so the total number of arcs with tail (i, r, I_i) for all i and I_i is $O(\sum_{i=r}^{R}(n - i)Q^2)$. Summing over r, the number of arcs connecting the nodes associated with the r^{th} and $(r + 1)^{th}$ recourse actions for all r is then $O(\sum_{r=1}^{R-1}\sum_{i=r}^{R}(n - i)Q^2)$, and since R is bounded by n, after some algebra this expression reduces to $O(n^3Q^2)$. The number of nodes connected to both s and t is bounded by nQ in each case, so the total number of arcs in $\mathcal{G}(\mathcal{P})$ is $O(n^3Q^2)$.$\square$

The previous lemmas imply that the adversarial problem can be solved in pseudopolynomial time for recourse policies for which the corresponding recourse conditions can also be evaluated in at most pseudopolynomial time.

Definition 8 (History independent recourse policy) *A recourse policy $\mathcal{P}$ is history independent, if for every $i, k \in \mathcal{T}$ such that $i < k$, recourse conditions $C^{r,r+1}$ can be evaluated independent of (or without considering) $d(1), \cdots, d(i - 1)$, for all $r = 1, \cdots, R - 1$.*

For history independent recourse policies, the corresponding state space does not need to include on-board inventory at departure (I_i) after a customer is served. Policies in which the on-board inventory needs to be considered in order to determine if there exists a demand realization $d \in \mathcal{U}$ for which the $(r+1)^{th}$ recourse can occur at customer k *given* that the r^{th} recourse occurred at customer i, are not history independent because on-board inventory is a function of $d(1), \cdots, d(i-1)$, together with $d(i)$. Recourse conditions for history independent policies take the form $C^1(i)$ and $C^{r,r+1}(i,k)$ for all customers i and k in the tour.

For history independent recourse policies the adversarial problem can be solved using a simpler directed graph $\mathcal{G}_0(\mathcal{P}) = (\mathcal{N}, \mathcal{A})$, which is now defined and illustrated for $n = 4$ in Figure 10. The set of nodes is defined as

$$\mathcal{N} = \{s\} \cup \{t\} \cup \{(r,i) \mid i \in \mathcal{T},\, r \in \{i, \cdots, R\}\}$$

The arc set $\mathcal{A}$ is defined as follows:

1. $(s,(1,i)) \in \mathcal{A}$ for every $i \in \mathcal{T}$ that satisfies $C^1(i)$. The cost of each arc is the net increment in travel time incurred when making a recourse action for customer i. Each arc is associated with a demand realization for which the first recourse is taken for i.

2. For $r = 1$ to $R-1$, $((r,i),(r+1,k)) \in \mathcal{A}$ for $i,k \in \mathcal{T}$ $(i < k)$ that satisfy $C^{r,r+1}(i,k)$. The cost of each arc is the net increment in travel time incurred when making a recourse action for customer k. Each arc is associated with a demand realization for which the r^{th} recourse action is taken for i and the $(r+1)^th$ recourse action is taken for k.

3. For all $(r,i) \in \mathcal{A}$ such that $indeg(r,i) > 0$ and the $outdeg(r,i) = 0$, then $((r,i),t) \in \mathcal{A}$ with cost 0.

Let $L(\mathcal{G}_0(\mathcal{P}))$ be the longest $s-t$ path in the network. If there does not exist an $s-t$ path, then by definition $L(\mathcal{G}_0(\mathcal{P})) = 0$.

Lemma 3 *Consider a fixed sequence $\mathcal{T}$ operated using recourse policy $\mathcal{P}$. If $\mathcal{P}$ is history independent, then $\Phi(\mathcal{T},\mathcal{P}) = L(\mathcal{G}_0(\mathcal{P}))$.*

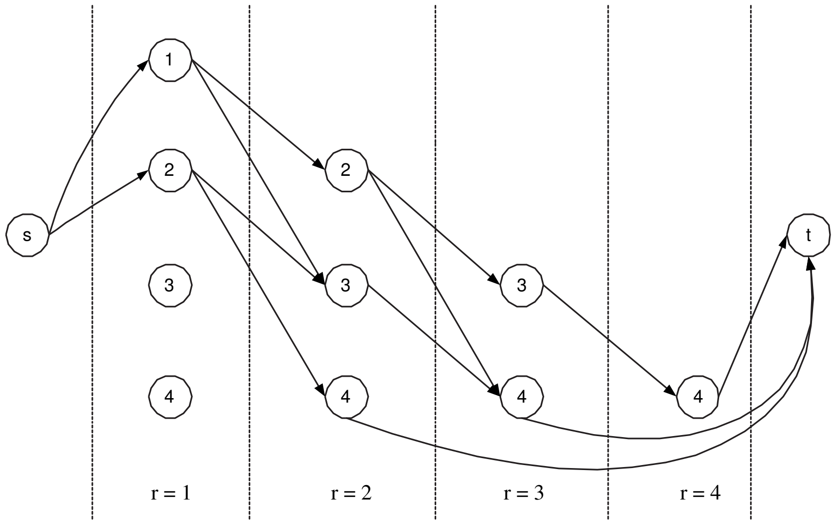

Figure 10: Illustration of $\mathcal{G}_0(\mathcal{P})$ for an instance with $n = R = 4$.

Proof. From the construction of network $\mathcal{G}_0(\mathcal{P})$, it is clear that every $d \in \mathcal{U}$ is associated with one and only one $s-t$ path in $\mathcal{G}_0(\mathcal{P})$, and that every $s-t$ path in the network is associated with at least one demand realization in the uncertainty space. In order to evaluate $\Phi(\mathcal{T}, \mathcal{P})$, the longest path on the network needs to be identified because this path is associated with a demand realization with maximum additional travel time; hence, $\Phi(\mathcal{T}, \mathcal{P}) = L(\mathcal{G}_0(\mathcal{P}))$. □

Lemma 4 *$L(\mathcal{G}_0(\mathcal{P}))$ can be calculated in $O(n^3)$.*

Proof. Similar to the proof of Lemma 2. □

History independent policies are attractive because by the previous lemmas, their adversarial problem can be solved in polynomial time provided that conditions $C^1(i)$ and $C^{r,r+1}(i,k)$ can be evaluated in polynomial time.

4.3 Recourse policies

In this section, a number of recourse policies are defined and their recourse conditions derived. In this context, it is important to consider how demand information is revealed over time. Three cases are considered:

1. *No advance demand information*: demand is revealed only after the vehicle arrives at a customer location.

2. *Total demand visibility*: demand for all customers is known before the vehicle departs from the depot for the first time.

3. *One customer ahead demand visibility*: before departing from a customer's location, but after the customer is served, the demand of the next customer in the tour is revealed.

All policies studied here can be classified according to the following categories:

1. *Splitting vs non-splitting policies*: in a splitting policy, when the vehicle arrives at a location and the demand is larger than the on-board inventory, the vehicle delivers the on-board inventory, restocks at the depot and delivers the remaining demand. In a non-splitting policy, when the vehicle arrives at a location and the demand is larger than the on-board inventory, the vehicle restocks at the depot before satisfying any part of the customer demand.

2. *Anticipatory vs non-anticipatory (myopic) policies*: recourse actions are taken at customer locations. The information used to decide if such an action should be taken determines whether a policy is classified as anticipatory or non-anticipatory. The term myopic is also used to refer to non-anticipatory policies. In non-anticipatory policies recourse actions are taken based only on local information; namely, the observed customer demand and the on-board inventory. Anticipatory policies also make use of additional information available to the decision maker, such as travel times and demand intervals $(\underline{d}, \overline{d})$ of the customers that are ahead in the tour to decide when a recourse action takes place.

4.3.1 No advance demand information

In this section, it is assumed that a customer demand is known only after the vehicle arrives at the corresponding customer location. Before proceeding consider the following useful lemma.

Lemma 5 *Consider vectors* $\underline{d}, \overline{d} \in \mathbb{Z}^n_+$ *such that* $\underline{d} \leq \overline{d}$ *and* $\overline{d} > \mathbf{0}$*. If*

$$\sum_{\ell=1}^{n-1} \underline{d}(\ell) \;\leq\; M \quad \text{ and } \quad \sum_{\ell=1}^{n} \overline{d}(\ell) \;\geq\; M+1,$$

where $M < \infty$ *is an integer constant, then there exists a vector* $d \in \mathbb{Z}^n_+$ *such that* $\underline{d} \leq d \leq \overline{d}$ *that satisfies*

$$\sum_{\ell=1}^{n-1} d(\ell) \;\leq\; M \quad \text{ and } \quad \sum_{\ell=1}^{n} d(\ell) \;\geq\; M+1$$

Proof. A constructive proof is provided. Initialize d as $d(\ell) = \underline{d}(\ell)$ for $\ell = 1, \cdots, n-1$ and $d(n) = \overline{d}(n)$. Clearly $\underline{d} \leq d \leq \overline{d}$ and $\sum_{\ell=1}^{n-1} d(\ell) \leq M$. If such d also satisfies $\sum_{\ell=1}^{n} d(\ell) \geq M+1$, then the desired vector is obtained. Else, since $\overline{d} > \mathbf{0}$ it follows that $\sum_{\ell=1}^{n-1} d(\ell) < M$ and that

$$\sum_{\ell=1}^{n-1} d(\ell) \;<\; \sum_{\ell=1}^{n} d(\ell). \tag{49}$$

Take any $d(\ell)$ such that $d(\ell) < \overline{d}(\ell)$ for $1 \leq \ell \leq n-1$ and perform the following update: $d(\ell) \leftarrow d(\ell) + 1$. Such an ℓ exists, otherwise $\sum_{\ell=1}^{n} \overline{d}(\ell) \geq M+1$ would be contradicted. After such an update, inequality (49) remains true, and furthermore $\underline{d} \leq d \leq \overline{d}$ and $\sum_{\ell=1}^{n-1} d(\ell) \leq M$. Therefore, if $\sum_{\ell=1}^{n} d(\ell) \geq M+1$ stop with the desired d; else, take any $d(\ell)$ such that $d(\ell) < \overline{d}(\ell)$ for $1 \leq \ell \leq n-1$ and repeat the process described in this paragraph. At every step the difference between $M+1$ and $\sum_{\ell=1}^{n} d(\ell)$ is reduced by 1, and condition $\sum_{\ell=1}^{n-1} d(\ell) \leq M$ remains true; hence, in a finite number of steps the desired vector is obtained. □

4.3.1.1 A non-anticipatory and splitting policy

Definition 9 $\mathcal{P}_0^s$ *is used to denote the following non-anticipatory and splitting recourse policy on a fixed sequence* $\mathcal{T}$*: a recourse action is triggered by customer* $i \in \mathcal{T}$ *if and only if when the vehicle arrives at* i*, it observes that* $d(i)$ *is strictly greater than on-board inventory. After satisfying as much of* $d(i)$ *as possible, the vehicle restocks at the depot, then returns to* i *and finishes satisfying the rest of the demand before proceeding.*

Recourse conditions for this policy are now derived. Observe that for this recourse policy, whenever a recourse action occurs the on-board inventory after the customer is served is always strictly less than Q. Also, a recourse action will never be taken at customer 1 because by assumption $d(i) \leq Q$ for all customers.

Proposition 5 (Recourse condition $C^{r,r+1}(i, I_i, k, I_k)$ for $\mathcal{P}_0^s$) *Consider a fixed sequence $\mathcal{T} = \{1, \cdots, n\}$ operated using recourse policy $\mathcal{P}_0^s$. Assume there exists a demand realization $d \in \mathcal{U}$ such that the r^{th} recourse occurs at customer $i \in \mathcal{T}$ leaving onboard inventory I_i. The $(r+1)^{th}$ recourse can occur for customer k $(k > i)$ with corresponding I_k if and only if*

$$\sum_{\ell=i+1}^{k-1} \underline{d}(\ell) \leq I_i \quad \text{and} \quad \sum_{\ell=i+1}^{k} \overline{d}(\ell) \geq I_i + 1;$$

for $I_k \in \{0, 1, \cdots, Q-1\}$ further bounded by expression

$$Q + I_i - \min\left\{I_i, \sum_{\ell=i+1}^{k-1} \overline{d}(\ell)\right\} - \overline{d}(k) \leq I_k \leq Q + I_i - \sum_{\ell=i+1}^{k} \underline{d}(\ell).$$

Proof.

$(\Rightarrow)$ A demand realization d that has the $(r+1)^{th}$ recourse action for customer k with I_k *given* that the r^{th} recourse action occurred at customer i with I_i must satisfy

$$\sum_{\ell=i+1}^{k-1} d(\ell) \leq I_i \quad \text{and} \quad \sum_{\ell=i+1}^{k} d(\ell) \geq I_i + 1;$$

with corresponding

$$I_k = I_i - \sum_{\ell=i+1}^{k} d(\ell) + Q.$$

Since $d \in \mathcal{U}$, $\underline{d}(\ell) \leq d(\ell)$ and $d(\ell) \leq \overline{d}(\ell)$ for all ℓ; summing over ℓ, the above implies

$$\sum_{\ell=i+1}^{k-1} \underline{d}(\ell) \leq \sum_{\ell=i+1}^{k-1} d(\ell) \leq I_i \quad \text{and} \quad \sum_{\ell=i+1}^{k} \overline{d}(\ell) \geq \sum_{\ell=i+1}^{k} d(\ell) \geq I_i + 1;$$

it also implies that

$$Q + I_i - \sum_{\ell=i+1}^{k} \overline{d}(\ell) \leq I_k \leq Q + I_i - \sum_{\ell=i+1}^{k} \underline{d}(\ell),$$

but since $\sum_{\ell=i+1}^{k-1} d(\ell) \leq I_i$, it follows that

$$Q + I_i - \min\left\{I_i, \sum_{\ell=i+1}^{k-1} \overline{d}(\ell)\right\} - \overline{d}(k) \leq I_k \leq Q + I_i - \sum_{\ell=i+1}^{k} \underline{d}(\ell).$$

($\Leftarrow$) The existence of a demand realization in $\mathcal{U}$ such that the $r+1$ recourse takes place for k given that r^{th} takes place for i with on-board inventory I_i follows directly from Lemma 5; I_k just needs to be bounded accordingly. $\square$

When the vehicle departs from the depot for the first time, system conditions are equivalent to having a recourse at customer 0 (*i.e.*, the depot) with $I_0 = Q$; the following result then follows directly from Proposition 5.

Proposition 6 (Recourse condition $C^1(i, I_i)$ for $\mathcal{P}_0^s$) *Consider a fixed sequence $\mathcal{T} = \{1, \cdots, n\}$ operated using recourse policy $\mathcal{P}_0^s$. Recourse conditions $C^1(i, I_i)$ are equivalent to $C^{r,r+1}(0, Q, i, I_i)$.*

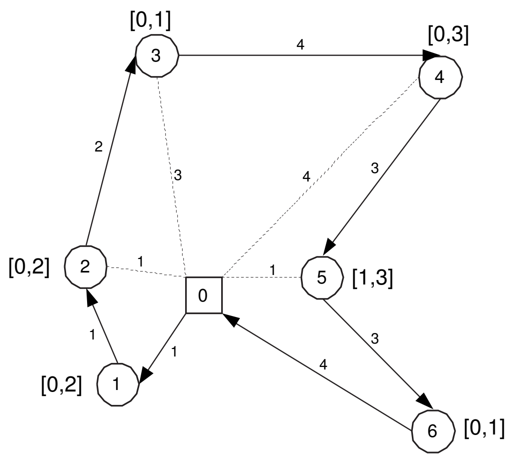

Figure 11: Tour operated using $\mathcal{P}_0^s$ and $Q = 3$.

Example. Consider the six customer tour in Figure 11, where relevant transit times are included on the corresponding arc. Assume that the tour is operated using recourse policy $\mathcal{P}_0^s$ with a vehicle of capacity $Q = 3$. Observe that $\mathrm{L}(\mathcal{T}) = 18$ and for demand realization $\overline{d}$, Table 8 shows the corresponding values that the state variables take under this realization. From the table it can be seen that $R = 3$ and $\phi(\mathcal{T}, \mathcal{P}_0^s, \overline{d}) = 12$.

The adversarial problem is solved using network $\mathcal{G}(\mathcal{P}_0^s)$, which for this example is shown in Figure 12. The longest $s-t$ path in the network corresponds to $\{s, (1,3,2), (2,4,2), (3,6,2), t\}$, which is associated with recourse actions taken at customers 3, 4 and 6, with cost 22. Observe then that the maximum net additional travel time due to recourse actions is not

Table 8: Values of the state variables for demand realization $\overline{d}$ when the tour is operated using recourse policy $\mathcal{P}_0^s$.

i	1	2	3	4	5	6
$d(i)$	2	2	1	3	3	1
r	0	1	0	2	3	0
I_i	1	2	1	1	1	0

achieved with demand realization $\overline{d}$.

The demand realizations $d^\star$ for which $\phi(\mathcal{T}, \mathcal{P}_0^s, d)$ achieves optimal cost 22, are $[2,1,1,3,2,1]$ and $[1,2,1,3,2,1]$, neither of which corresponds to extreme points of the uncertainty space $\mathcal{U}$. Table 9 shows the values of the state variables for these realizations. Observe that for both of these demand outcomes the state variables for customer 4 take values $r = 2$ and $I_4 = 2$, so that if the demand for customer 5 were to take value $\overline{d}(5) = 3$, then the third recourse would take place at this customer, no more recourse actions could take place, with a resulting net increment equal to 16. On the other hand, if demand for customer 5 were to take value $\underline{d}(5) = 1$, no additional recourse actions are needed so the resulting net increment would equal 14.

Table 9: Values of the state variables for demand realization $[2,1,1,3,2,1]$ and $[1,2,1,3,2,1]$ respectively when the tour is operated using recourse policy $\mathcal{P}_0^s$.

i	1	2	3	4	5	6	i	1	2	3	4	5	6
$d(i)$	2	1	1	3	2	1	$d(i)$	1	2	1	3	2	1
r	0	0	1	2	0	1	r	0	0	1	2	0	1
I_i	1	0	2	2	0	2	I_i	2	0	2	2	0	2

From Proposition 5, it follows that $\mathcal{P}_0^s$ is not a history independent recourse policy because I_i is a function of $d(1), \cdots, d(i-1)$. In order to solve the adversarial problem the longest $s - t$ path on $\mathcal{G}(\mathcal{P}_0^s)$ needs to be identified. From Lemmas 1 and 2 it follows that this can be done on pseudopolynomial time. Interestingly, a special case of $\mathcal{P}_0^s$ remains polynomially solvable.

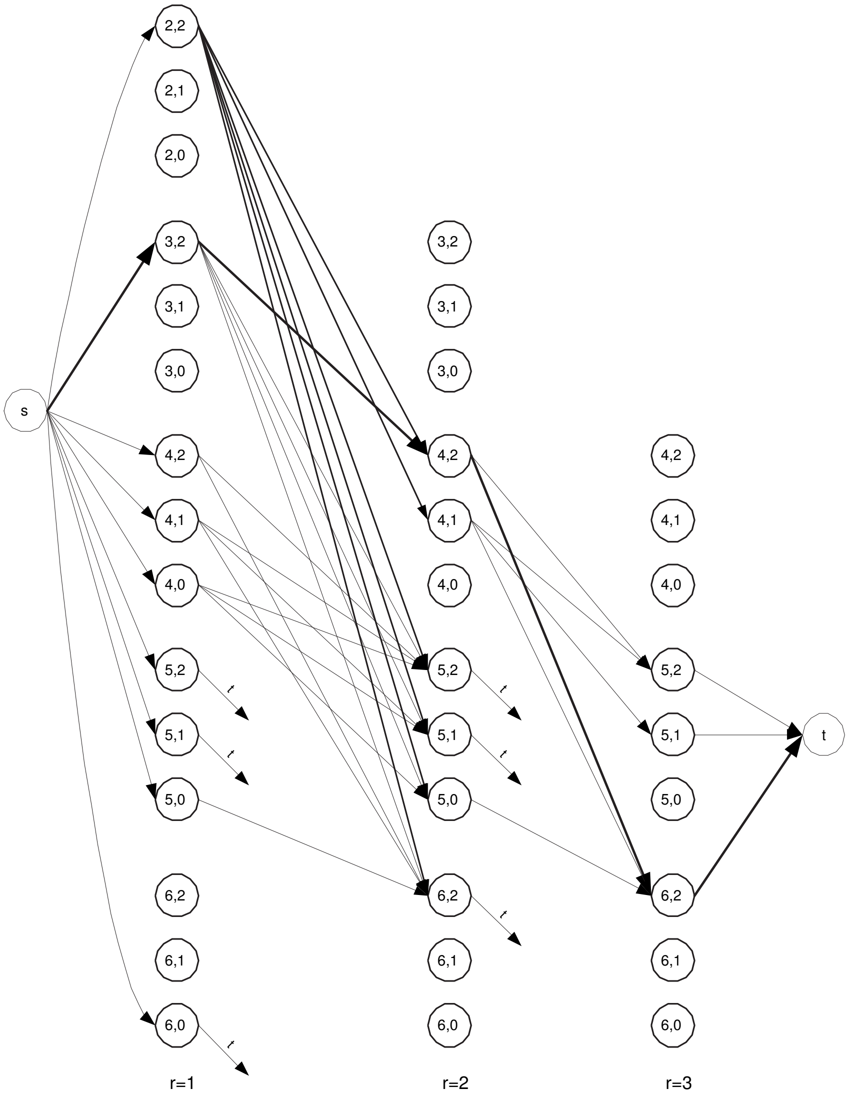

Figure 12: Network $\mathcal{G}(\mathcal{P}_0^s)$.

Proposition 7 *Consider a fixed sequence $\mathcal{T} = \{1, \cdots, n\}$ operated using recourse policy $\mathcal{P}_0^s$. For instances such that $R \leq 2$, there exists a demand realization $d \in \mathcal{U}$ such that the first recourse occurs for $i \in \mathcal{T}$ and the second recourse occurs for $k \in \mathcal{T}$ such that $i < k$ if and only if*

$$\sum_{\ell=1}^{i-1} \underline{d}(\ell) \leq Q \quad \textit{and} \quad \sum_{\ell=1}^{i} \overline{d}(\ell) \geq Q + 1;$$

$$\max\left\{1, \sum_{\ell=1}^{i} \underline{d}(\ell) - Q\right\} + \sum_{\ell=i+1}^{k-1} \underline{d}(\ell) \leq Q \;\; \textit{and} \;\; \min\left\{\overline{d}(i), \sum_{\ell=1}^{i} \overline{d}(\ell) - Q\right\} + \sum_{\ell=i+1}^{k} \overline{d}(\ell) \geq Q+1.$$

Proof. First recourse for i follows from Proposition 6. To show the second part, let Δ_i be the part of the demand at customer i that was satisfied after the restock.

($\Rightarrow$) Such a demand realization d must satisfy

$$\Delta_i + \sum_{\ell=i+1}^{k-1} d(\ell) \le Q \quad \text{and} \quad \Delta_i + \sum_{\ell=i+1}^{k} d(\ell) \ge Q+1$$

Observe that $\Delta_i \ge 1$ in order for a recourse to have occurred at i. Also, since $d \in \mathcal{U}$, $\underline{d} \le d \le \overline{d}$; further, $\Delta_i \ge \sum_{\ell=1}^{i} \underline{d}(\ell) - Q$, $\Delta_i \le \overline{d}(i)$ and $\Delta_i \le \sum_{\ell=1}^{i} \overline{d}(\ell) - Q$. This, and the inequalities above imply

$$\max\left\{1, \sum_{\ell=1}^{i} \underline{d}(\ell) - Q\right\} + \sum_{\ell=i+1}^{k-1} \underline{d}(\ell) \le \Delta_i + \sum_{\ell=i+1}^{k-1} d(\ell) \le Q \quad \text{and}$$

$$\min\left\{\overline{d}(i), \sum_{\ell=1}^{i} \overline{d}(\ell) - Q\right\} + \sum_{\ell=i+1}^{k} \overline{d}(\ell) \ge \Delta_i + \sum_{\ell=i+1}^{k} d(\ell) \ge Q+1$$

($\Leftarrow$) Follows from Lemma 5 making $n = k - i + 1$ and $M = Q$; where $d(1)$ in the lemma is associated with Δ_i, $d(n)$ in the lemma is associated with $d(k)$ and shifting all other entries accordingly; further making $\underline{d}(1) = \max\{1, \sum_{\ell=1}^{i} \underline{d}(\ell) - Q\}$ and $\overline{d}(1) = \min\{\overline{d}(i), \sum_{\ell=1}^{i} \overline{d}(\ell) - Q\}$ in the lemma. □

The following theorem summarizes the results of this section.

Theorem 4 *Consider a fixed sequence $\mathcal{T} = \{1, \cdots, n\}$ operated using recourse policy $\mathcal{P}_0^s$. Instances for which $R \le 2$, $\Phi(\mathcal{T}, \mathcal{P}_0^s)$ can be evaluated in polynomial time by solving a longest s-t path problem on $\mathcal{G}_0(\mathcal{P}_0^s)$. In general, $\Phi(\mathcal{T}, \mathcal{P}_0^s)$ can be evaluated in pseudopolynomial time by solving a longest s-t path problem on $\mathcal{G}(\mathcal{P}_0^s)$.*

Proof. Proposition 7 implies that for instances with $R \le 2$, although $\mathcal{P}_0^s$ is not in general a history independent recourse policy, $\Phi(\mathcal{T}, \mathcal{P}_0^s) = L(\mathcal{G}_0(\mathcal{P}_0^s))$ because the recourse conditions can be evaluated without explicitly considering on-board inventory after the first recourse, hence lemmas 3 and 4 apply in this case. Building network $\mathcal{G}_0(\mathcal{P}_0^s)$ is also done in polynomial time because for any arc in this network, conditions in Proposition 7 can be evaluated in polynomial time.

The same argument can be used for the general case to show that $\Phi(\mathcal{T}, \mathcal{P}_0^s)$ can be evaluated in pseudopolynomial time, by means of lemmas 1 and 2, applying propositions 5 and 6 on network $\mathcal{G}(\mathcal{P}_0^s)$. $\square$

4.3.1.2 A non-anticipatory and non-splitting policy

Definition 10 $\mathcal{P}_0^n$ *is used to denote the following non-anticipatory and non-splitting recourse policy on a fixed sequence* $\mathcal{T}$*: a recourse action is taken for customer* $i \in \mathcal{T}$ *if and only if when the vehicle arrives at* i *it observes that* $d(i)$ *is strictly greater than the on-board inventory. Before satisfying any part of* $d(i)$*, the vehicle restocks at the depot returns to* i *and satisfies the entire demand.*

Under this policy since $d(i) \leq Q$ a customer's demand is never split between two vehicle restocks. Also, the first recourse will never take place at the first customer in the tour. The following observation is useful.

Observation 1 *Consider a fixed sequence* $\mathcal{T} = \{1, \cdots, n\}$ *operated using recourse policy* $\mathcal{P}_0^n$*, and assume a demand realization* d *such that there is a recourse at customer* $i \in \mathcal{T}$*; then, after* i *is served the on-board inventory is always* $Q - d(i)$*.*

Based on the previous observation, $\mathcal{P}_0^n$ seems to be a history independent recourse policy because after a recourse action is taken, say at customer i, it is just as if the vehicle had left the depot and i was the first customer in the tour. It seems like the system status is independent of the events that occurred before a restock, and therefore $\Phi(\mathcal{T}, \mathcal{P}_0^n)$ could be evaluated in polynomial time (see Lemma 4). Unfortunately, this is not true, consider the following result.

Lemma 6 *Consider a fixed sequence* $\mathcal{T} = \{1, \cdots, n\}$ *operated using recourse policy* $\mathcal{P}_0^n$*, and assume a demand realization* $d \in \mathcal{U}$ *such that the* $(r-1)^{th}$ *recourse action occurred for* $j \in \mathcal{T}$*. If the* r^{th} *recourse action occurs for* $i \in \mathcal{T}$ $(j < i)$ *then*

$$d(i) \;\geq\; \underline{d}(i/j) \;\equiv\; \max\left\{1, Q + 1 - \sum_{\ell=j}^{i-1} \overline{d}(\ell)\right\}.$$

Proof. For $\mathcal{P}_0^n$ such demand realization d must satisfy

$$\sum_{\ell=j}^{i-1} d(\ell) \leq Q \tag{50}$$

$$\sum_{\ell=j}^{i} d(\ell) \geq Q+1 \tag{51}$$

Since $d \in \mathcal{U}$ then $\underline{d}(\ell) \leq d(\ell) \leq \overline{d}(\ell)$ and summing over ℓ

$$\sum_{\ell=j}^{i-1} d(\ell) \leq \sum_{\ell=j}^{i-1} \overline{d}(\ell)$$

The above expression together with (50) implies

$$\sum_{\ell=j}^{i-1} d(\ell) \leq \min\left\{Q, \sum_{\ell=j}^{i-1} \overline{d}(\ell)\right\}$$

From (51)

$$\begin{aligned} d(i) &\geq Q+1-\sum_{\ell=j}^{i-1} d(\ell) \\ &\geq Q+1-\min\left\{Q, \sum_{\ell=j}^{i-1} \overline{d}(\ell)\right\} \\ &\geq \max\left\{1, Q+1-\sum_{\ell=j}^{i-1} \overline{d}(\ell)\right\} \end{aligned}$$

□

The previous result implies that $\mathcal{P}_0^n$ is not a history independent recourse policy because to determine if there exists a demand realization in the uncertainty set where the $(r+1)^{th}$ recourse occurs for k, given that the r^{th} recourse occurred for i, information provided by $d(1), \cdots, d(i-1)$ needs to be considered in order to determine for which customer the $(r-1)^{th}$ recourse action actually occurred. This implies that one way in which the adversarial problem can be solved for $\mathcal{P}_0^n$ is to derive recourse conditions that are a function of I_i and I_k (like the conditions in Proposition 5) and solve a longest $s-t$ path problem on network $\mathcal{G}$, which by lemmas 1 and 2 would be done in pseudopolynomial time.

An alternative way to address the adversarial problem for $\mathcal{P}_0^n$ is to derive recourse conditions that can be used to determine for which customers the $(r+1)^{th}$ recourse action can takes place *conditioning on* the customers for which the r^{th} and $(r-1)^{th}$ recourse actions occurred. As will be shown later, under this approach the adversarial problem can be solved in polynomial time for recourse policy $\mathcal{P}_0^n$.

For recourse policy $\mathcal{P}_0^n$, the recourse conditions used here take the form $C^{r-1,r,r+1}(j,i,k)$ for customers j, i and k in the tour such that $j < i < k$, and for values of r such that $r \geq 1$. This conditions are used to determine if there exists a demand realization where the $(r+1)^{th}$ recourse occurs for k, *given* that the r^{th} recourse occurred for i and that the $(r-1)^{th}$ recourse occurred for j. When $r=1$, recourse conditions can always be evaluated using $j=1$, because by Observation 1, under $\mathcal{P}_0^n$ the state of the system when the vehicle departs from customer 1 is the same as if a recourse had taken place for this customer; that is, $I_1 = Q - d(1)$. For the same reason, $\underline{d}(i/1)$ in Lemma 6 is associated with a lower-bound on $d(i)$ if the first recourse takes place for customer i.

Proposition 8 (Recourse condition $C^1(i)$ for $\mathcal{P}_0^n$) *Consider a fixed sequence $\mathcal{T} = \{1, \cdots, n\}$ operated using recourse policy $\mathcal{P}_0^n$. There exists a demand realization $d \in \mathcal{U}$ such that the first recourse occurs for $i \in \mathcal{T}$ if and only if*

$$\sum_{\ell=1}^{i-1} \underline{d}(\ell) \leq Q \quad \textit{and} \quad \sum_{\ell=1}^{i} \overline{d}(\ell) \geq Q+1.$$

Proof.

($\Rightarrow$) Such demand realization d must satisfy

$$\sum_{\ell=1}^{i-1} d(\ell) \leq Q \quad \text{and} \quad \sum_{\ell=1}^{i} d(\ell) \geq Q+1.$$

Since $d \in \mathcal{U}$, $\underline{d}(\ell) \leq d(\ell)$ and $d(\ell) \leq \overline{d}(\ell)$ for all ℓ; summing over ℓ, the above implies

$$\sum_{\ell=1}^{i-1} \underline{d}(\ell) \leq \sum_{\ell=1}^{i-1} d(\ell) \leq Q \quad \text{and} \quad \sum_{\ell=1}^{i} \overline{d}(\ell) \geq \sum_{\ell=1}^{i} d(\ell) \geq Q+1.$$

($\Leftarrow$) Follows from Lemma 5 making $n = i$ and $M = Q$. $\square$

Proposition 9 (Recourse condition $C^{r-1,r,r+1}(j,i,k)$ for $\mathcal{P}_0^n$) *Consider a fixed sequence $\mathcal{T} = \{1, \cdots, n\}$ operated using recourse policy $\mathcal{P}_0^n$. Assume there exists a demand realization*

$d \in \mathcal{U}$ such that the $(r-1)^{th}$ recourse occurs for j and the r^{th} occurs for $i > j$ for $r > 0$. The $(r+1)^{th}$ recourse can occur for $k > i$ if and only if

$$\max\left\{\underline{d}(i/j), \underline{d}(i)\right\} + \sum_{\ell=i+1}^{k-1} \underline{d}(\ell) \leq Q \quad \textit{and} \quad \sum_{\ell=i}^{k} \overline{d}(\ell) \geq Q+1.$$

Proof.

($\Rightarrow$) Since there is a recourse for i, by Observation 1 the on-board inventory after i is served is $Q - d(i)$; hence, such demand realization d must satisfy

$$\sum_{\ell=i}^{k-1} d(\ell) \leq Q \quad \text{and} \quad \sum_{\ell=i}^{k} d(\ell) \geq Q+1$$

Since $d \in \mathcal{U}$, $d(\ell) \leq \overline{d}(\ell)$ for all ℓ, therefore $Q+1 \leq \sum_{\ell=i}^{k} d(\ell) \leq \sum_{\ell=i}^{k} \overline{d}(\ell)$ which corresponds to the second inequality in the proposition. Also, $\underline{d}(\ell) \leq d(\ell)$ for all ℓ; further, by Lemma 6, $\underline{d}(i/j) \leq d(i)$ and therefore $\max\{\underline{d}(i/j), \underline{d}(i)\} + \sum_{\ell=i+1}^{k-1} \underline{d}(\ell) \leq \sum_{\ell=i}^{k-1} d(\ell) \leq Q$.
($\Leftarrow$) Follows from Lemma 5 by making $M = Q$, $n = k-i+1$ and associating $d(1)$ in Lemma 5 with $d(i)$ in the proposition, $d(n)$ in Lemma 5 with $d(k)$ in the Proposition, and shifting all other entries accordingly; just need to further update $\underline{d}(1) = \max\{\underline{d}(i/j), \underline{d}(i)\}$, which follows from Lemma 6. □

Propositions 8 and 9 state the recourse conditions C^1 and $C^{r-1,r,r+1}$ for $\mathcal{P}_0^n$. Although this policy is not history independent, observe that the only element in the proposition bringing information from the past is $\underline{d}(i/j)$. In order to determine in which customer the $r+1$ recourse can occur, only information about where the r and $r-1$ recourses needs to be considered (not about any previous recours actions). Also, observe that it is not necessary to keep track of the inventory level after customers i or j are served.

It is possible to find instances of the VRPSD where the $r+1$ recourse can occur for k given the that the r recourse occurred for i and the $r-1$ occurred for j_1, but it might not be the case if the $r-1$ occurred for j_2. Observe that when $r = 1$, the situation illustrated above cannot occur because, by Observation 1, the result of Proposition 9 can be used by making $j = 1$ for all i. The following example illustrates such a case.

Example. Consider the six customer tour $\mathcal{T}$ in Figure 13 and assume that it is operated using recourse policy $\mathcal{P}_0^n$ with a vehicle of capacity $Q = 3$. It is easy to check that customers 2 to 6 satisfy the conditions of Proposition 8, and so there are demand realizations for which the first recourse occurs at their locations.

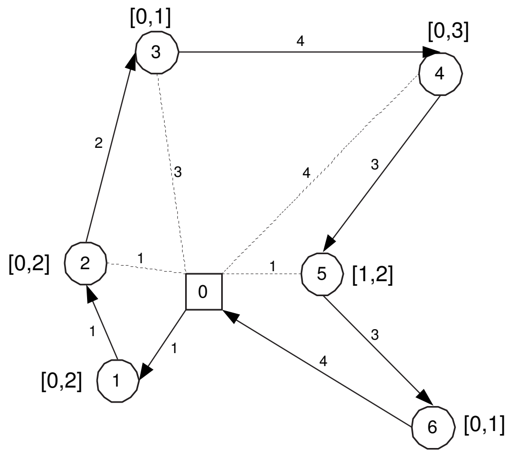

Figure 13: Tour operated using $\mathcal{P}_0^n$ and $Q = 3$.

Observe that $\underline{d}(2/1) = 2 = \overline{d}(2)$, and recall that under our notation $\underline{d}(2/1)$ it is associated with the first recourse taking place for customer 2. Therefore, if the first recourse occurs for 2 then $d(2) = 2$. Similarly, since $\underline{d}(3/1) = 1$, if the first recourse occurs for 3 then $d(3) = 1$. Evaluating recourse conditions $C^{0,1,2}(1, 2, 4)$ from Proposition 9, since $\underline{d}(2/1) + \underline{d}(3) = 2$ and $\sum_{\ell=2}^{4} \overline{d}(\ell) = 6$, then there exists a demand realization in the uncertainty set for which the first recourse occurs for 2 and the second for 4. Also, $\sum_{\ell=3}^{4} \overline{d}(\ell) = 4$, so there exists a demand realization in which the first recourse occurs for 3 and the second at 4.

Now consider $C^{1,2,3}(2, 4, 6)$ and observe that $\underline{d}(4/2) = 1$; then, by Proposition 9 since $\underline{d}(4/2) + \underline{d}(\ell) = 2$ and $\sum_{\ell=4}^{6} \overline{d}(\ell) = 5$, there exists a demand realization such that the first recourse occurs for 2, the second for 4 and the third for 6. On the other hand, consider $C^{1,2,3}(3, 4, 6)$ and observe $\underline{d}(4/3) = 3$; hence, there does not exist a demand realization such that the first recourse occurs for 3, the second for 4 and the third for 6. Observe then that the plausibility of having the third recourse taking place at customer 6 given the second recourse at customer 4 depends on which customer the first recourse occurred.

It is now clear that $\Phi(\mathcal{T}, \mathcal{P}_0^n)$ cannot be evaluated using network $\mathcal{G}_0$; still, since a limited part of the history of the process needs to be traced to evaluate $C^{r-1,r,r+1}$, a new network is defined. Let $\mathcal{G}_1(\mathcal{P}_0^n) = (\mathcal{N}, \mathcal{A})$ denote this network, which is now defined and illustrated in Figure 14 for the instance in the previous example. The set of nodes is defined as

$$\mathcal{N} = \{s\} \cup \{t\} \cup \{(1, i/1) \,|\, i \in \mathcal{T} \setminus \{1\}\} \cup \mathcal{N}'$$

where $(1, i/1)$ is associated with having the first recourse at customer i, and

$$\mathcal{N}' = \big\{(r, i/j) \,|\, r \in \{2, \cdots, R\},\, i \in \{r+1, \cdots, n\},\, j \in \{r, \cdots, i-1\}\big\}$$

where $(r, i/j)$ is associated with having the r^{th} recourse at customer i given that the $(r-1)^{th}$ took place at customer j. The arc set $\mathcal{A}$ is defined as follows:

1. $(s, (1, i/1)) \in \mathcal{A}$ for every $i \in \mathcal{T}$ that satisfies Proposition 8. The cost of each arc is the net increment in traveled time incurred when making a recourse action for customer i.

2. For $r = 1$ to $R - 1$, $((r, i/j), (r+1, k/i)) \in \mathcal{A}$ for $i, j, k \in \mathcal{T}$ such that $j < i < k$ that satisfy Proposition 9. The cost of each arc is the net increment in traveled time incurred when making a recourse action for customer k.

3. For all $(r, i/x) \in \mathcal{A}$ such that $indeg(r, i/x) > 0$ and the $outdeg(r, i/x) = 0$, then $((r, i/x,), t) \in \mathcal{A}$ with cost 0.

Let $L(\mathcal{G}_1(\mathcal{P}_0^n))$ be the longest s-t path in the network. If there does not exist an $s - t$ path, then by definition $L(\mathcal{G}_1(\mathcal{P}_0^n)) = 0$.

Theorem 5 *Consider a fixed sequence $\mathcal{T} = \{1, \cdots, n\}$ operated using recourse policy $\mathcal{P}_0^n$. For instances such that $R \leq 2$, $\Phi(\mathcal{T}, \mathcal{P}_0^n)$ can be evaluated in $O(n^3)$ by solving a longest s-t path problem on $\mathcal{G}_0(\mathcal{P}_0^n)$. In general, $\Phi(\mathcal{T}, \mathcal{P}_0^n)$ can be evaluated in $O(n^5)$ by solving a longest s-t path problem on $\mathcal{G}_1(\mathcal{P}_0^n)$.*

Proof. Observe that although $\mathcal{P}_0^n$ is not history independent, when evaluating the conditions in Proposition 9 on i for $r = 1$, since I_1 always equals $Q - d(1)$, by Observation 1 it is clear

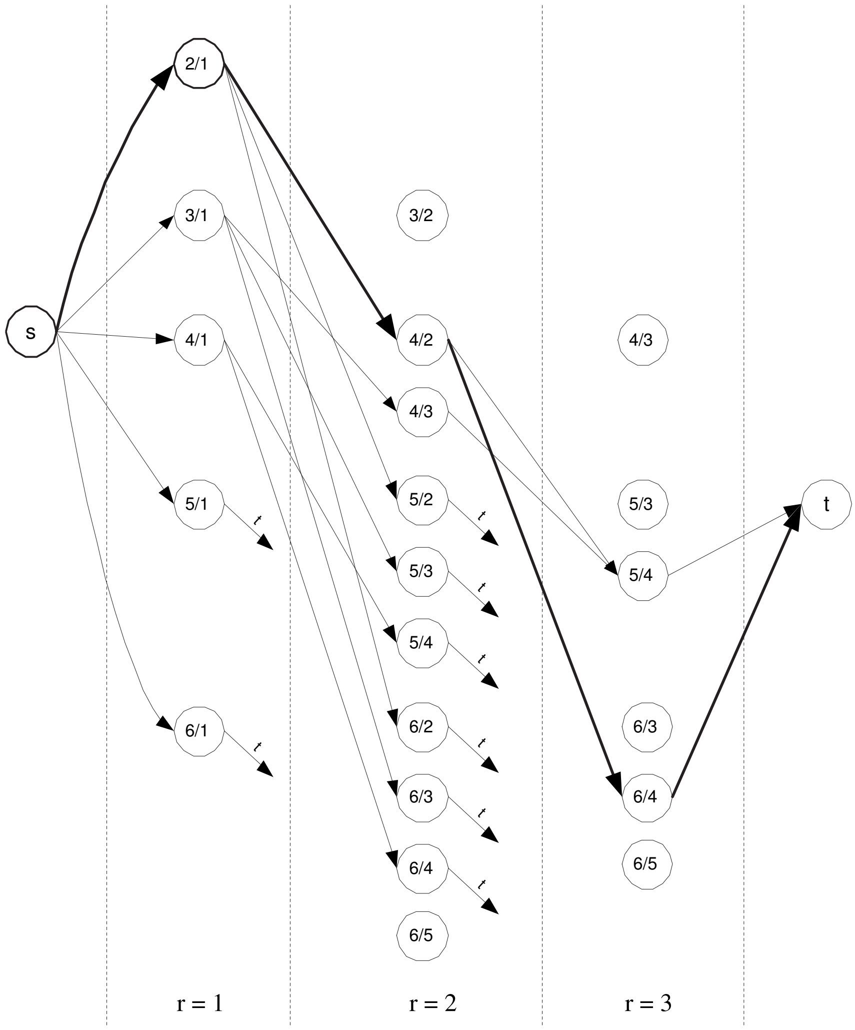

Figure 14: Network $\mathcal{G}_1(\mathcal{P}_0^n)$

that $\underline{d}(i/j)$ only needs to be evaluated for $j = 1$ independent of the demand history between customers 1 and $i-1$; hence $\mathcal{G}_0(\mathcal{P}_0^n)$ can be used to solve adversarial problem for tours where $R \leq 2$. $\mathcal{G}_0(\mathcal{P}_0^n)$ is an acyclic network, so the longest s-t path can be calculated in $O(|\mathcal{A}|)$. The number of $(s, (i, 1))$ and $((i, r), t)$ arcs is $O(n)$ in each case; further, the number of $((i, 1), (k, 2))$ is $O(n^2)$; construction of network $\mathcal{G}_0(\mathcal{P}_0^n)$ is done in $O(n^3)$ because recourse conditions of proposition 9 can be checked on $O(n)$ for each arc, therefore $\Phi(\mathcal{T}, \mathcal{P}_0^n)$ can be evaluated in $O(n^3)$.

In general, $\Phi(\mathcal{T}, \mathcal{P}_0^n)$ can be calculated using $\mathcal{G}_1$. Observe that the number of nodes in the r^{th} layer of the network for $r > 1$ is bounded by $\sum_{i=1}^{n-r} i = \frac{(n-r)(n-r+1)}{2}$; each node $(i/j, r)$ is connected to at most $n-i$ nodes in the $r+1$ layer. Therefore, the number of arcs between the r and $r+1$ layers is $O(n^3)$. R is bounded by n, so summing the number of arcs over all r, the resulting number of these arcs is $O(n^4)$. The number of arcs with tail in s is bounded by n. The number of nodes in the network is clearly bounded by n^3 so the number of arcs with head t is also bounded by n^3. The total number of arcs in the network is then $O(n^4)$, and checking recourse conditions of propositions 8 and 9 can be done in $O(n)$ for each arc, therefore $\Phi(\mathcal{T}, \mathcal{P}_0^n)$ can be evaluated in $O(n^5)$. □

Example. Consider again the tour in Figure 13, observe that $\mathrm{L}(\mathcal{T}) = 18$, and that for demand realization $\overline{d}$, Table 10 shows the corresponding value for the state variables. From this table, $R = 3$ and $\phi(\mathcal{T}, \mathcal{P}_0^n, \overline{d}) = 12$.

Table 10: Values of the state variables for demand realization $\overline{d}$ when the tour is operated using recourse policy $\mathcal{P}_0^n$.

i	1	2	3	4	5	6
$d(i)$	2	2	1	3	2	1
r	0	1	0	2	3	0
I_i	1	1	0	0	1	0

Network $\mathcal{G}_1(\mathcal{P}_0^n)$, illustrated in Figure 14, is now used to solve the adversarial problem. From Proposition 8, it is clear that there is a demand realization in which the first recourse action can occur in one of the customers between 2 and 6. There is an arc connecting s and each of the associated nodes in $\mathcal{G}_1$. By using the recourse conditions $C^{0,1,2}(1, 2, 3)$ of Proposition 9, $\overline{d}(2)+\overline{d}(3) = 3 < Q+1$; it can be concluded that *given* that the first recourse was taken for customer 2, the second recourse cannot occur for customer 3.
Consider $C^{0,1,2}(1, 2, 4)$, observe $\overline{d}(2)+\overline{d}(3)+\overline{d}(4) = 6 \geq Q+1$ and $\underline{d}(2/1)+\underline{d}(3) = 2$ where $\underline{d}(2/1) = 2$; therefore, there is a demand realization in which the second recourse occurs for customer 4 given that the first occurred for customer 2; this is incorporated in $\mathcal{G}_1$ with an arc that connects node $(1, 2/1)$ with node $(2, 4/2)$. By evaluating these conditions, it can be seen that this is also true for customers 5 and 6.

Given that the second recourse can occur for customer 4, it is now determined in which customers can the third recourse take place. Consider recourse condition $C^{1,2,3}(2,4,5)$ and $C^{1,2,3}(3,4,5)$; observe that $\underline{d}(4/2) = 1$ and $\underline{d}(4/3) = 3$; also, $\overline{d}(4) = 3$ and $\overline{d}(4) = 2$; therefore condition $\sum_{\ell=4}^{5} \overline{d}(\ell) = 5 \geq Q+1$ is satisfied. It can be concluded that the third recourse can occur for customer 5 independently of whether the first recourse action took place for customer 2 or 3. This is incorporated in $\mathcal{G}_1$ by adding arcs $((2,4/2),(3,5/4))$ and $((2,4/3),(3,5/4))$. Consider now customer 6, and the corresponding recourse condition $C^{1,2,3}(2,4,6)$ and $C^{1,2,3}(3,4,6)$; observe that $\sum_{\ell=4}^{6} \overline{d}(\ell) \geq Q+1$ and $\underline{d}(4/2)+\underline{d}(5) = 2 \leq Q$, but $\underline{d}(4/3)+\underline{d}(5) = 4 > Q$. In this case, given that the second recourse occurs for customer 4, the third recourse can take place for customer 6 only if the first occurred for customer 2, and not if it occurred for customer 3. In this case, only arc $((2,4/2),(3,6/4))$ is added into $\mathcal{G}_1$.

The longest $s-t$ path in $\mathcal{G}_1(\mathcal{P}_0^n)$ corresponds to $\{s,(1,2/1),(2,4/2),(3,6/4),t\}$ with an associated cost of 18. The following observations are important. First, it does not correspond to the cost obtained for demand realization $\overline{d}$. Second, demands realizations d for which $\phi(\mathcal{T},\mathcal{P}_0^n,d)$ achieves cost 18 are $[2,2,0,2,1,1]$, $[2,2,1,1,2,1]$ and $[2,2,1,2,1,1]$, none of which correspond to an extreme point of the uncertainty space.

4.3.1.3 Anticipatory and non-splitting policies

Definition 11 *$\mathcal{P}_1^n$ is used to denote the following one customer ahead anticipatory and non-splitting recourse policy on a fixed sequence $\mathcal{T} = \{1,\cdots,n\}$: a recourse action is taken at customer $i-1 \in \mathcal{T}$ for customer $i \in \mathcal{T}$ if and only if after the vehicle services customer $i-1$ it observes that on-board inventory is strictly less than $\overline{d}(i)$. The vehicle restocks at the depot and then resumes the tour at customer i.*

In this policy a recourse action is undertaken after a vehicle services a customer when there is the possibility that when the vehicle arrives at the next customer's location a restock must take place. In order to anticipate this, the vehicle deviates to the depot, restocks and resumes the tour by traveling directly to the next customer's location. The net increment in travel time due to taking a recourse action at customer $i-1$ is $l(i-1,0)+l(0,i)-l(i-1,i)$. Under triangular inequality on the travel time, the total travel time would be reduced by anticipating the stock-out if indeed a recourse actions was to occur; on the other hand, an unnecessary recourse could have taken place if it was not the case.

Observe that by the definition of this policy, given that a recourse occurred at customer $i-1$ for customer i, the location where the next recourse would occur, say k, depends only on $d(i), \cdots, d(k)$, therefore the following observation can be made.

Observation 2 *$\mathcal{P}_1^n$ is a history independent recourse policy.*

Proposition 10 (Recourse condition $C^1(i)$ for $\mathcal{P}_1^n$) *Consider a fixed sequence $\mathcal{T} = \{1, \cdots, n\}$ operated using recourse policy $\mathcal{P}_1^n$. There exists a demand realization $d \in \mathcal{U}$ such that the first recourse occurs at $i-1 \in \mathcal{T}$ for $i \in \mathcal{T}$ if and only if*

$$\sum_{\ell=1}^{j-1} \underline{d}(\ell) + \overline{d}(j) \le Q \quad \forall j = 2, \cdots, i-1 \quad \textit{and} \quad \sum_{\ell=1}^{i} \overline{d}(\ell) \ge Q+1.$$

Proof.

($\Rightarrow$) Such demand realization d must satisfy

$$\sum_{\ell=1}^{j-1} d(\ell) + \overline{d}(j) \le Q \quad \forall j = 2, \cdots, i-1 \quad \text{and} \quad \sum_{\ell=1}^{i} d(\ell) \ge Q+1$$

Since $d \in \mathcal{U}$, $\underline{d}(\ell) \le d(\ell) \le \overline{d}(\ell)$ for all ℓ, then the above expressions imply

$$\sum_{\ell=1}^{j-1} \underline{d}(\ell)+\overline{d}(j) \le \sum_{\ell=1}^{j-1} d(\ell)+\overline{d}(j) \le Q \;\; \forall j = 2, \cdots, i-1 \quad \text{and} \quad \sum_{\ell=1}^{i} \overline{d}(\ell) \ge \sum_{\ell=1}^{i} d(\ell) \ge Q+1.$$

($\Leftarrow$) The result follows from Lemma 5 by making $n = i$, $M = Q$ and $\underline{d}(n-1) = \overline{d}(n-1)$ for $j = i-1$. By the policy definition, if $\sum_{\ell=1}^{j-1} \underline{d}(j) > Q$ for any $j < i-1$ then the first recourse must occur before customer $i-1$ for i. $\square$

Proposition 11 (Recourse condition $C^{r,r+1}(i,k)$ for $\mathcal{P}_1^n$) *Consider a fixed sequence $\mathcal{T} = \{1, \cdots, n\}$ operated using recourse policy $\mathcal{P}_1^n$. Assume there exists a demand realization*

$d \in \mathcal{U}$ such that the r^{th} recourse occurs at $i-1$ for i. The $(r+1)^{th}$ recourse can occur at $k-1$ for k if and only if

$$\sum_{\ell=i}^{j-1} \underline{d}(\ell) + \overline{d}(j) \leq Q \quad \forall j = 2, \cdots, k-1 \qquad \textit{and} \qquad \sum_{\ell=i}^{k} \overline{d}(\ell) \geq Q+1.$$

Proof.

From Observation 2, it is clear that when a recourse is taken for customer i, the situation (in terms of where the next recourse will occur) is identical to the one when the vehicle departs the depot for the first time. The result then follows from Proposition 10. □

Theorem 6 *Consider a fixed sequence $\mathcal{T} = \{1, \cdots, n\}$ operated using recourse policy $\mathcal{P}_1^n$. $\Phi(\mathcal{T}, \mathcal{P}_1^n)$ can be evaluated in $O(n^4)$ by solving a longest s-t path problem on $\mathcal{G}_0(\mathcal{P}_1^n)$.*

Proof. $\mathcal{P}_1^n$ is a history independent recourse policy, so the result follows from Lemma 3 and Lemma 4, further observing that the conditions in propositions 10 and 11 can be checked in $O(n)$ for each arc in $\mathcal{G}_0(\mathcal{P}_1^n)$. □

Example. Consider the tour in Figure 13 but operated using recourse policy $\mathcal{P}_1^n$. Observe that the net increment in travel time due to taking a recourse at customer 1 for customer 2 is 1; at customer 2 for 3 is 2; at customer 3 for 4 is 3 ; at customer 4 for 5 is 2 and at customer 5 for 6 is 2. Clearly under $\mathcal{P}_1^n$ a recourse action would never be taken for customer 1 or at customer 6.

Table 11: Values of the state variables for demand realization $\overline{d}$ when the tour is operated using recourse policy $\mathcal{P}_1^n$.

i	1	2	3	4	5	6
$d(i)$	2	2	1	3	2	1
r	0	1	0	2	3	0
I_i	1	1	0	0	1	0

From Table 11 it can be seen that $\phi(\mathcal{T}, \mathcal{P}_1^n, \overline{d}) = 6$ and $R = 3$. Network $\mathcal{G}_0(\mathcal{P}_1^n)$ is used to determine the optimal solution to the adversarial problem. Figure 15 shows $\mathcal{G}_0(\mathcal{P}_1^n)$ for this example. A longest $s-t$ path is $\{s,3,4,6,t\}$ with cost 7.

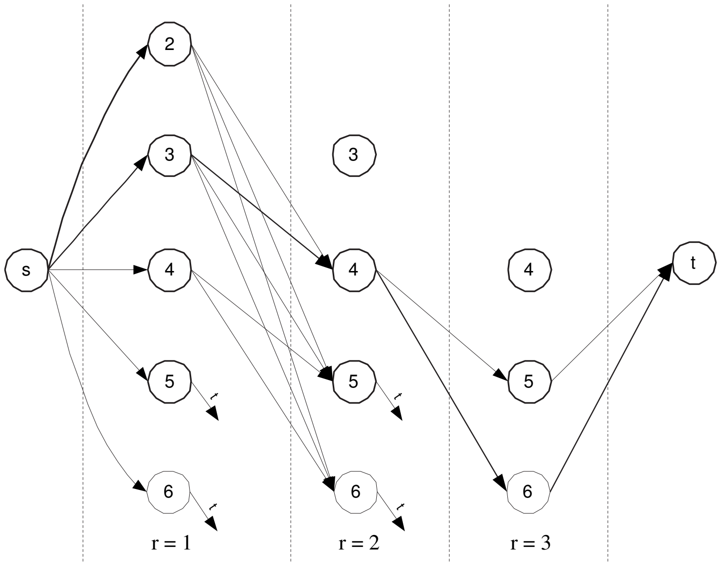

Figure 15: Network $\mathcal{G}_0(\mathcal{P}_1^n)$

Up to this point, all recourse policies considered in this research are simple decision rules based on very local information; namely, take a recourse action now because there is not enough on-board inventory to satisfy the current customer's demand, or because there is a chance of not being able to satisfy the next customer's demand. Recourse policies that use an optimization approach, and that take into account more global information are now considered. It is assumed that after each customer is served, a recourse plan is dynamically generated based on on-board inventory, the travel time matrix $(l(i,j))$ and an estimate of the demand of the remaining customers in the tour. A recourse plan specifies at which customers will recourse actions be taken assuming that demand realization will be equal to demand estimates. The objective is to determine an optimal recourse plan, which is the one that minimizes total traveled time over the remaining customers. This plan is dynamically updated, if needed, after each customer is served using the newly available information. Only non-splitting recourse plans are considered, but all results are easily extendible to the splitting case.

In order to formalize this concept, assume that customer $i-1$ in the tour has just been served and that on-board inventory is I_{i-1}. Let

$$\mathcal{T}_i \;=\; \{j \in \mathcal{T} \mid\; i \leq j\}$$

represent the set of remaining customers in the tour, which will be considered in order to determine a recourse plan. Let $\hat{d}_i \in \mathbb{Z}_+^{n-i+1}$ represent the vector of demand estimates of the customers in $\mathcal{T}_i$, where $\hat{d}_i(j)$ is the estimate of the demand for customer $j \in \mathcal{T}_i$. Observe that for a given value of vector $\hat{d}_i$ there are a finite number of subsets of customers at which non-splitting recourse actions that satisfy all estimated demand can take place. Let $R(I_{i-1}, \hat{d}_i)$ represent the set of all such subsets of customers in $\mathcal{T}_i$ assuming on-board inventory I_{i-1} and demand $\hat{d}_i$. Let $\phi'(\mathcal{T}_i, r)$ denote the total traveled time between $i-1$ and the final arrival to the depot after all customers have been served under recourse plan $r \in R(I_{i-1}, \hat{d}_i)$.

We are then interested in solving problem

$$\mathbf{ORP}(\mathcal{T}_i, I_{i-1}, \hat{d}_i) \;=\; \min_{r \in R(I_{i-1}, \hat{d}_i)} \;\phi'(\mathcal{T}_i, r) \tag{52}$$

referred to as the Optimal Recourse Plan problem, where $r^\star$ is used to denote an optimal solution.

Following an approach similar to the one proposed by Beasley [4] for the deterministic VRP, the problem (52) can be solved as a shortest path problem on a directed network. Let $G(I_{i-1}, \hat{d}_i) = (\mathcal{N}, \mathcal{A})$, where $\mathcal{N} = \{s\} \cup \{t\} \cup \mathcal{T}_i$, and arc set $\mathcal{A}$ defined as

1. $(s, j) \in \mathcal{A}$ for all $j \in \mathcal{T}_i$ s.t. $\sum_{\ell=i}^{j} \hat{d}_i(\ell) \leq I_{i-1}$, with arc-cost $= \sum_{\ell=i-1}^{j-1} l(\ell, \ell+1)$; each of these arcs is associated with taking the next recourse action at customer j.

2. $(s, j) \in \mathcal{A}$ for all $j \in \mathcal{T}_i$ s.t. $\sum_{\ell=i}^{j} \hat{d}_i(\ell) \leq Q$, arc-cost $= l(i-1, 0) + l(0, i) + \sum_{\ell=i}^{j-1} l(\ell, \ell+1)$; each of these arcs is associated with taking a recourse action at customer $i-1$ and then a recourse actions at customer j.

3. Similarly, $(j, k) \in \mathcal{A}$ for all $j, k \in \mathcal{T}_i$, s.t. $j < k$ and $\sum_{\ell=j+1}^{k} \hat{d}_i(\ell) \leq Q$, arc-cost $= l(j, 0) + l(0, j+1) + \sum_{\ell=j+1}^{k-1} l(\ell, \ell+1)$. Each of these arcs are associated with taking a recourse actions at customer j and the next recourse at customer k.

4. Finally, add an arc that connects node $n \in \mathcal{T}_i$ to t with cost $l(n, 0)$.

By the construction of the network, it is clear that each $s-t$ path is associated with a recourse plan $r \in R(I_{i-1}, \hat{d}_i)$. Observe that for the recourse policy under consideration it does not make any sense to take a recourse at customer n, so a recourse action is considered at every customer in the tour, except at this customer. An optimal recourse plan ($r^\star$) is then associated with the shortest $s-t$ path in $G(I_{i-1}, \hat{d}_i)$.

The concept of an optimal recourse plan is now used to define a new recourse policy for the case of no advance demand information. It corresponds to a generalization of $\mathcal{P}_1^n$, where after serving a customer instead of just examining the next customer, an optimal recourse plan is determined considering the remaining customers in the tour assuming the maximum demand realization for each of them.

Definition 12 *$\mathcal{P}_n^n$ is used to denote the following n customers ahead anticipatory and non-splitting recourse policy on a fixed sequence $\mathcal{T} = \{1, \cdots, n\}$: for each customer $i \in \mathcal{T}$ prior to departing from $i-1$, solve* $\mathbf{ORP}(\mathcal{T}_i, I_{i-1}, \hat{d}_i)$*, for $\hat{d}_i(j) = \overline{d}(j)$ for all $j \in \mathcal{T}_i$ and follow the optimal recourse plan $r^\star$.*

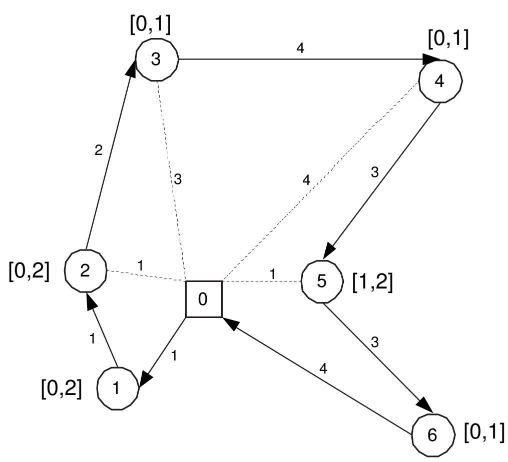

Figure 16: Tour operated using $\mathcal{P}_n^n$ with $Q = 4$.

Example. Consider the tour in Figure 16 operated using P_n^n with a vehicle of capacity $Q = 4$. When departing from the depot for the first time, problem $\mathbf{ORP}(\mathcal{T}_1, Q, \overline{d})$ is solved by identifying the shortest $s-t$ path in $G(Q, \overline{d})$. Figure 17(a) shows this network; the shortest $s-t$ path is $\{s,1,4,6,t\}$, hence $r^\star$ states taking recourse actions at customers 1 and 4, with a total travel time of 21.

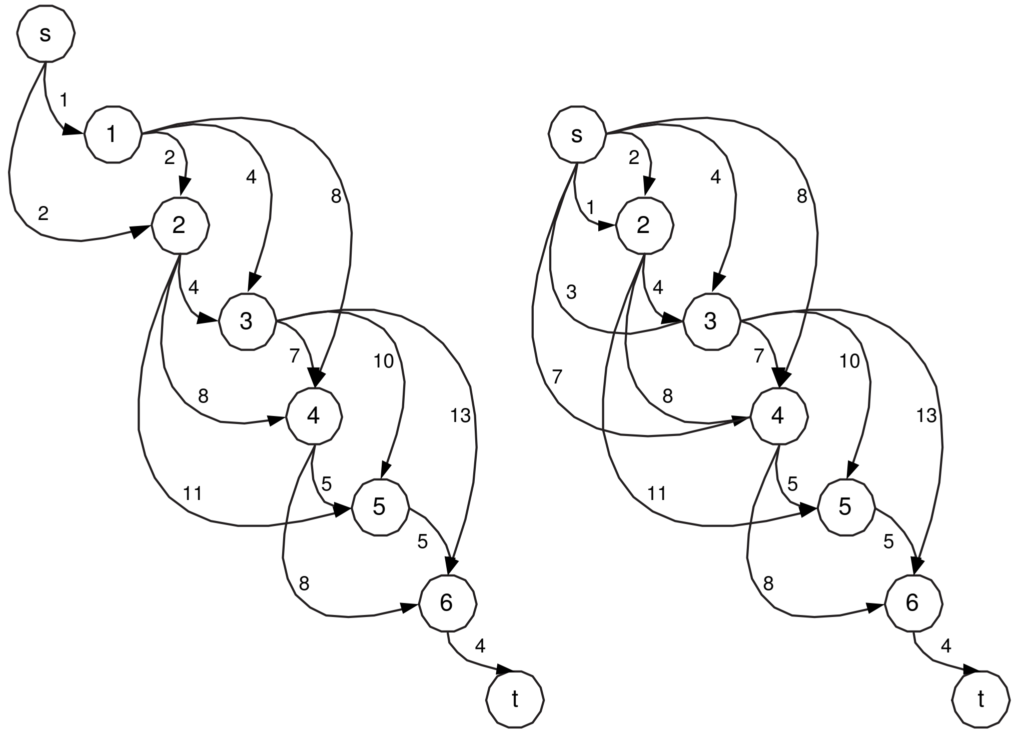

Figure 17: (a) Network $G(Q,\overline{d})$ for $\mathcal{T}_1$. (b) Network $G(I_1,\overline{d})$ for $\mathcal{T}_2$ using $I_1 = 4$ after observing $d(1) = 0$.

.

Assume that when the vehicle reaches the first customer it observes $d(1) = 0$; then for $\mathcal{T}_2$ the shortest $s-t$ path in $G(I_1, \hat{d}_1)$ using $I_1 = 4$ needs to be identified in order to determine the corresponding optimal plan $r^\star$ starting at customer 1 for demand estimate $\hat{d}_j = \overline{d}(j)$ for all $j \in \mathcal{T}_2$. Figure 17(b) shows this network; the shortest $s-t$ path is $\{s,4,6,t\}$. Observe then that given the newly available demand information, $r^\star$ states taking recourse action only at customers 4, with a total travel time of 19 from the location of customer 1. It took 1 time unit to travel to customer 1, so the total travel time if this recourse plan was followed would be 20.

This process is then repeated at each of the remaining customers in the tour, dynamically updating the recourse plan in order to try to reduce the total traveled time. Observe that in this case there was a reduction of 1 unit in total travel time for the new recourse plan generated at customer 1 when compared to the initial plan generated at the depot.

In order to implement recourse policy $\mathcal{P}_n^n$, the first step is to solve an shortest $s-t$ path problem in $G(Q,\overline{d})$. This problem can be solved in $O(nQ)$ because $G(Q,\overline{d})$ is an acyclic network with $n+1$ nodes, where the maximum number of arcs with tail in any node is Q, and with $outdgeg(t)=0$. If a shortest $s-t$ problem is solved from scratch at every customer in the tour, the total solution time would be $O(n^2Q)$.

Consider networks $G(I_{i-1},\overline{d}_i)$ and $G(I_i,\overline{d}_{i+1})$, where $\overline{d}_i(j)=\overline{d}(j)$ and $\overline{d}_{i+1}(k)=\overline{d}(k)$ for all $j\in\mathcal{T}_i$ and $k\in\mathcal{T}_{i+1}$ respectively. Observe that $G(I_{i-1},\overline{d}_i)$ and $G(I_i,\overline{d}_{i+1})$ are almost identical; the main difference is that $G(I_i,\overline{d}_{i+1})$ does not include node i; also, node s in $G(I_i,\overline{d}_{i+1})$, which is associated with node i in $G(I_{i-1},\overline{d}_i)$, is the tail of some additional arcs. These additional arcs are associated with not taking a recourse action at customer i. This implies that in order to solve the Optimal Recourse Plan problem (52) at customer i there is no point in making all calculation from scratch because most of the problem data in the previous stage, at customer $i-1$, is the same.

Finding the shortest $s-t$ path in $G(Q,\overline{d})$ can be accomplished by means of a labeling algorithm, where the label of node $j\in\mathcal{T}$ represents the minimum travel time from j to t. From the argument in the previous paragraph, except for node s, the label of all other nodes in $G(I_1,\overline{d}_2)$ do not change with respect to the labels calculated in $G(Q,\overline{d})$. Calculating the label for s requires $O(Q)$ operations, corresponding to evaluating the additional arcs associated with not taking a recourse action at i. This argument is also true for any $G(I_{i-1},\overline{d}_i)$ and $G(I_i,\overline{d}_{i+1})$ for $i>2$; therefore, calculations required throughout the whole process can then be performed in $O(nQ+nQ)=O(nQ)$.

Proposition 12 *Consider a fixed sequence $\mathcal{T}=\{1,\cdots,n\}$ operated using recourse policy $\mathcal{P}_n^n$. An optimal solution to the adversarial problem is always $d=\overline{d}$.*

Proof. For any $i\in\mathcal{T}$, by the construction of network $G(I_{i-1},\hat{d})$, $\phi'(\mathcal{T}_i,r^\star)$ is a non-increasing function on I_{i-1}, because as I_{i-1} decreases, the arc set on this network reduces, further reducing the number of $s-t$ paths in $G(I_{i-1},\hat{d})$. I_{i-1} decreases as $d(i-1)$ increases; it then follows that $\mathbf{ORP}(\mathcal{T}_i,I_{i-1},\hat{d})$ achieves its maximum value by setting $d(i-1)=\overline{d}(i-1)$.

Now consider the optimal recourse plan $r^\star$ obtained from solving $\mathbf{ORP}(\mathcal{T}_i, I_{i-1}, \overline{d})$ and observe that $r^\star$ remains feasible (*i.e.*, all demands can be satisfied) for any demand realization $d \in \mathcal{U}$ the vehicle actually sees as it visits the customers in $\mathcal{T}_i$; hence the cost of $\phi'(\mathcal{T}_i, r^\star)$ is always attainable if the vehicle follows $r^\star$. Therefore, the optimal recourse plan $r^\star$ obtained from solving $z_0^\star = \mathbf{ORP}(\mathcal{T}_1, Q, \overline{d})$ is always feasible and $z_0^\star$ is an upper-bound on the actual total travel time on the tour. Observe that this bound is tight if $d(j) = \overline{d}(j)$ is chosen for all subsequent problems associated with $j \in \mathcal{T}$.

All that remains to show is that $z_0^\star$, obtained by setting $d = \overline{d}$, corresponds to the maximum value of $\phi'(\mathcal{T}_1, r^\star)$. This follows from the construction of $G(Q, \hat{d})$, because the arc set of the network reduces as $\hat{d}(j)$ increases for every $j \in \mathcal{T}$. □

4.3.2 Total advance demand visibility

In this section it is assumed that before the vehicle departs the depot for the first time, the demand of all customers in the tour is known. Total visibility of demand allows for an optimal recourse plan.

Definition 13 *$\mathcal{P}_{nv}^n$ is used to denote the following n customer ahead with perfect demand visibility and non-splitting recourse policy on a fixed sequence $\mathcal{T} = \{1, \cdots, n\}$: Once demand realization d is observed, follow the optimal recourse plan $r^\star$ obtained by solving* $\mathbf{ORP}(\mathcal{T}_1, Q, d)$.

The adversarial problem in this case has a simple solution. The following result is quite intuitive.

Proposition 13 *Consider a fixed sequence $\mathcal{T} = \{1, \cdots, n\}$ operated using recourse policy $\mathcal{P}_{nv}^n$. An optimal solution to the adversarial problem is always $d = \overline{d}$.*

Proof. Parallel to the proof of Proposition 12. □

4.3.3 Partial advance demand information

4.3.3.1 One customer ahead

In this section it is assumed that after a vehicle finishes servicing a customer, before departing there is perfect visibility of the demand of the next customer in the tour. This corresponds to an intermediate case between the the no advance demand information and the total demand visibility cases.

Definition 14 *$\mathcal{P}_{1v}^n$ is used to denote the following one customer ahead anticipatory with perfect demand visibility and non-splitting recourse policy on a fixed sequence $\mathcal{T} = \{1, \cdots, n\}$: a recourse action is taken at customer $i-1 \in \mathcal{T}$ for customer $i \in \mathcal{T}$ if and only if after the vehicle services customer $i-1$ it observes that on-board inventory is strictly less than $d(i)$. The vehicle restocks at the depot and resumes the tour at customer i.*

Observation 3 *Consider a fixed sequence $\mathcal{T} = \{1, \cdots, n\}$ and a demand realization $d \in \mathcal{U}$. For d there is a recourse action at i under $\mathcal{P}_0^n$ if and only if for d there is a recourse action at $i-1$ for i under $\mathcal{P}_{1v}^n$.*

Theorem 7 *Proposition 8 and Proposition9 are also valid for $\mathcal{P}_{1v}^n$. Further, Theorem 5 also holds for $\mathcal{P}_{1v}^n$.*

Proof. Follows from Observation 3. $\square$

4.4 Analysis of recourse policies

We are now interested in analyzing and even assessing the benefits of having total or partial demand visibility over no advance demand information. Intuitively, it makes sense to believe that better decisions can always be made when more information is available. We now focus our attention on proving or disproving the validity of this intuitive result from a worst-case perspective by comparing the value of function $\mathcal{L}(\mathcal{T}, \mathcal{P})$ for all non-splitting recourse policies defined in the previous section.

From the definition of recourse policy $\mathcal{P}_{nv}^n$, since the decision maker has total visibility of all customer demands before the vehicle departs the depot, it then follows that $\mathcal{L}(\mathcal{P}_{nv}^n)$ is a lower-bound on the value of $\mathcal{L}(\mathcal{T}, \mathcal{P})$ for any non-splitting recourse policies $\mathcal{P}$ considered in this research. Equivalently, since in $\mathcal{P}^n_{nv}$ the decision maker has the best attainable information of demand and total flexibility to create a recourse plan, no other policy can perform better. The following observation is then obvious.

Observation 4 *Consider a fixed sequence $\mathcal{T}$. For instances with triangular inequality on the travel time*

$$\begin{aligned}
\mathcal{L}(\mathcal{T}, \mathcal{P}^n_{nv}) &\leq \mathcal{L}(\mathcal{T}, \mathcal{P}^n_0), \\
\mathcal{L}(\mathcal{T}, \mathcal{P}^n_{nv}) &\leq \mathcal{L}(\mathcal{T}, \mathcal{P}^n_1), \\
\mathcal{L}(\mathcal{T}, \mathcal{P}^n_{nv}) &\leq \mathcal{L}(\mathcal{T}, \mathcal{P}^n_n), \\
\mathcal{L}(\mathcal{T}, \mathcal{P}^n_{nv}) &\leq \mathcal{L}(\mathcal{T}, \mathcal{P}^n_{1v})
\end{aligned}$$

We begin by considering recourse policies $\mathcal{P}^n_{1v}$ and $\mathcal{P}^n_0$, and our interest is to determine if there exists a dominant relation of the former over the latter since $\mathcal{P}^n_{1v}$ assumes partial visibility of the demand process while $\mathcal{P}^n_0$ does not. The following proposition establishes the existence of such a relation.

Proposition 14 *Consider a fixed sequence $\mathcal{T}$. For instances with triangular inequality on travel time*

$$\mathcal{L}(\mathcal{T}, \mathcal{P}^n_{1v}) \leq \mathcal{L}(\mathcal{T}, \mathcal{P}^n_0)$$

Proof. Follows from Theorem 7 under triangular inequality on the travel time because $\mathcal{G}_1(\mathcal{P}^n_0)$ and $\mathcal{G}_1(\mathcal{P}^n_{1v})$ are identical expect for the costs on the arcs. By Observation 3, for the same demand realization if a recourse takes place at i under $G(\mathcal{P}^n_0)$ a recourse will be taken at $i-1$ for i; the additional traveled time in the first case is $l(i,0)+l(0,i)$, and in the second is $l(i-1,0)+l(0,i)l(i-1,i)$. Subtracting the latter expression from the $l(i,0)+l(0,i)$, the result is $l(i-1,i)+l(i,0)-l(i-1,i)$, which by triangular inequality has to be ≥ 0. □

Proposition 14 implies that $\mathcal{P}^n_{1v}$ dominates $\mathcal{P}^n_0$ from a worst-case perspective. That is, given $\mathcal{T}$, the value of function $\mathcal{L}(\mathcal{T}, \mathcal{P}^n_{1v})$ will be less than or equal to the value of $\mathcal{L}(\mathcal{T}, \mathcal{P}^n_0)$.

This result seems to support the idea that more information always leads to better decisions. The following results contradicts this intuitive idea.

Proposition 15 *Consider a fixed sequence $\mathcal{T}$, for instances such that $R \leq 1$ when the tour is operated using P_1^n and P_{1v}^n then*

$$\mathcal{L}(\mathcal{T}, \mathcal{P}_1^n) \;\leq\; \mathcal{L}(\mathcal{T}, \mathcal{P}_{1v}^n)$$

Proof. From theorems 5 and 7, $\Phi(\mathcal{T}, \mathcal{P}_{1v}^n)$ can be determined using network $\mathcal{G}_0$ because $R \leq 1$. $\mathcal{P}_1^n$ is a history independent recourse policy, so by Lemma 3 and 4, $\Phi(\mathcal{T}, \mathcal{P}_1^n)$ can also be determine using $\mathcal{G}_0$. Comparing recourse conditions $C^1(i)$ for both policies, observe that any node i satisfying Proposition 10 also satisfies Proposition 8 and not viceversa, hence the arcs set $(s, (i, 1))$ in $\mathcal{G}_0$ under $\mathcal{P}_1^n$ is a subset of the corresponding arcs set under $\mathcal{P}_{1v}^n$ with the same arc cost; therefore $\Phi(\mathcal{T}, \mathcal{P}_1^n) \leq \Phi(\mathcal{T}, \mathcal{P}_{1v}^n)$ and the result follows. □

Proposition 15 establishes an interesting and counterintuitive result. For any tour with at most one recourse action under both policies, $\mathcal{P}_1^n$ dominates over $\mathcal{P}_{1v}^n$; this implies that in this case a policy that requires less information always performs better. This type of result actually raises questions regarding the true benefit of perfect information in a worst-case context, and naturally leads to questions like, what is the gap between $\mathcal{L}(\mathcal{T}, \mathcal{P}_{nv}^n)$ and $\mathcal{L}(\mathcal{T}, \mathcal{P}_n^n)$? The following result gives the answer.

Theorem 8 *For any fixed sequence $\mathcal{T}$*

$$\mathcal{L}(\mathcal{T}, \mathcal{P}_{nv}^n) \;=\; \mathcal{L}(\mathcal{T}, \mathcal{P}_n^n)$$

Proof. From Proposition 13 it follows that $\mathcal{L}(\mathcal{T}, \mathcal{P}_{nv}^n) = \mathbf{ORP}(\mathcal{T}_1, Q, \overline{d})$. From Proposition 12 it follows that $\mathcal{L}(\mathcal{T}, \mathcal{P}_n^n) = \mathbf{ORP}(\mathcal{T}_1, Q, \overline{d})$. □

This result implies that for the case of no advance demand information, recourse policy $\mathcal{P}_n^n$ always achieves the same objective value than $\mathcal{P}_{nv}^n$, which assumes perfect demand visibility over the demands of all customers in the tour. Therefore, from a worst-case perspective the value of perfect demand visibility is zero.

Theorem 9 *Consider a fixed sequence $\mathcal{T}$, for instances with triangular inequality on travel time such that $R \leq 1$ when the tour is operated using $\mathcal{P}_0^n$, $\mathcal{P}_1^n$ and $\mathcal{P}_{1v}^n$ then*

$$\mathcal{L}(\mathcal{T}, \mathcal{P}_{nv}^n) \ = \ \mathcal{L}(\mathcal{T}, \mathcal{P}_n^n) \ \leq \ \mathcal{L}(\mathcal{T}, \mathcal{P}_1^n) \ \leq \ \mathcal{L}(\mathcal{T}, \mathcal{P}_{1v}^n) \ \leq \ \mathcal{L}(\mathcal{T}, \mathcal{P}_0^n)$$

Example. For $R > 1$ the previous result is not true. Consider the tour $\mathcal{T}$ in Figure 18 for a vehicle with capacity $Q = 5$. Observe that $\mathrm{L}(\mathcal{T}) = 8$.

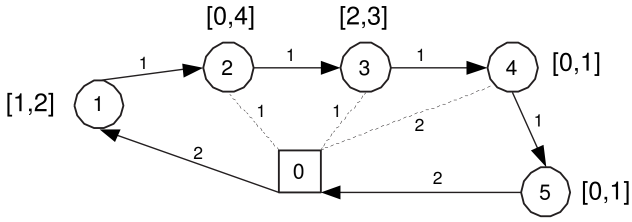

Figure 18: Tour $\mathcal{T}$ for $Q = 5$.

By evaluating the tour for demand realization $\overline{d}$, it can be seen that $R = 2$ if the tour is evaluated with recourse policy $\mathcal{P}_0^n$, $\mathcal{P}_1^n$ or $\mathcal{P}_{1v}^n$; therefore, for all three policies the adversarial problem can be solved using network $\mathcal{G}_0$. Figure 19 shows the corresponding $\mathcal{G}_0$ for each of these policies. Observe that $\Phi(\mathcal{T}, \mathcal{P}_0^n) = 4$, $\Phi(\mathcal{T}, \mathcal{P}_{1v}^n) = 3$ and $\Phi(\mathcal{T}, \mathcal{P}_1^n) = 5$; therefore $\mathcal{L}(\mathcal{T}, \mathcal{P}_{1v}^n) \ < \ \mathcal{L}(\mathcal{T}, \mathcal{P}_0^n) \ < \ \mathcal{L}(\mathcal{T}, \mathcal{P}_1^n)$. This example not only shows that the previous Theorem is not valid when there is more than one recourse, it also shows that there are instances in which anticipatory policy $\mathcal{P}_1^n$ actually performs worst than myopic policy $\mathcal{P}_0^n$.

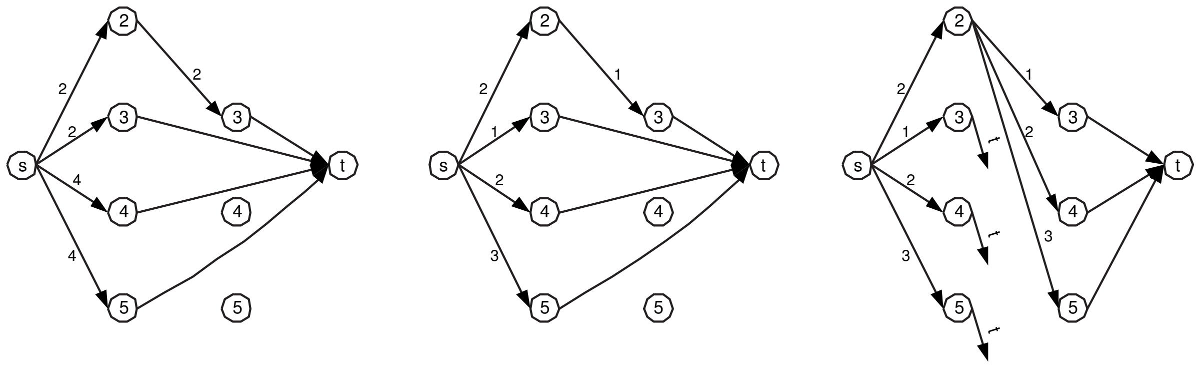

Figure 19: (a) $\mathcal{G}_0(\mathcal{P}_0^n)$ (b) $\mathcal{G}_0(\mathcal{P}_{1v}^n)$ (b) $\mathcal{G}_0(\mathcal{P}_1^n)$.

4.5 Solution approach

In this section, we develop heuristic methods for solving both RVRPSD and VRPSDDC. Developing an exact solution approach is beyond the scope of this dissertation. Nonetheless, we present a simple result on how to solve single vehicle instances of the RVRPSD for recourse policies $\mathcal{P}^n_{nv}$ and $\mathcal{P}^n_n$.

Let L_{VRP} denote the optimal cost of the deterministic capacitated VRP on network $G = (V_0, A)$ with vehicles of capacity Q, using cost matrix $(l(i,j))$ and customer demand $\overline{d}$.

Theorem 10

$$\mathcal{L}(\mathcal{T}^\star, \mathcal{P}^n_{nv}) \;=\; \mathcal{L}(\mathcal{T}^\star, \mathcal{P}^n_n) \;=\; L_{VRP}$$

where $\mathcal{T}^\star$ is the optimal solution of the single vehicle RVRPSD.

Proof. Consider a tour $\mathcal{T}$ operated using policy $\mathcal{P}^n_n$. From Proposition 12 it follows $\mathcal{L}(\mathcal{T}, \mathcal{P}^n_n) = \mathbf{ORP}(\mathcal{T}_1, Q, \overline{d})$. Clearly, $\mathbf{ORP}(\mathcal{T}_1, Q, \overline{d})$ is the cost of a feasible solution to the associated deterministic VRP with demand vector $\overline{d}$; therefore,

$$L_{VRP} \;\leq\; \mathcal{L}(\mathcal{T}, \mathcal{P}^n_n). \tag{53}$$

Take the optimal solution to the deterministic VRP and build a tour $\mathcal{T}$ by arbitrarily sequencing tours in the optimal solution and deleting all edges that connect to the depot except for one in the first and last tours in the sequence; then connect any degree one vertices between consecutive tours in the sequence. It is clear that $L_{VRP} = \mathbf{ORP}(\mathcal{T}_1, Q, \overline{d}) = \mathcal{L}(\mathcal{T}, \mathcal{P}^n_n)$, which implies that the bound in expression (53) is tight. The result for $\mathcal{P}^n_{nv}$ follows from Theorem 8. □

Thus, for these two policies, optimal robust solutions can be found by solving a deterministic vehicle routing problem.

4.5.1 A Tabu Search Heuristic

We propose a Tabu Search heuristic for solving instances of the RVRPSD with a single vehicle, and instances of the VRPSDDC with multiple vehicles. The heuristic closely follows the ideas presented in Gendreau et al. [32], which computational studies show performs quite well on expectation minimization problems with stochastic demands and customers.

Initial solution

For the single vehicle RVRPSD, the initial solution is simply a random permutation of the customers. For the VRPSDDC, on the other hand, the initial solution simply places each customer in its own tour, identical to the initialization of the savings heuristic of Clarke and Wright [19] for the deterministic vehicle routing problem.

Neighborhood structure

The typical $N(p, q, x)$ neighborhood is used, where for a given solution x, the neighborhood includes all tours that are obtained by removing in turn one of q randomly selected customers, and inserting each of them either immediately after or immediately before each of its p nearest neighbors.

Tabu moves

If vertex (customers or depot) is moved in iteration ν, then any move involving that vertex, either among the q randomly selected or among the p nearest neighbors is tabu until iteration $\nu + \theta$ where θ is randomly selected in interval $[N - 5, N]$.

Aspiration criteria

The search process moves from one iteration to the next considering only nontabu solutions in the neighborhood associated with the current solution, unless a tabu solution in the neighborhood improves the best solution found thus far.

Steps of the tabu search algorithm

STEP 1: (*Initialization*)

Construct a randomly generated solution x, let T^* be the cost of x and let $x^* = x$. Set $p = \min\{N - 1, 5\}$ and $q = \min\{N - 1, 5\}$. Set $t_0 = 0$, $t_1 = 0$ and $t_2 = 30N$, where t_0 is the iteration counter, t_1 is the number of iterations in which the best solution found thus far has not improved and t_2 is the maximum number of iterations allowed without one of such improvements.

STEP 2: (*Neighborhood search*)

Set $t_0 = t_0 + 1$ and $t_1 = t_1 + 1$. Consider all moves in $N(p, q, x)$ and build a LIST where all moves are sorted in nondecreasing order of their cost. Let y be the first move in the list, if it is not a tabu move or it improves the best solution found thus far then let $x = y$; else continue examining LIST until such solution is found.

STEP 3: (*Incumbent update*)

If $T(x) < T^*$, set $T^* = T(x)$, $x^* = x$ and $t_1 = 0$. Else, set $t_1 = t_1 + 1$. If $t_1 < t_2$, go to STEP 2; otherwise go to STEP 4.

STEP 4: (*Intensification or termination*)

If $t_2 = 30N$, set $t_1 = 0$, $t_2 = 10N$, $p = \min\{N - 1, 10\}$, $q = N - 1$ and go to STEP 2. Otherwise; stop, x^* is the best solution found.

Note that in cases where the tabu search heuristic is used to solve single vehicle RVRPSD instances, evaluation of the objective function cost $T(x)$ requires evaluation of the maximum tour duration $\mathcal{L}(\mathcal{T}, \mathcal{P})$ using the techniques described in earlier sections, given recourse policy $\mathcal{P}$. When solving multiple vehicle VRPSDDC instances, evaluation of $\mathcal{L}(\mathcal{T}_k, \mathcal{P})$ is conducted for each vehicle tour $\mathcal{T}_k$ as necessary in order to guarantee solution feasibility with respect to the duration constraint. In this case, the objective function is calculated directly by expectation. Unlike the approach in Gendreau et al. [32], the exact expected cost of the complete solution is determined for each potential move.

4.6 Computational Study

A computational study was conducted in order to obtain insights about the characteristics of robust minimax solutions obtained using different recourse policies for single vehicle instances of the Robust Vehicle Routing Problem with Stochastic Demands, and to understand the potential expected cost impacts of imposing hard duration constraints on stochastic routing problems by solving Vehicle Routing Problem with Stochastic Demands and Duration Constraints instances.

The computational study is therefore divided into two parts. In the first part we focus our attention on the single vehicle RVRPSD and proceed by:

1. Performing a comparison of solutions obtained using recourse policy $\mathcal{P}_0^n$ for the robust and expectation minimization approaches for single tour problems.

2. Performing a comparison of solutions obtained with our robust approach for myopic recourse policy $\mathcal{P}_0^n$ and anticipatory recourse policy $\mathcal{P}_n^n$ for single tour problems.

In the second part, we consider the multiple vehicle VRPSDDC and proceed by performing a comparison of expected solution cost for unconstrained instances, to those solved with more restrictive duration constraints.

All heuristic runs were performed on a computer running LINUX with dual 2.4GHz Intel Pentium processors and with 2GB of memory.

4.6.1 Part I: Single Vehicle RVRPSD

In this part of the computational study three factors are considered in the experimental design:

1. Number of customers. Three levels were considered: 10, 20 and 30.

2. Demand variability. Customers have homogeneous demand uncertainty intervals, *i.e.*, $\underline{d}(i)$ and $\overline{d}(i)$ is the same for any customer i. Parameter $\underline{d}(i)$ was set to 1 in all instances, and three different levels of $\overline{d}(i)$ were considered in order to capture different levels of demand variability: $\overline{d}(i) = 5, 11$ and 17. For comparisons using the expectation minimization approach, demand was assumed to be discrete uniform on the intervals, independent and identically distributed for all customers. For this type of stochastic demand, Ak and Erera [2] propose efficient methods that can be used to evaluate the objective function; such algorithms were employed in this study.

3. Average number of failures per tour. The value of Q was set using the following expression
$$Q = \frac{n}{r}\frac{\underline{d}(i) + \overline{d}(i)}{2},$$
where the value of r was set at $1, 2, 3$. The value of r can be interpreted as an approximation to the number of recourse actions if each customer demanded the

mean value. Observe that for this computational study Q is always integer given the values used to model demand and the number of customers.

A total of 30 base instances were randomly generated for this computational study, which corresponds to ten instances for each of the three levels of number of customers described above. These instances were obtained by randomly generating customer locations uniformly in $[0, 100]^2$ with the depot always located at point $(50, 50)$. Travel times were assumed equal to the Euclidean distance between points. All combinations of r and $\overline{d}(i)$ were considered for each of the ten instances to obtain a total of 270 instances.

4.6.1.1 Comparison of robust and expectation minimization approaches

First, we investigate the performance of recourse policy $\mathcal{P}_0^n$ under both the robust minimax and expectation minimization approaches. The objective is to understand the benefits and shortcomings of each approach, and to obtain insight into their key differences.

The Tabu search algorithm described in the previous section was used to solve each of the 270 instances for both the robust and the expectation minimization approaches; the only difference being the function called to evaluate the objective function. Further, for both approaches all instances were solved 4 different times starting from different randomly generated solutions, and the best of the four solutions was chosen for this analysis; therefore, a total of 2160 solution runs were performed for this study.

Let $\mathcal{T}_R$ and $\mathcal{T}_E$ represent the best solutions found under the robust and expectation minimization approaches respectively for a particular instance. We use the term 'robust tour' to refer to $\mathcal{T}_R$. In order to study and analyze the characteristics of these tours, we use four key metrics or performance measures. For all instances and for all factors considered in this computational study, we calculate:

BestRob $\equiv$ $\mathcal{L}(\mathcal{T}_R, \mathcal{P}_0^n)$: maximum cost of the robust solution.

BestExp $\equiv$ $\mathcal{L}_E(\mathcal{T}_E, \mathcal{P}_0^n)$: expected cost of $\mathcal{T}_E$.

MeanBestRob $\equiv$ $\mathcal{L}_E(\mathcal{T}_R, \mathcal{P}_0^n)$: expected cost of the robust solution.

MaxBestExp $\equiv$ $\mathcal{L}(\mathcal{T}_E, \mathcal{P}_0^n)$: maximum cost of $\mathcal{T}_E$.

It is easy to see that if $\mathcal{T}_R$ and $\mathcal{T}_E$ are the optimal solutions for the robust and expectation minimization approaches respectively, the following relations hold:

$$\mathcal{L}_E(\mathcal{T}_E, \mathcal{P}_0^n) \ \leq \ \mathcal{L}_E(\mathcal{T}_R, \mathcal{P}_0^n) \ \leq \ \mathcal{L}(\mathcal{T}_R, \mathcal{P}_0^n) \ \leq \ \mathcal{L}(\mathcal{T}_E, \mathcal{P}_0^n). \tag{54}$$

In this computational study we are interested in quantifying the differences in value of these four expressions. Observe that the difference between $\mathcal{L}(\mathcal{T}_E, \mathcal{P}_0^n)$ and $\mathcal{L}_E(\mathcal{T}_E, \mathcal{P}_0^n)$ corresponds to the maximum deviation from the expected cost for tour $\mathcal{T}_E$; that is, this expression captures the difference between the average and the worst-case cost for this tour. On the other hand, the difference between $\mathcal{L}_E(\mathcal{T}_R, \mathcal{P}_0^n)$ and $\mathcal{L}_E(\mathcal{T}_E, \mathcal{P}_0^n)$ captures the average cost increment of the robust tour over $\mathcal{T}_E$.

The average value of $\mathcal{L}_E(\mathcal{T}_E, \mathcal{P}_0^n)$ among all instances considered in this study was 507.3. The corresponding value for $\mathcal{L}(\mathcal{T}_E, \mathcal{P}_0^n)$ was 699.2, which implies that the percentage cost increment for $\mathcal{T}_E$ between the expected value and its maximum value is on average 37.8%. The average values of $\mathcal{L}_E(\mathcal{T}_R, \mathcal{P}_0^n)$ and $\mathcal{L}(\mathcal{T}_R, \mathcal{P}_0^n)$ were 522.7 and 677.7 respectively. Observe then that on average, the expected cost of $\mathcal{T}_R$ is only 3% larger than the expected cost of $\mathcal{T}_E$.

Figure 20 shows the average values of $\mathcal{L}_E(\mathcal{T}_E, \mathcal{P}_0^n)$, $\mathcal{L}_E(\mathcal{T}_R, \mathcal{P}_0^n)$, $\mathcal{L}(\mathcal{T}_R, \mathcal{P}_0^n)$ and $\mathcal{L}(\mathcal{T}_E, \mathcal{P}_0^n)$ over all instances, summarized by number of customers. Among all three levels of number of customers, the maximum percentage difference between $\mathcal{L}(\mathcal{T}_E, \mathcal{P}_0^n)$ and $\mathcal{L}_E(\mathcal{T}_E, \mathcal{P}_0^n)$ was 52.6% for 10 customers; for 20 and 30 customers this value was 33.3% and 31.3% respectively.

Comparing figures 20, 21 and 22, it can be observed that among the three factors considered in this computational study and their corresponding levels, the average number of failures per tour (r) was found to have the highest impact on the four performance measures that we consider here. Further, increasing the value of parameter r was found to have a larger effect on $\mathcal{L}$ than on $\mathcal{L}_E$. That is, by increasing the value of r by one unit, the increment found on the value of $\mathcal{L}$ was higher than on $\mathcal{L}_E$. Thus, by increasing the number of possible failures in a tour, the maximum cost appears to be more sensitive than the average cost.

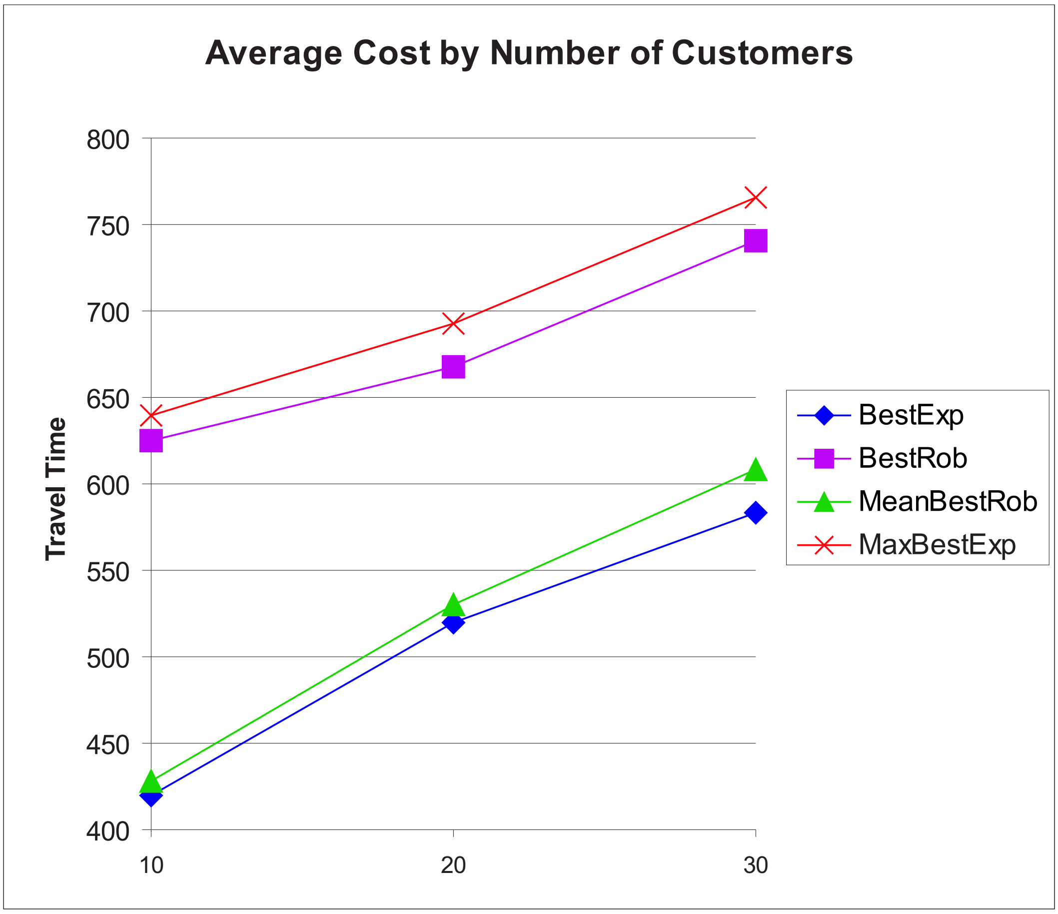

Figure 20: Average cost of all solutions over the 270 instances for each level of number of customers.

Increasing the demand variability seems to have the least effect on cost among the three factors considered. Further, function $\mathcal{L}_E$ is much less sensitive than $\mathcal{L}$ to an increase in customer demand variability. Also, $\mathcal{L}$ is much more sensitive when $\overline{d}(i)$ is increased from 5 to 11, than when it is increased from 11 to 17. This might be explained partially by our previous observation that for recourse policy $\mathcal{P}_0^n$, the worst-case demand does not necessarily correspond to an extreme point of the uncertainty space. Still, it is intuitively clear that increasing the size of the demand window might allow more flexibility to the adversary to find the worst-case demand that further increases the maximum cost. On the other hand, increasing the number of customers was found to have a slightly higher effect on $\mathcal{L}_E$ than on $\mathcal{L}$, especially between 10 and 20 customers.

In order to compare the characteristics of both solutions, $\mathcal{T}_R$ and $\mathcal{T}_E$, consider now the

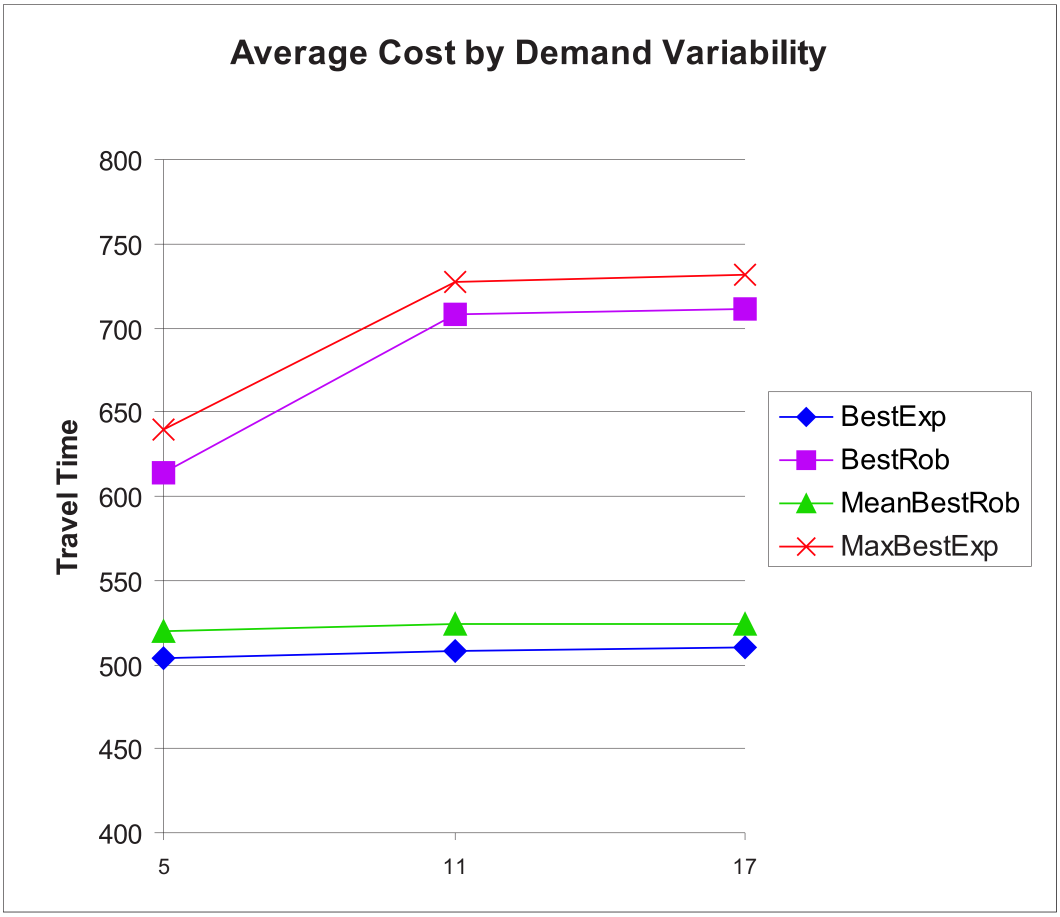

Figure 21: Average cost of all solutions over the 270 instances for each level of number of demand variability.

following performance measures:

1. WCE - Worst case of expected solution, defined by

$$WCE = \frac{\mathcal{L}(\mathcal{T}_E, \mathcal{P}_0^n) - \mathcal{L}_E(\mathcal{T}_E, \mathcal{P}_0^n)}{\mathcal{L}_E(\mathcal{T}_E, \mathcal{P}_0^n)},$$

which corresponds to the maximum potential percentage increment in cost from the expected cost of the best solution found in the expectation minimization approach.

2. WCR - Worst case of the robust solution, defined as

$$WCR = \frac{\mathcal{L}(\mathcal{T}_R, \mathcal{P}_0^n) - \mathcal{L}_E(\mathcal{T}_R, \mathcal{P}_0^n)}{\mathcal{L}_E(\mathcal{T}_E, \mathcal{P}_0^n)},$$

which corresponds to the maximum potential percentage increment in cost of the robust solution with respect to its expected cost.

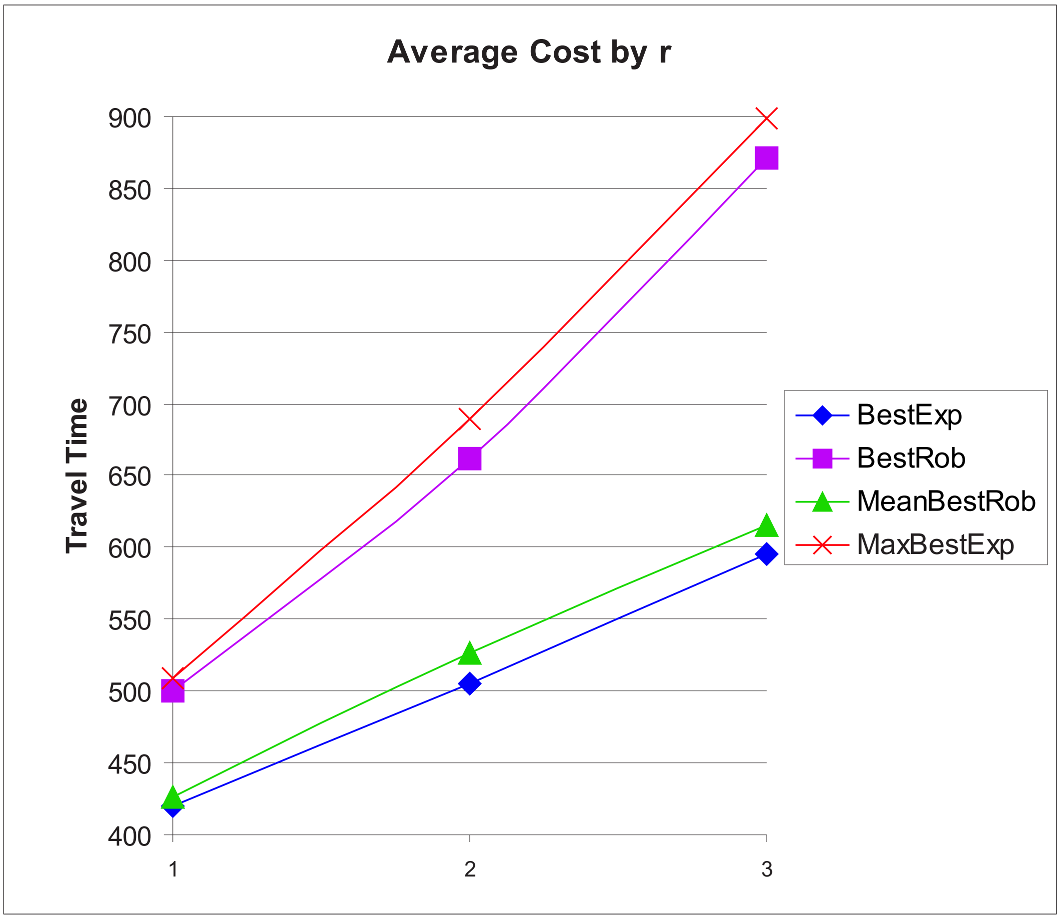

Figure 22: Average cost of all solutions over the 270 instances for each level of parameter r.

3. POR - Price of robustness, defined by

$$POR = \frac{\mathcal{L}_E(\mathcal{T}_R, \mathcal{P}_0^n) - \mathcal{L}_E(\mathcal{T}_E, \mathcal{P}_0^n)}{\mathcal{L}_E(\mathcal{T}_E, \mathcal{P}_0^n)},$$

which corresponds to the percentage increment in cost of the expected value of the best robust tour with respect to the best solution of the expectation minimization approach. This is the average price to pay for having a robust solution that is less sensitive to worst-case demand.

4. WRE - Worst case over the robust tour of the expectation minimization solution, defined by

$$WRE = \frac{\mathcal{L}(\mathcal{T}_E, \mathcal{P}_0^n) - \mathcal{L}(\mathcal{T}_R, \mathcal{P}_0^n)}{\mathcal{L}_E(\mathcal{T}_E, \mathcal{P}_0^n)},$$

which corresponds to the maximum percentage increment over the max cost for the robust solution of the solution obtained using the expectation minimization approach. In other words, it is the potential price to pay in the worst case over the robust solution for using a solution that on average performs well.

Observe that each of these performance measures is computed using the same denominator ($\mathcal{L}_E(\mathcal{T}_E, \mathcal{P}_0^n)$), and thus they are readily comparable. The following three tables summarize these average performance indices across all 270 instances.

nCust	WCE	WCR	POR	WRE
10	48.5	42.8	2.2	3.6
20	32.2	25.4	2.2	4.6
30	30.4	22.1	4.3	4.2

width	WCE	WCR	POR	WRE
5	26.3	18.5	3.0	4.8
11	42.3	35.6	3.0	3.7
17	42.6	36.2	2.8	3.8

r	WCE	WCR	POR	WRE
1	21.5	17.8	1.7	2.2
2	37.1	27.7	3.9	5.6
3	52.5	44.8	3.1	4.6

The average value of WCE overall instances is 37%. Among the four measures considered in this study, WCE showed a surprisingly high value. For all three factors, it was always observed that its average value was over 20%. For 10 customers and $r = 3$, WCE took values of over 100% for the larger values of demand variability. Furthermore, observe that unlike $\mathcal{L}_E$ (see Figure 21), the value of WCE appears to be very sensitive to demand variability. This implies that solutions obtained using the traditional expectation minimization approach can experience worst-case demand realizations with considerable deviations from its expected cost.

The average value of WCR overall instances is 30%, seven units smaller than WCE. In the tables above, it can be seen that on average, for each of the three factors considered in this study, WCR is always at least five units smaller than WCE. This consistency across factors suggests that worst-case deviations from the expected cost are significantly less in the robust solutions than in the solutions obtained using expectation minimization.

The average value of POR is 3%, which implies that on average the expected cost deviation of the tours developed under the robust approach with respect to the cost of the best expectation minimization solution is not substantial, or at least much less considerable than the deviation from the mean cost for the worst-case demand for the tours obtained under the expectation minimization approach. Interestingly, it was always observed that on average $POR < WRE$, which implies that when using the expectation minimization approach, the potential price to pay in the worst case over the robust solution is on average higher than the price of robustness. Of course, one must be careful when interpreting such results since the price of robustness is paid each period, while the worst-case price is paid only during worst-case periods. Without knowledge of the decision-maker's attitude toward risk, it is not clear which solution is preferable.

4.6.1.2 Comparison of myopic and anticipatory robust approaches

In this section recourse policy $\mathcal{P}_n^n$ is compared to $\mathcal{P}_0^n$ in order to assess the benefit, if any, of anticipatory policies over myopic policies. The same test instances from the previous section were used. From Theorem 10, it follows that for single vehicle problems, an optimal tour for RVRPSD with $\mathcal{P}_n^n$ can be obtained by solving a deterministic VRP. Solutions for $\mathcal{P}_n^n$ were obtained using the optimal partitioning approach used in deterministic VRPs (see Beasley [4]).

As defined in the previous section, $\mathcal{T}_R$ represents the best solution obtained for recourse policy $\mathcal{P}_0^n$. Now, let $\mathcal{T}_R^a$ be the best solution obtained for anticipatory recourse policy $\mathcal{P}_n^n$. In addition to two of the performance measures already defined, $\mathcal{L}(\mathcal{T}_R, \mathcal{P}_0^n)$ and $\mathcal{L}_E(\mathcal{T}_R, \mathcal{P}_0^n)$, we are now also interested in quantifying and comparing the values of $\mathcal{L}(\mathcal{T}_R^a, \mathcal{P}_n^n)$, and $\mathcal{L}_E(\mathcal{T}_R^a, \mathcal{P}_n^n)$. Just as in the previous section, let:

BestRobAnt $\equiv$ $\mathcal{L}(\mathcal{T}_R^a, \mathcal{P}_n^n)$: maximum cost of the robust anticipatory solution.

MeanBestRobAnt $\equiv$ $\mathcal{L}_E(\mathcal{T}_R^a, \mathcal{P}_n^n)$: expected cost of the robust solution under anticipatory policy $\mathcal{P}_n^n$.

Table 12: Detailed average percentage increment on expected travel time

No. cust	Avg. Failures r	Demand Var.	$\mathcal{L}(\mathcal{T}_R, \mathcal{P}_0^n)$	$\mathcal{L}_E(\mathcal{T}_E, \mathcal{P}_0^n)$	$\mathcal{L}_E(\mathcal{T}_R, \mathcal{P}_0^n)$	$\mathcal{L}(\mathcal{T}_E, \mathcal{P}_0^n)$	WCE	WCR	POR	WRE
10	1	5	382.6	323.6	329.7	393.8	21.8	16.4	1.8	3.6
		11	401.2	326.0	333.0	404.8	24.3	21.2	2.1	1.1
		17	401.2	326.8	333.6	404.8	24.1	21.0	2.0	1.1
	2	5	505.8	417.7	424.0	525.9	26.0	19.6	1.5	5.0
		11	631.8	421.1	437.1	656.8	56.0	46.2	3.8	5.9
		17	631.8	422.2	438.0	656.8	55.6	45.9	3.8	5.9
	3	5	631.8	507.4	518.6	648.9	27.8	22.3	2.2	3.2
		11	1018.1	514.2	520.6	1033.9	101.1	96.9	1.3	3.2
		17	1018.1	516.5	522.8	1033.9	100.2	95.8	1.3	3.2
20	1	5	507.9	428.8	435.9	522.0	21.7	16.8	1.7	3.3
		11	514.5	432.9	437.2	523.4	21.0	18.0	1.3	2.0
		17	515.4	435.1	436.0	528.6	21.6	18.4	0.9	3.0
	2	5	623.1	515.6	531.2	660.6	28.1	17.8	3.0	7.3
		11	673.6	518.8	531.8	698.5	34.8	27.3	2.6	4.9
		17	673.8	520.0	535.3	698.9	34.4	26.5	3.1	4.8
	3	5	748.7	604.4	618.5	782.0	29.5	21.5	2.4	5.6
		11	877.2	608.6	623.4	909.2	49.6	41.7	2.4	5.4
		17	876.2	610.0	627.0	907.6	48.8	40.9	2.8	5.2
30	1	5	589.0	500.4	512.4	599.0	19.8	15.3	2.5	2.1
		11	591.9	504.8	511.7	603.9	19.7	15.9	1.6	2.4
		17	599.3	503.4	512.1	602.5	19.8	17.4	1.8	1.1
	2	5	702.9	573.9	610.4	740.3	29.0	16.1	6.4	6.8
		11	742.9	580.5	621.3	769.5	32.5	20.9	7.1	4.5
		17	772.4	582.4	605.6	798.8	37.2	28.7	4.0	4.8
	3	5	835.8	662.7	697.1	880.5	32.9	20.9	5.2	6.8
		11	919.8	667.7	702.7	944.3	41.4	32.5	5.2	3.9
		17	912.2	670.5	705.3	948.1	41.4	30.9	5.2	5.4

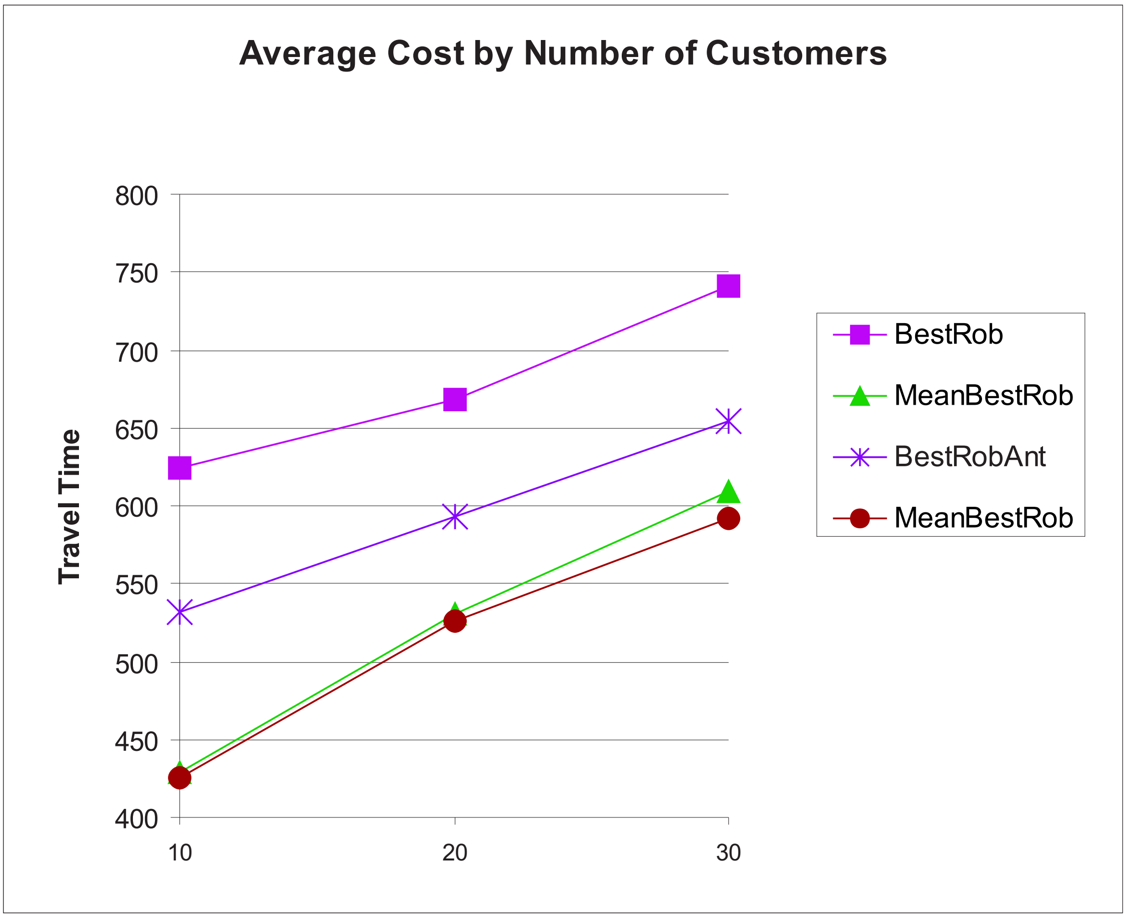

Figure 23: Average cost of all solutions over the 270 instances for each level of number of customers.

The value of $\mathcal{L}_E(\mathcal{T}_R^a, \mathcal{P}_n^n)$ is difficult to calculate because of the complexity of the recourse policy, which potentially implies solving an optimization problem at each customer location in the tour. Thus, a discrete event simulation model was developed in order to estimate the value of $\mathcal{L}_E(\mathcal{T}_R^a, \mathcal{P}_n^n)$. The best tour found for $\mathcal{P}_n^n$ using optimal partitioning in each of the 270 instances was simulated. For each tour $\mathcal{T}_R^a$, a total of 1800 simulation replications were used to calculate $\hat{\mathcal{L}}_E(\mathcal{T}_R^a, \mathcal{P}_n^n)$, an estimate of $\mathcal{L}_E(\mathcal{T}_R^a, \mathcal{P}_n^n)$. Each replication corresponds to the simulation of a vehicle traversing the entire tour once. As before, all customers' demands were assumed independent and identically distributed and discrete uniform over the interval in which demand takes values.

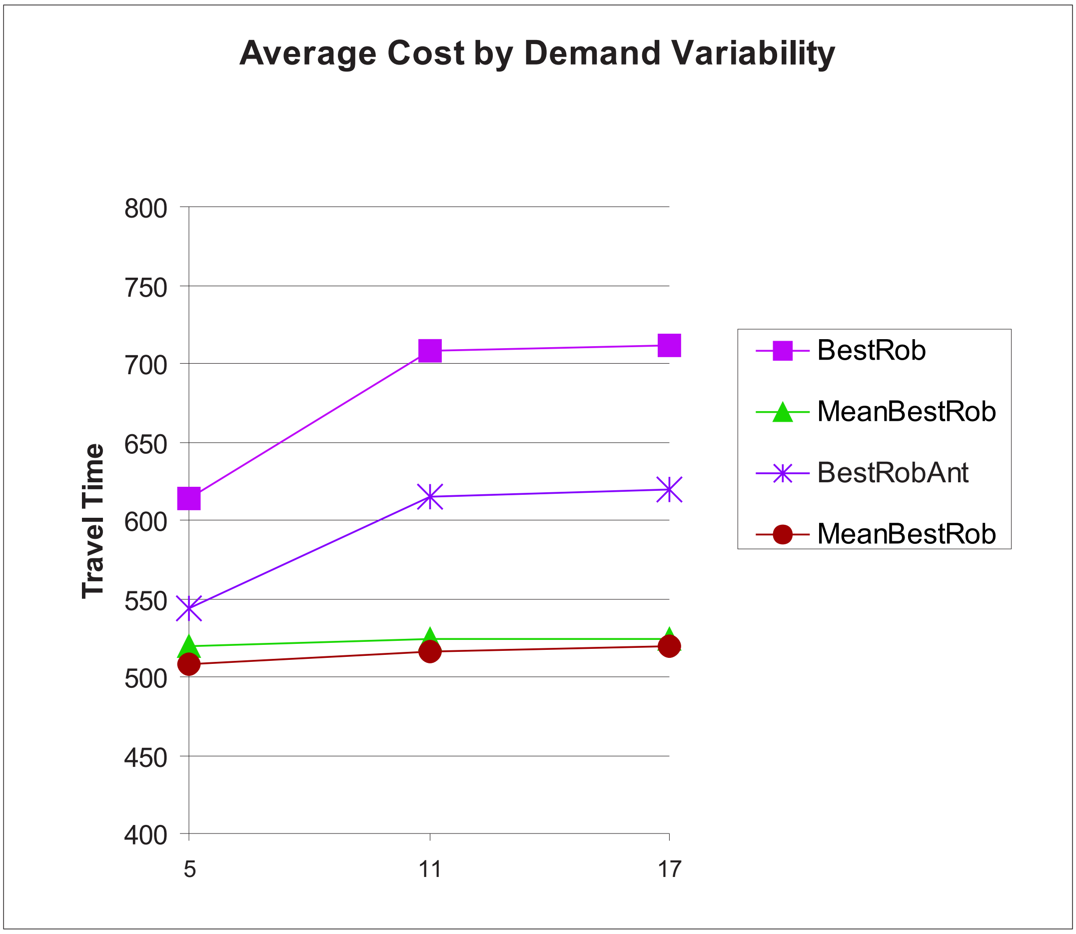

Figure 24: Average cost of all solutions over the 270 instances for each level of demand variability.

In order to understand better the differences between the myopic and anticipatory policies, we calculate the following metrics:

1. The percentage worst-case difference of the two policies

$$\text{PWD} = \frac{\mathcal{L}(\mathcal{T}_R, \mathcal{P}_0^n) - \mathcal{L}(\mathcal{T}_R^a, \mathcal{P}_n^n)}{\mathcal{L}(\mathcal{T}_R^a, \mathcal{P}_n^n)}$$

2. The percentage average case difference

$$\text{PAD} = \frac{\mathcal{L}_E(\mathcal{T}_R, \mathcal{P}_0^n) - \hat{\mathcal{L}}_E(\mathcal{T}_R^a, \mathcal{P}_n^n)}{\hat{\mathcal{L}}_E(\mathcal{T}_R^a, \mathcal{P}_n^n)}$$

PWD captures the worst-case difference in cost of the myopic policy with respect to the anticipatory policy. On the other hand, PAD captures this relation for the average case.

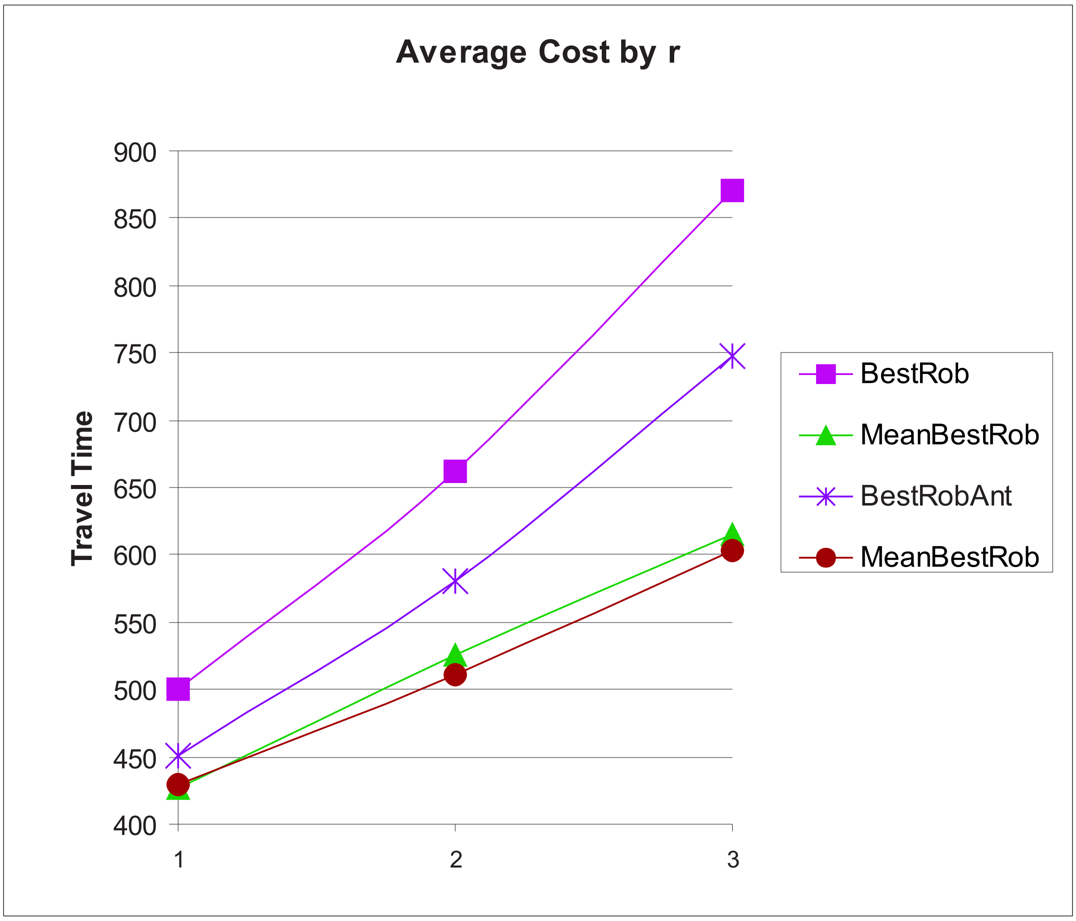

Figure 25: Average cost of all solutions over the 270 instances for each level r.

If $\mathcal{T}_R$ and $\mathcal{T}_R^a$ were the optimal solutions for recourse policies $\mathcal{P}_0^n$ and $\mathcal{P}_n^n$ respectively, it is clear that PWD ≥ 0, but no such relation is clear for PAD.

In our results from the computational experiment the average values across all instances for $\hat{\mathcal{L}}_E(\mathcal{T}_R^a, \mathcal{P}_n^n)$ and $\mathcal{L}(\mathcal{T}_R^a, \mathcal{P}_n^n)$ were 514.5 and 592.9 respectively. In turn, average values for $\mathcal{L}_E(\mathcal{T}_R, \mathcal{P}_0^n)$ and $\mathcal{L}(\mathcal{T}_R, \mathcal{P}_0^n)$ were 522.7 and 677.7 respectively; these results suggest that on average anticipatory policy $\mathcal{P}_n^n$ outperforms myopic policy $\mathcal{P}_0^n$ in both expected and maximum cost. Although the difference on average cost is relatively small, the reduction on maximum cost is quite significant. Average values for PWD and PAD overall instances were 14.31% and 1.38% respectively. This is an interesting result. In practice, it is very useful to reduce the gap between the average and maximum values as much as possible, since this greatly improves the predictability of the operational problem and its cost. Details of all results can be found in Table 13. Although on average over all 270 instances the value of PAD is positive, in some instances the the myopic policy actually performed better (*i.e.*, PAD < 0). This suggests that in terms of average-case, the anticipatory policy does not dominate the myopic policy.

Average results by number of customers are summarized in Figure 23. Interestingly, the gap between the worst and the average case reduces as the number of customers increases. This gap is 106.4 for 10 customers (25% with respect to the average case), and reduces almost by 40% for 30 customers. This suggests that anticipatory policies perform better as the number of customers in the tour increases. This is reasonable. To reduce costs, the anticipatory policy determines whether or not to take a recourse action at each customer in the tour by considering both its current capacity status and the travel times and demand intervals of the customers that remain to be visited. When the number of customers increases, more information and more potential decision points become available to dynamically adjust the recourse plan as the tour is traversed, and thus the advantages of using an anticipatory over a myopic policy becomes more important.

Figure 24 summarizes results by all levels of demand variability. For this factor, as well as for number of customers, on average the anticipatory policy performs better than the myopic policy. For the anticipatory policy, the average gap between the worst and the average case is only 7% for the lowest level of demand variability. This percentage increases to over 19% for the two highest levels. Although the anticipatory policy outperformed the myopic, it is very sensitive to demand variability because at every customer location the minimum cost recourse plan is determined assuming that the demand for the remaining customers in the tour will take the highest value (*i.e.*, $\overline{d}(i)$). This suggests that this gap can further be reduced and in general the average performance of the anticipatory policy can be improved by considering ways to address the conservatism of the recourse plan.

Among the three factors considered in this study, the average number of recourses r was found to have the highest effect on the performance of the anticipatory policy. For $r = 1$, the gap between the worst and the average case is less than 5%. It increases up to 24% for $r = 3$. Interestingly, on average for $r = 1$ the myopic policy actually outperformed the anticipatory policy.

Table 13: Detailed average results for anticipatory policy $\mathcal{T}_R^a$

No. cust.	Avg. Failures r	Demand Var.	$\mathcal{L}(\mathcal{T}_R^a, \mathcal{P}_n^n)$	$\hat{\mathcal{L}}_E(\mathcal{T}_R^a, \mathcal{P}_n^n)$	PWD	PAD
10	1	5	349.6	328.0	9.4%	0.5%
		11	362.8	333.0	10.6%	0.0%
		17	362.8	334.7	10.6%	-0.3%
	2	5	441.8	410.9	14.5%	3.2%
		11	541.9	419.1	16.6%	4.3%
		17	541.9	422.8	16.6%	3.6%
	3	5	541.9	503.0	16.6%	3.1%
		11	819.4	531.0	24.2%	-1.9%
		17	819.4	541.4	24.2%	-3.4%
20	1	5	457.8	441.7	10.9%	-1.3%
		11	470.9	443.0	9.3%	-1.3%
		17	470.9	443.9	9.4%	-1.8%
	2	5	558.0	521.4	11.7%	1.9%
		11	586.1	521.7	14.9%	1.9%
		17	586.1	524.3	15.0%	2.1%
	3	5	664.0	602.9	12.8%	2.6%
		11	772.5	615.6	13.6%	1.3%
		17	772.5	620.8	13.4%	1.0%
30	1	5	521.1	512.5	13.0%	0.0%
		11	527.9	515.1	12.1%	-0.7%
		17	535.5	517.5	11.9%	-1.1%
	2	5	627.4	586.9	12.0%	4.0%
		11	651.3	591.3	14.1%	5.1%
		17	686.0	596.2	12.6%	1.6%
	3	5	727.7	664.1	14.9%	5.0%
		11	805.2	671.2	14.2%	4.7%
		17	805.2	677.2	13.3%	4.1%

4.6.2 Part II: Multiple Vehicle VRPSDDC

In this section, we conduct a computational study of the Vehicle Routing Problem with Stochastic Demands and Duration Constraints (VRPSDDC). The objective of this study is to assess the effect of including duration constraints in the vehicle routing problem with stochastic demands. We are interested in quantifying the potential increases in total expected system travel time when duration constraints are enforced on a per tour basis, as well as understanding how different factors, like number of customers, vehicle capacity and demand variability can affect this performance. For simplicity, we restrict this study to the myopic recourse policy $\mathcal{P}_0^n$.

For each instance, we proceed by first solving a base unconstrained problem. Then, the instance is solved again by including constraints on the maximum duration of any individual vehicle tour, and the percentage difference in the value of the expected travel time of the two solutions is calculated for comparison. The value of the maximum duration (*i.e.*, the right-hand side of the duration constraint) is determined for each instance individually by identifying the maximum tour duration in the solution to the unconstrained problem, and then multiplying this value by a nonnegative reduction factor $\alpha < 1$.

In this computational exercise, four factors were considered:

1. *Number of customers*: five levels were considered; namely $20, 40, 60, 80$ and 100 customers.

2. *Demand variability*: customer demand was assumed discrete uniform, independent and identically distributed for all customers; hence the demand of any customer takes values over the same interval, *i.e.*, $\underline{d}(i)$ and $\overline{d}(i)$ is the same for any customer i. In this computational study $\underline{d}(i)$ was set to 1 in all instances, and three different levels of $\overline{d}(i)$ were considered in order to capture different levels of demand variability: $\overline{d}(i) = 5, 11$ and 17.

3. *Vehicle capacity*: three levels of vehicle capacity were considered: Q_1, Q_2 and Q_3. As in the previous computational study, the vehicle capacity was determined using expression
$$Q = \frac{n}{r} \frac{\underline{d}(i) + \overline{d}(i)}{2}.$$
In this context, the value of r can be interpreted as a tight lower bound on the number of vehicles required in a solution of the deterministic Vehicle Routing Problem (VRP) obtained by setting each customer demand equal to the corresponding expected value. When modifying the values of n and $\overline{d}(i)$, factors 1 and 2 in this computational study, the vehicle capacity is set so that the value of r remains constant. In that sense, instances generated using the same value of r, for different combinations of number of

customers and demand variability, are comparable; and differences in their solutions should, at least in part, be accounted to the stochastic component of the problem. The values of Q_1, Q_2 and Q_3 are determined by substituting in this expression $r = 5, 4$ and 3 respectively. Observe then that $Q_1 < Q_2 < Q_3$.

4. *Reduction factor* (α): three levels were considered, $0.75, 0.85$ and 0.95.

A total of 50 base instances were randomly generated for this study; this corresponds to 10 instances for each of the five levels of number of customers. The instances were obtained by randomly generating customer locations uniformly in $[0, 100]^2$ with the depot located at point $(50, 50)$. Again, travel times between points are determined using the Euclidean distance function. For each base instance, all combinations of levels of demand variability, vehicle capacity and reduction factor were considered; this generated 1350 different instances. Furthermore, for each base instance, an unconstrained problem was also solved for each level of demand variability and vehicle capacity. Thus, a total of 1800 different problem instances were solved for this computational study.

All instances were solved using the tabu search heuristic algorithm described in a previous section. In each iteration of the algorithm, primal feasibility was enforced. The initial solution considered by this algorithm, where each customer is assigned to a different vehicle, must satisfy the duration constraint; otherwise the problem is clearly infeasible. At any iteration, solutions in LIST that violate the maximum duration constraint are discarded. Each of the 1800 instances were solved ten different times using each time a different initial seed for the random number generator; among these ten solutions only the one with minimum expected travel time was selected for the subsequent result analysis.

The numerical results obtained show that duration constraints can significantly impact the structure of the tours in a best solution to the VRPSD, and furthermore may substantially increase the value of the total expected travel time of a solution. With respect to the unconstrained problem, the average percentage increase in expected travel time was found to be 6.6%. Although this average increase may not be very large, for some combinations of parameters considered in this study, average increases of over 12% were identified, and specific instances with increments of over 25% were found. Table 14 provides a list of general statistics of the results of this study summarized by reduction factor. Table 15 presents a more detailed summary of overall results.

Table 14: General statistics on the increment of expected travel time with respect to the unconstrained problem for all levels of reduction factor α.

Reduction Factor (α)	Mean	Max	St.dev.
0.75	10.1%	25.7%	5.6%
0.85	6.4%	20.1%	4.4%
0.95	3.2%	13.7%	3.0%

Number of customers was found to be a factor with a very important effect on the expected travel time when duration constraints are enforced. For problems with 20 customers, results show an average percentage increment of 2.5%. For 40 customers, this value jumped-up to 6.1%; this non-decreasing trend continued for 60 and 80 customers with values 7.8% and 8.3% respectively. Interestingly, for 100 customers this value decreased to 8.1%. Figure 26 shows the interactions found by number of customers and reduction factor on the percentage increment of expected travel time. It is interesting to observe that in the case with tightest duration constraints (*i.e.*, $\alpha = 0.75$) the increment reaches its maximum at 60 customers, and then decreases for 80 and 100 customers. These results suggest that in terms of expected travel time, systems where customers are geographically more concentrated per unit area may be less affected by enforcing duration constraints than systems with more geographically dispersed customers.

A consistent trend was identified between expected travel time increment and vehicle capacity: as the value of Q increases, the average percentage increment was also observed to increase. For levels Q_1, Q_2 and Q_3 the observed increment was $6.3\%, 6.5\%$ and 6.9% respectively. This suggests that in general fleets with larger vehicle capacity are more negatively affected when duration constraints are imposed as compared to fleets with lower vehicle capacity. This result has a clear intuitive explanation; when duration constraints are enforced, vehicle capacity might no longer be a binding constraint. Having a larger

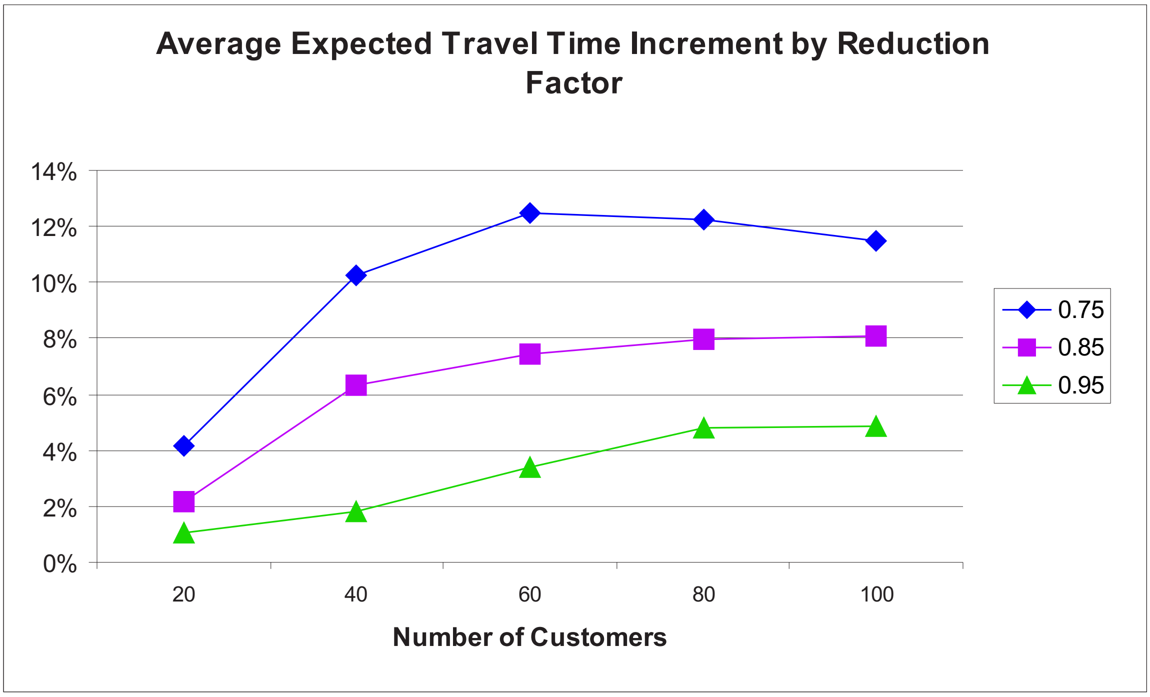

Figure 26: Average percentage increment in expected travel time for each level of reduction factor and number of customers.

vehicle capacity might no longer work in your benefit because it might not be possible to utilize the vehicles fully since by doing so, duration constraints may be violated.

A very interesting result follows from further considering number of customers and reduction factor in the analysis; this is illustrated in Figure 27. For specific circumstances, results indicate that systems operated using fleets with higher vehicle capacity may indeed be less affected by duration constraints than their counterparts. Specifically, in systems with tighter duration constraints and customers that are more geographically concentrated, the average increment in expected travel time due to duration constraints is less significant when operated with fleets with larger vehicle capacity. From Figure 27 it is clear that in the cases given by $\alpha = 0.75$, for 80 and 100 customers, a lower increment is observed when operated with vehicles of capacity Q_3 than with vehicles of capacity Q_1.

These results also imply that for systems with tighter duration constraints where customers are less concentrated per unit area, fleets with lower vehicle capacity are less affected by the introduction of durations constraints. In Figure 27 consider the cases given

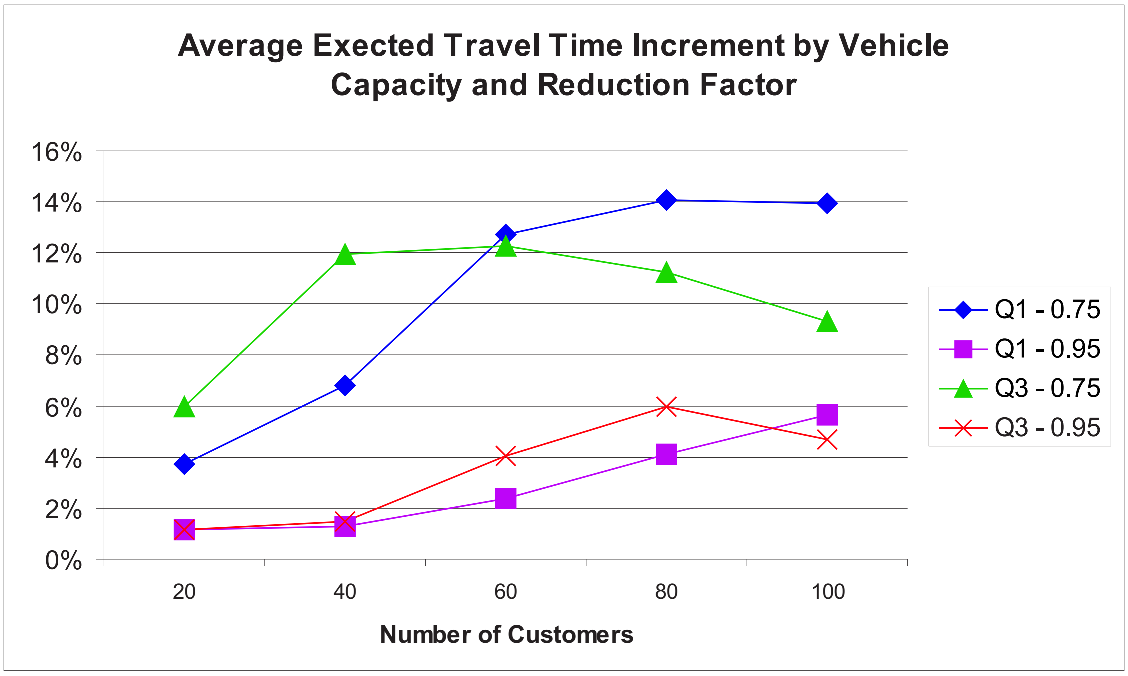

Figure 27: Average percentage increment in expected travel time for each level of reduction factor for vehicle capacity and demand variability.

by $\alpha = 0.75$ for 20 and 40 customers, Q_1 has a more moderate increment than with Q_3.

Although the idea of benefiting from fleets with higher vehicle capacity in the presence of duration constraints might seem counterintuitive, there is actually a very simple explanation. The reason for preferring vehicles with larger capacity when more customers need to be served is that although duration constraints tend to mitigate the benefits of this type of vehicle, when customers are highly concentrated per unit area, more can be packed into shorter duration tours to take advantage of the larger vehicle capacity. If this is not the case, and there are not enough nearby customers, this advantage of higher capacity cannot be exploited.

In our results, it was observed that increasing demand variability tends to also increase the average expected increment in total expected solution travel time. For levels of $\overline{d}(i)$ equals to 5, 11 and 17, the corresponding increment was $5.9\%, 6.8\%$ and 7% respectively. Still, when compared to number of customers, both vehicle capacity and demand variability seem to have a less substantial effect on transit times. Figure 28 shows the interaction or each of these factors by reduction factor α.

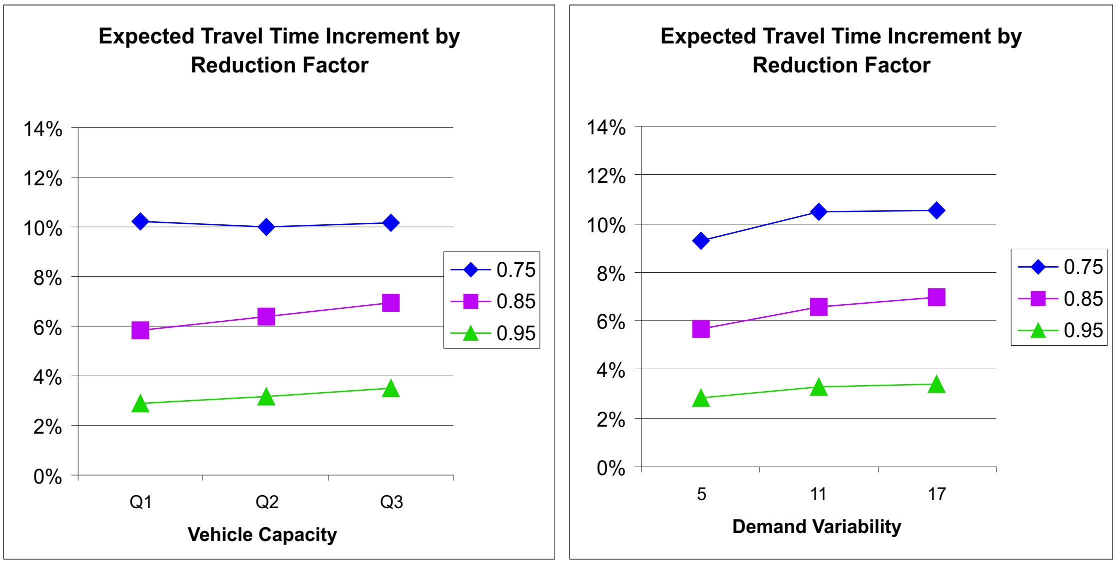

Figure 28: Average percentage increment in expected travel time for each level of reduction factor for vehicle capacity and demand variability.

Results for the first three levels of number of customers, $20, 40$ and 60, were pooled together into a single group called 'Low number of customers'; the remaining results were pooled into another group called 'High number of customers'. Results for these two groups are summarized in Figure 29. As previously observed, this figure shows that for the problems with higher customer density, it is advantageous to use vehicles with higher capacity; while for the other group, it is better to use vehicles with lower capacity. Interestingly, this figure also shows that the group with lower number of customers is less sensitive to increased demand variability than the group with larger number of customers. This can be explained at least in part by the fact that, as previously illustrated, when customers are not very concentrated geographically, it is not possible to use the vehicle's capacity to the maximum limit. Thus, each vehicle has excess capacity that can be used as a buffer to hedge against demand variability. Observe that this also explains why high density customer group, which can make a better utilization of vehicle capacity, is more sensitive when demand variability increases: it lacks the capacity buffer.

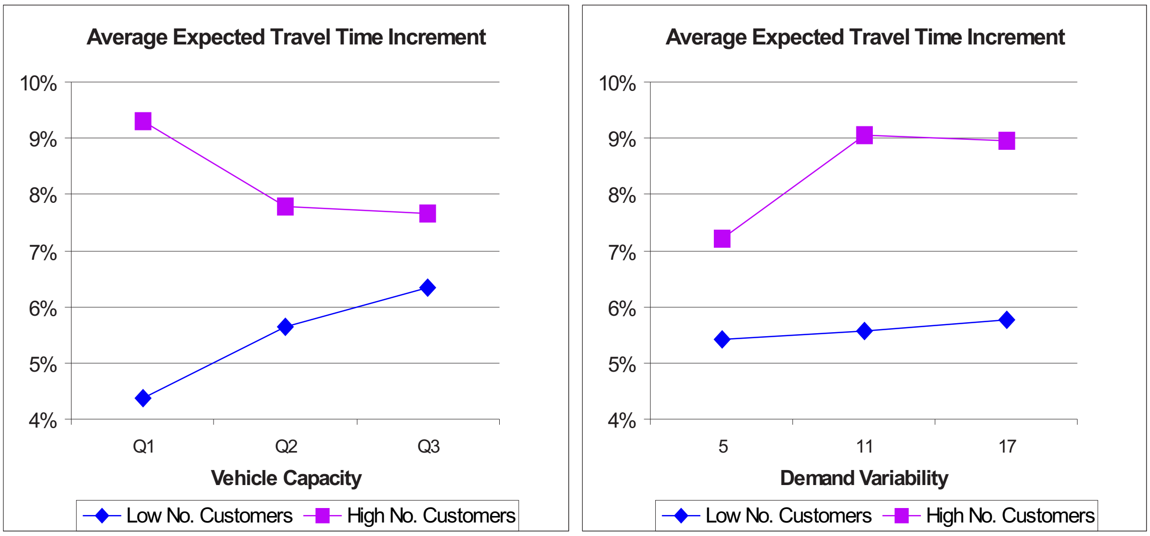

Figure 29: Average percentage increment in expected travel time for each level of reduction factor for for high and low number of customers.

Table 15: Detailed average percentage increment on expected travel time

No. customers	Vehicle capacity	Demand variability	$\alpha = 0.75$	$\alpha = 0.85$	$\alpha = 0.95$
20	Q_1	5	4.89%	1.89%	1.25%
		11	3.50%	1.66%	1.09%
		17	2.74%	1.69%	1.12%
	Q_2	5	5.05%	3.04%	1.05%
		11	1.47%	1.01%	0.81%
		17	1.57%	1.05%	0.83%
	Q_3	5	6.63%	3.28%	1.63%
		11	5.85%	2.95%	0.92%
		17	5.54%	2.87%	0.89%
40	Q_1	5	8.24%	4.28%	0.82%
		11	6.06%	3.76%	1.49%
		17	6.12%	3.26%	1.61%
	Q_2	5	10.56%	6.70%	2.85%
		11	12.79%	7.70%	2.82%
		17	12.60%	7.25%	2.40%
	Q_3	5	10.14%	6.88%	1.01%
		11	13.05%	8.26%	1.76%
		17	12.69%	9.06%	1.66%
60	Q_1	5	10.93%	3.88%	2.02%
		11	14.22%	6.26%	2.75%
		17	13.02%	7.11%	2.29%
	Q_2	5	11.25%	6.80%	4.00%
		11	11.05%	6.77%	3.12%
		17	15.06%	9.07%	3.92%
	Q_3	5	10.92%	7.88%	4.28%
		11	12.67%	9.19%	4.54%
		17	13.13%	9.79%	3.41%
80	Q_1	5	13.86%	8.62%	4.08%
		11	14.42%	9.06%	4.29%
		17	13.96%	8.84%	4.04%
	Q_2	5	8.41%	6.25%	3.24%
		11	12.14%	6.33%	3.92%
		17	13.93%	9.39%	5.56%
	Q_3	5	10.33%	6.08%	5.10%
		11	12.12%	8.75%	5.69%
		17	11.20%	8.52%	7.23%
100	Q_1	5	11.35%	7.74%	5.19%
		11	15.82%	10.12%	6.08%
		17	14.60%	9.60%	5.60%
	Q_2	5	8.87%	6.93%	3.39%
		11	12.55%	8.38%	5.33%
		17	12.37%	9.17%	4.02%
	Q_3	5	8.41%	4.62%	2.84%
		11	9.91%	8.18%	4.91%
		17	9.64%	7.87%	6.34%

CHAPTER V

CONCLUSIONS AND FUTURE RESEARCH

This dissertation investigates effective planning and control for fleet management in freight transportation.

Chapter 2 presents a large scale dynamic model for the management of a global fleet of tank containers. A computational study demonstrates that realistic size instances of this model can be solved with commercially available optimization software. The computational experiments also suggest that substantial cost savings and improvements in equipment utilization can be attained. Further analysis, however, reveals that in practice these savings may not be realized because demand uncertainty has been ignored, which may lead to situations where costly actions are required to deal with unexpected changes in demand.

The study described in Chapter 2 serves as motivation for the main focus of the thesis: the development of methodologies for generating fleet management plans that allow cost-effective control in the presence of demand uncertainty. We investigate approaches based on robust optimization as opposed to more traditional expectation minimization models.

Chapter 3 presents a robust optimization approach for generating empty repositioning plans. Starting from a deterministic model currently in use at some freight carriers, we develop a model that explicitly incorporates demand uncertainty in a simple, yet flexible manner, and that can easily be understood and implemented by practitioners. Risk aversion is also included in a way that we believe expert fleet managers can easily relate to. Computational evidence suggests that the approach is tractable, generates plans with a clear intuitive interpretation of how inventory is managed to hedge against demand uncertainty, and that the price that has to be paid for robustness remains within reasonable limits. The approach, although based on sophisticated mathematical optimization concepts, has simple input requirements, and allows for easy interpretation of results. Therefore, we believe it may be readily accepted by freight carriers.

Chapter 4 presents robust optimization techniques for the development of fixed routes for local delivery operations under various control policies. Computational experiments demonstrate that our robust optimization approach generates solutions with expected costs that compare favorably to those obtained with traditional expectation minimization techniques, but also that perform much better in worst-case scenarios. This is especially true when anticipatory recourse policies are considered. We further illustrate how the techniques can be applied to enforce tour duration constraints, which are often important in practical routing applications but are usually ignored by the research literature. A computational study shows that considering tour duration constraints can have a significant effect on the expected cost of the fixed routes.

The findings and results of this dissertation suggest a few interesting areas for further research.

Our findings clearly demonstrate the benefits of incorporating redundant capacity (or buffers) into fleet management plans to hedge against demand uncertainty thereby allowing cost-effective control. This redundant capacity, of course, comes at a price: the price of robustness. It will be interesting to study a kind of "dual problem" in which we try to maximize the robustness of a plan given a certain robustness budget. If we are willing to accept an increase in planned cost of say two percent, how much robustness can we achieve? A deeper analysis of how robustness is achieved in different contexts, *i.e.*, what hedging mechanism are used in different situations, may lead to insights that can be captured in simple rules of thumb that may be of value for practitioners.

Our computational experiments with the vehicle routing problem with stochastic demands show that the use of anticipatory strategies within robust optimization approaches can significantly reduce worst-case recovery costs at a minimal increase in expected costs. Almost all of the research in the area of vehicle routing problem with stochastic demand has focused on simple myopic policies. Our efforts suggest that much may be gained by considering anticipatory strategies and this should be a fruitful area of further research.

REFERENCES

[1] Abrache, J., Crainic, T. G., and Gendreau, M., "A new decomposition algorithm for the deterministic dynamic allocation of empty containers," Tech. Rep. CRT-99-49, Centre de Recherche sur les Transport, 1999.

[2] Ak, A. and Erera, A., "A paired-vehicle recourse strategy for the vehicle routing problem with stochastic demand," *Transportation Science*, vol. to appear, 2006.

[3] Atamturk, A. and Zhang, M., "Two-stage robust network flow and design under demand uncertainty," *Operations Research*, vol. to appear, 2006.

[4] Beasley, J., "Route first - cluster second methods for vehicle routing," *Omega*, vol. 11, pp. 403–408, 1983.

[5] Ben-Tal, A., Goryashko, A., Guslitzer, E., and Nemirovski, A., "Adjustable robust solutions of uncertain linear programs," *Mathematical Programming*, vol. 99, pp. 351–376, 2004.

[6] Ben-Tal, A. and Nemirovski, A., "Robust convex optimization," *Mathematics of Operations Research*, vol. 23, pp. 769–805, 1998.

[7] Ben-Tal, A. and Nemirovski, A., "Robust solutions of linear programming problems contaminated with uncertain data," *Mathematical Programming*, vol. 88, pp. 411–424, 2000.

[8] Bertsimas, D., "A vehicle routing problem with stochastic demand," *Operations Research*, vol. 40, pp. 574–585, 1992.

[9] Bertsimas, D., Pachamanova, D., and Sim, M., "Robust linear optimization under general norms," *Operations Research Letters*, vol. 32, pp. 510–516, 2004.

[10] Bertsimas, D. and Sim, M., "Robust discrete optimization and network flows," *Mathematical Programming*, vol. 98, pp. 49–71, 2003.

[11] Bertsimas, D. and Sim, M., "The price of robustness," *Operations Research*, vol. 52, pp. 35–53, 2004.

[12] Bertsimas, D. and Simchi-Levi, D., "The new generation of vehicle routing research: robust algorithms addressing uncertainty," *Operations Research*, vol. 44, pp. 286–304, 1994.

[13] Bertsimas, D. and Thiele, A., "A robust optimization approach to supply chain management," in *Proceedings of the 10th International Integer Programming and Combinatorial Optimization (IPCO) Conference*, pp. 86–100, 2004.

[14] Bertsimas, D. and Thiele, A., "A robust optimization approach to inventory theory," *Operations Research*, vol. 54, pp. 150–168, 2006.

[15] Bourbeau, B., Crainic, T., and Gendron, B., "Branch-and-bound parallelization strategies applied to a depot location and container fleet management problem," *Parallel Computing*, vol. 26, pp. 27–46, 2000.

[16] Cheung, R. and Chen, C., "A two-stage stochastic network model and solution methods for the dynamic empty container allocation problem," *Transportation Science*, vol. 32, pp. 142–162, 1998.

[17] Cheung, R. and Powell, W., "Models and algorithms for distribution problems with uncertain demands," *Transportation Science*, vol. 30, pp. 43–59, 1996.

[18] Choong, S. T., Cole, M. H., and Kutanoglu, E., "Empty container management for intermodal transportation networks," *Transportation Research Part E*, vol. 38, pp. 423–438, 2002.

[19] Clarke, C. and Wright, J., "Scheduling of vehicles from a central depot to a number of delivery points," *Operations Research*, vol. 12, pp. 568–581, 1964.

[20] Cook, W., Cunningham, W., Pulleyblank, W., and Schrijver, A., *Combinatorial Optimization*. New York: John Wiley and Sons, 1998.

[21] Crainic, T. G., "Service network design in freight transportation," *European J. Operational Research*, vol. 122, pp. 272–288, 2000.

[22] Crainic, T. G., "A survey of optimization models for long-haul freight transportation," in *Handbooks of Transportation Science* (Hall, R. W., ed.), pp. 451–516, Kluwer, 2003.

[23] Crainic, T. G., Gendreau, M., and Dejax, P., "Dynamic and stochastic models for the allocation of empty containers," *Operations Research*, vol. 41, pp. 102–126, 1993.

[24] Crainic, T. G. and Laporte, G., "Planning models for freight transportation," *European J. Operational Research*, vol. 97, pp. 409–438, 1997.

[25] Crainic, T., Gendreau, M., and Dejax, P., "Dynamic stochastic models for the allocation of empty containers," *Operations Research*, vol. 41, pp. 102–126, 1993.

[26] Dror, M., Laporte, G., and Trudeau, P., "Vehicle routing with stochastic demands: properties and solutions frameworks," *Transportation Science*, vol. 23, pp. 166–176, 1989.

[27] Dror, M. and Trudeau, P., "Stochastic vehicle routing with modified savings algorithm," *European Journal of Operational Research*, vol. 23, pp. 228–235, 1986.

[28] Erera, A., Morales, J., and Savelsbergh, M., "Global intermodal tank container management for the chemical industry," *Transportation Research Part E: Logistics and Transportation Review*, vol. 41, pp. 551–566, 2005.

[29] Frantzeskakis, L. and Powell, W., "A successive linear approximation procedure for stochastic, dynamic vehicle allocation problems," *Transportation Science*, vol. 24, pp. 40–57, 1990.

[30] Gendreau, M., Laporte, G., and Seguin, R., "An exact algorithm for the vehicle routing problem with stochastic demands and customers," *Transportation Science*, vol. 29, pp. 143–155, 1995.

[31] Gendreau, M., Laporte, G., and Seguin, R., "Stochastic vehicle routing," *European Journal of Operational Research*, vol. 88, pp. 3–12, 1996.

[32] Gendreau, M., Laporte, G., and Seguin, R., "A tabu search heuristic for the vehicle routing problem with stochastic demands and customers," *Operations Research*, vol. 44, pp. 469–477, 1996.

[33] Godfrey, G., , and Powell, W., "An adaptive dynamic programming algorithm for single-period fleet management problems I: Single period travel times," *Transportation Science*, vol. 36, pp. 21–39, 2002.

[34] Godfrey, G. and Powell, W., "An adaptive dynamic programming algorithm for single-period fleet management problems II: Multiperiod travel times," *Transportation Science*, vol. 36, pp. 40–54, 2002.

[35] Jansen, B., Swinkels, P., Teeuwen, G., van Antwerpen de Fluiter, B., and Fleuren, H., "Operational planning of a large-scale multi-modal transportation system," *European J. Operational Research*, vol. 156, pp. 41–53, 2004.

[36] Jordan, W. and Turnquist, M., "A stochastic, dynamic network model for railroad car distribution," *Transportation Science*, vol. 17, pp. 123–145, 1983.

[37] Kouvelis, P. and Yu, G., *Robust Discrete Optimization and Its Applications*, vol. 14 of *Nonconvex Optimization and Its Applications*. Dordrecht, Germany: Kluwer Academic, 1997.

[38] Laporte, G., Louveaux, F., and Hamme, L., "An integer l-shaped algorithm for the capacitated vehicle routing problem with stochastic demands," *Operations Research*, vol. 50, pp. 415–423, 2002.

[39] Laporte, G., Louveaux, F., and Mercure, H., "Models and exact solutions for a class of stochastic location-routing problems," *European Journal of Operational Research*, vol. 39, pp. 71–78, 1989.

[40] Misra, S., "Linear programming of empty wagon disposition," *Rail International*, vol. 3, pp. 151–158, 1972.

[41] Powell, W., "A stochastic formulation of the dynamic vehicle allocation problem," *Transportation Science*, vol. 20, pp. 117–129, 1986.

[42] Powell, W., "A stochastic formulation of the dynamic assignment problem, with an application to truckload motor carriers," *Transportation Science*, vol. 30, pp. 195–219, 1996.

[43] Powell, W., "Dynamic models of transportation operations," in *Handbooks in Operations Research and Management Science: Supply Chain Management* (Graves, S. and de Tok, T. A. G., eds.), pp. 677–756, Amsterdam: Elsevier, 2003.

[44] Powell, W. and Carvalho, T., "Real-time optimization of containers and flatcars for intermodal operations," *Transportation Science*, vol. 32, pp. 108–126, 1998.

[45] Powell, W., Jaillet, P., and Odoni, A., "Stochastic and dynamic networks and routing," in *Network Routing* (Ball, M., Magnanti, T., Monma, C., and Nemhauser, G., eds.), vol. 8 of *Handbooks in Operations Research and the Management Sciences*, pp. 141–295, Amsterdam: Elsevier Science, 1995.

[46] Powell, W. and Topaloglu, H., "Fleet management," in *Applications of Stochastic Programming* (Wallace, S. and Ziemba, W., eds.), Math Programming Society SIAM Series in Optimization, 2005.

[47] Savelsbergh, M. and Goetschalckx, M., "A comparison of the efficiency of fixed versus variable vehicle routes," *Journal of Business Logistics*, vol. 16, pp. 163–187, 1995.

[48] Shen, W. S. and Khoong, C. M., "A DSS for empty container distribution planning," *Decision Support Systems*, vol. 15, pp. 75–82, 1995.

[49] Soyster, A., "Convex programming with set-inclusive constraints and applications to inexact linear programming," *Operations Research*, vol. 21, pp. 1154–1157, 1973.

[50] Stewart, W. and Golden, B., "Stochastic vehicle routing: a comprehensive approach," *European Journal of Operational Research*, vol. 14, pp. 371–385, 1983.

[51] Stolt-Nielsen Transportation Group, "Private communication." 2004.

[52] Topaloglu, H. and Powell, W., "An algorithm for approximating piecewise linear concave functions from sample gradients," *Operations Research Letters*, vol. 31, pp. 66–76, 2003.

[53] Topaloglu, H. and Powell, W., "Dynamic programming approximations for stochastic, time-staged integer multicommodity flow problems," *Informs Journal on Computing*, vol. to appear, 2004.

[54] White, W., "Dynamic transshipment networks: an algorithm and its application to the distribution of empty containers," *Networks*, vol. 2, pp. 211–236, 1972.

[55] Yang, W., Mathur, K., and Ballou, R., "Stochastic vehicle routing problem with restocking," *Transportation Science*, vol. 34, pp. 99–112, 2000.