Université de Montréal

Optimization models and methods for real-time transportation planning in forestry

par
Amine Amrouss

Département d'informatique et de recherche opérationnelle
Faculté des arts et des sciences

Thèse présentée à la Faculté des des arts et des sciences
en vue de l'obtention du grade de Philosophiæ Doctor (Ph.D.)
en informatique

Avril, 2016

RÉSUMÉ

Lors du transport du bois de la forêt vers les usines, de nombreux événements imprévus peuvent se produire, événements qui perturbent les trajets prévus (par exemple, en raison des conditions météo, des feux de forêt, de la présence de nouveaux chargements, etc.). Lorsque de tels événements ne sont connus que durant un trajet, le camion qui accomplit ce trajet doit être détourné vers un chemin alternatif. En l'absence d'informations sur un tel chemin, le chauffeur du camion est susceptible de choisir un chemin alternatif inutilement long ou pire, qui est lui-même "fermé" suite à un événement imprévu. Il est donc essentiel de fournir aux chauffeurs des informations en temps réel, en particulier des suggestions de chemins alternatifs lorsqu'une route prévue s'avère impraticable. Les possibilités de recours en cas d'imprévus dépendent des caractéristiques de la chaîne logistique étudiée comme la présence de camions auto-chargeurs et la politique de gestion du transport. Nous présentons trois articles traitant de contextes d'application différents ainsi que des modèles et des méthodes de résolution adaptés à chacun des contextes.

Dans le premier article, les chauffeurs de camion disposent de l'ensemble du plan hebdomadaire de la semaine en cours. Dans ce contexte, tous les efforts doivent être faits pour minimiser les changements apportés au plan initial. Bien que la flotte de camions soit homogène, il y a un ordre de priorité des chauffeurs. Les plus prioritaires obtiennent les volumes de travail les plus importants. Minimiser les changements dans leurs plans est également une priorité. Étant donné que les conséquences des événements imprévus sur le plan de transport sont essentiellement des annulations et/ou des retards de certains voyages, l'approche proposée traite d'abord l'annulation et le retard d'un seul voyage, puis elle est généralisée pour traiter des événements plus complexes. Dans cette approche, nous essayons de re-planifier les voyages impactés durant la même semaine de telle sorte qu'une chargeuse soit libre au moment de l'arrivée du camion à la fois au site forestier et à l'usine. De cette façon, les voyages des autres camions ne seront pas modifiés. Cette approche fournit aux répartiteurs des plans alternatifs en quelques secondes.

De meilleures solutions pourraient être obtenues si le répartiteur était autorisé à apporter plus de modifications au plan initial. Dans le second article, nous considérons un contexte où un seul voyage à la fois est communiqué aux chauffeurs. Le répartiteur attend jusqu'à ce que le chauffeur termine son voyage avant de lui révéler le prochain voyage. Ce contexte est plus souple et offre plus de possibilités de recours en cas d'imprévus. En plus, le problème hebdomadaire peut être divisé en des problèmes quotidiens, puisque la demande est quotidienne et les usines sont ouvertes pendant des périodes limitées durant la journée. Nous utilisons un modèle de programmation mathématique basé sur un réseau espace-temps pour réagir aux perturbations. Bien que ces dernières puissent avoir des effets différents sur le plan de transport initial, une caractéristique clé du modèle proposé est qu'il reste valable pour traiter tous les imprévus, quelle que soit leur nature. En effet, l'impact de ces événements est capturé dans le réseau espace-temps et dans les paramètres d'entrée plutôt que dans le modèle lui-même. Le modèle est résolu pour la journée en cours chaque fois qu'un événement imprévu est révélé.

Dans le dernier article, la flotte de camions est hétérogène, comprenant des camions avec des chargeuses à bord. La configuration des routes de ces camions est différente de celle des camions réguliers, car ils ne doivent pas être synchronisés avec les chargeuses. Nous utilisons un modèle mathématique où les colonnes peuvent être facilement et naturellement interprétées comme des itinéraires de camions. Nous résolvons ce modèle en utilisant la génération de colonnes. Dans un premier temps, nous relaxons l'intégralité des variables de décision et nous considérons seulement un sous-ensemble des itinéraires réalisables. Les itinéraires avec un potentiel d'amélioration de la solution courante sont ajoutés au modèle de manière itérative. Un réseau espace-temps est utilisé à la fois pour représenter les impacts des événements imprévus et pour générer ces itinéraires. La solution obtenue est généralement fractionnaire et un algorithme de branch-and-price est utilisé pour trouver des solutions entières. Plusieurs scénarios de perturbation ont été développés pour tester l'approche proposée sur des études de cas

provenant de l'industrie forestière canadienne et les résultats numériques sont présentés pour les trois contextes.

Mots clés: temps réel, transport, foresterie, recherche opérationnelle, génération de colonnes.

ABSTRACT

When wood is transported from forest sites to mills, several unforeseen events may occur, events which perturb planned trips (e.g., because of weather conditions, forest fires, or the occurrence of new loads). When such events take place while the trip is under way, the truck involved must be rerouted to an alternative itinerary. Without relevant information on such alternative itineraries, the truck driver may choose a needlessly long one or, even worse, an itinerary that may itself be "closed" by an unforeseen event (the same event as for the original itinerary or another one). It is thus critical to provide drivers with real-time information, in particular, suggestions of alternative itineraries, when the planned one cannot be performed. Recourse strategies to deal with unforeseen events depend on the characteristics of the studied supply chain, such as the presence of auto-loaders and the management policy of forestry transportation companies. We present three papers dealing with three different application contexts, as well as models and solution methods adapted to each context.

In the first paper, we assume a context where truck drivers are provided a priori with the whole weekly plan. In this context, every effort must be made to minimize the changes in the initial plan. Although the fleet of trucks is homogeneous, there is a priority ranking of the truck drivers. The priority drivers are ensured the highest workloads. Minimizing the changes in their plans is also a priority. Since the consequences of unforeseen events on transportation are cancellations and/or delaying of some trips, the proposed approach deals first with single cancellations and single delayed trips and builds on these simple events to deal with more complex ones. In this approach, we try to reschedule the impacted trips within the same week in such a way that a loader is free at the truck arrival time both at the forest site and at the mill. In this way, none of the other trips will be impacted or changed. This approach provides the dispatchers with alternative plans in a few seconds.

Better solutions could be found if the dispatcher is allowed to make more changes to the original plan. In the second paper, we assume a context where only one trip at a time is communicated to the drivers. The dispatcher waits until the truck finishes its trip before revealing the next trip. This context is more flexible and provides more recourse possibilities. Also, the weekly problem can be divided into daily problems since the demand is daily and the mills are open only for limited periods in the day. We use a mathematical programming model based on a time-space network representation to react to disruptions. Although the latter can have different impacts on the initial transportation plan, one key characteristic of the proposed model is that it remains valid for dealing with all the unforeseen events, regardless of their nature. Indeed, the impacts of such events are reflected in the time-space network and in the input parameters rather than in the model itself. The model is solved for the current day each time an unforeseen event is revealed.

In the last paper, the fleet of trucks is heterogeneous, including trucks with onboard loaders. The route configuration of the latter is different than the regular truck routes, since they do not have to be synchronized with the loaders. We use a mathematical model where the columns can be easily and naturally interpreted as truck routes. We solve this model using column generation. As a first step, we relax the integrality of the decision variables and consider only a subset of feasible routes. The feasible routes with a potential to improve the solution are added iteratively to the model. A time-space network is used both to represent the impacts of unforeseen events and to generate these routes. The solution obtained is generally fractional and a heuristic branch-and-price algorithm is used to find integer solutions. Several disruption scenarios were developed to test the proposed approach on case studies from the Canadian forest industry and numerical results are presented for the three contexts.

Keywords: Real-time, transportation, forestry, operational research, column generation.

CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS

DSS	Decision Support System
DVRP	Dynamic Vehicle Routing Problem
LP	Linear Programming
LTSP	Log-Truck Scheduling Problem
MIP	Mixed Integer Programming
OR	Operational Research
VRP	Vehicle Routing Problem
VRPTW	Vehicle Routing Problem with Time Windows

“The way we see the problem is the problem.”

— Stephen Covey

REMERCIEMENTS

Il m'est particulièrement agréable d'exprimer ma reconnaissance envers toutes les personnes dont l'intervention au cours de cette thèse a permis son aboutissement.

Mes remerciements vont d'abord à Bernard Gendron et Michel Gendreau, pour m'avoir accueilli au sein de leur équipe et pour la confiance qu'ils m'ont accordée tout au long de ce travail. Merci pour votre encadrement, vos conseils judicieux et pour vos qualités humaines et scientifiques. Ma reconnaissance va également à Nizar El Hachemi pour ses encouragements et ses conseils qui m'ont facilité l'atteinte des objectifs de cette thèse.

Je remercie le réseau VCO pour m'avoir offert de travailler sur ce projet et pour leur soutien financier. Merci aussi à FPInnovations pour m'avoir accordé un stage, aidé à comprendre l'industrie forestière canadienne et fourni les données nécessaires pour compléter ce travail. Je remercie également le CIRRELT pour m'avoir offert des conditions de travail favorables à la réalisation de ce travail.

Je suis aussi reconnaissant à tout le corps professoral et administratif du DIRO et du CIRRELT pour tous les efforts déployés durant toutes ces années, ainsi qu'à tous les étudiants ayant contribué de près ou de loin au bon déroulement de cette thèse.

Je termine en exprimant ma profonde gratitude à toute ma famille pour leurs encouragements et leur patience. Sara, grâce à ton soutien, j'ai pu mener à terme ce travail.

CHAPTER 1

INTRODUCTION

Well-managed forests contribute to the conservation of the wildlife, the economic and the social well-being of millions of people around the world. According to the U.N. Food and Agriculture Organization (FAO) [46], wood is the main energy source for about a third of the world population. Forests provide also many direct and indirect employment opportunities to a large share of the world population. They are estimated at 13.2 million for the formal sector and 41 million for the informal sector. The formal sector involves wood production and processing and the pulp and paper industry. The informal sector consists of seasonal and part-time jobs performed by farmers and self-employed workers. Many countries include these jobs in agriculture, hunting, and fishing statistics rather than in forestry ones. The organization estimates also that forests cover over 4 billion hectares. However, the forest area has decreased from 31.6% of the global land area in 1990 to 30.6% in 2015. Therefore, a lot of effort was dedicated to planting trees, which are estimated at 7 percent of the total forest area in 2015. Moreover, different management plans were developed for about 52% of the forest area by 2010. This includes both plans intended to wildlife conservation and to wood production.

Canada has the third largest forest area in the world after Russia and Brazil. In 2014, the Canadian forest sector revenues attained 58 billion dollars according to the Forest Products Association of Canada (FPAC) [47]. This is partly due to the increasing interest in research and development, which was allocated over 200 million dollars in 2012. However, recent years have been difficult for the Canadian forest industry. It was highly affected by the global economic downturn, the collapse of the US housing market and the development of electronic media to the detriment of newsprint. Therefore, the FPAC 2020 vision was born with the aim of developing innovative forest products, reducing the environmental footprint and creating employment opportunities.

Transportation accounts for an important proportion of wood procurement costs and environmental impact. For example, trucking costs are estimated at 30% of the total cost in Quebec [40, 105]. The total cost includes stumpage, harvesting, transportation, silviculture and administration expenses. Also, log-trucks travel empty from mills to forest sites and therefore lose half of their transportation capacity. If the geographical positions of supply and demand points allow a flow of products going in opposite directions, one must take advantage of backhauling opportunities. In addition to empty driven trips, the unproductive activities include waiting. Loading and unloading operations are performed by loaders at forest sites and mills that are usually operated only for a specific period of the day. To avoid creating queues at the loaders, the synchronization of the trucks with the loaders must be taken into account.

In some Canadian provinces, the forest industry must also compete with oil and mining industries to retain the truck drivers. In addition to cost reduction and efficiency improvement, this leads to higher demand for accommodating drivers by proposing good quality routes. This means, for instance, that one must avoid routes that alternate trips to the south and to the north or those that alternate days with high and low workloads. Although it could be hard to meet all these objectives, some forest companies still rely on experienced dispatchers to derive transportation plans. This manual planning is laborious and stressful for the dispatcher. Therefore, more and more companies adopted decision support systems in the recent years. In addition to the graphical interfaces and the analysis capabilities that these systems provide to the managers, they generally use operational research models and methods to optimize the transportation. This resulted in lower costs and better quality routes.

Deriving schedules for trucks to transport different wood products between forest sites and wood mills at a minimal cost is known in the literature as the log-truck scheduling problem (LTSP). The main objective is to satisfy the demand for wood products at the wood mills, while not exceeding the supply at the forest sites. Some of the objectives presented above may or may not be taken into account depending on the application context. For example, the fleet of trucks could include trucks that are equipped with onboard loaders. These trucks do not need to be synchronized with loaders, but their usage is limited. Indeed, the onboard loaders decrease their transportation capacity and consequently their cost-efficiency. Also, some trucks may have specific constraints. For example, some roads cannot be used by some classes of trucks. Some drivers may have also preferences with regards to the visited sites. Every effort should be made to satisfy these needs especially for senior drivers.

The transportation plans obtained through the resolution of the LTSP are vulnerable to unforeseen events. The nature of the unforeseen events that arise in the forest industry is distinct from what can be found in the literature on similar problems in other sectors. We have drawn up a list of the most frequent unforeseen events. The list includes unforeseen events that are likely to appear at the forest sites, those involving trucks and road networks and those that occur at the mills. The manual reactions to these events are generally inefficient and time consuming. A more responsive and quicker decision system could improve the performance of the transportation fleet. More and more log-trucks are equipped with onboard computers and communication technologies that enable to capture the impacts of these events in real-time. In this thesis, we propose models and methods to generate, in a short computational time, new transportation plans in reaction to these impacts.

This thesis has been realized in collaboration with FPInnovations. FPInnovations is a not-for-profit research center that creates scientific and technical solutions to support the Canadian forest sector's competitiveness. Created in 2007 through the merger of Forintek, the Forest Engineering Research Institute of Canada (FERIC), the Pulp and Paper Research Institute of Canada (Paprican), and the Canadian Wood Fibre Centre, it brings together stakeholders with diversified backgrounds and acts as an innovation hub involving the industry, the government and universities. Its Value Maximization and Decision Support (VMDS) research group develops decision support systems to help industry members to generate more value through implementing value chain concepts. The latter include supply chain agility, which is considered, by many specialists, as a requirement for any modern supply chain. This is especially true given that the need and the request for timed deliveries are increasing in the forest industry. Within this framework, FPInnovations has raised the interest of forest companies for the development of methods and models to improve the responsiveness of the transportation systems to disruptions.

Two recent theses [37, 106] dealing with tactical and operational transportation planning in forestry designated the real-time rescheduling problem as one of the important research extensions of their work. These theses were realized in collaboration with FPInnovations and the work presented in [106] led to the development of *Truck Scheduler*. This tool optimizes transportation by minimizing both transportation (loaded and empty trips) costs and queuing times for loading and unloading the trucks. It produces transportation plans that satisfy a set of constraints such as demand satisfaction, increasing drivers' satisfaction and balancing the workload between different weeks. This tool takes several parameters (e.g., demand, supply, etc.) as an input. In practice, however, these parameters could be inaccurate since the decision maker might not have full knowledge of them and will then use estimates to derive the transportation plan. Using the same tool to reschedule transportation is not suited for a real-time application, first because it is time-consuming and second, it is often required to minimize the modifications to the initial transportation plan. The goal of this thesis is to develop fast algorithms that are able to handle the constraints necessary to deliver transportation plans in real-time.

To gain a better understanding of the forest operations, the author of this thesis was granted a four-month internship at FPInnovations. The first application context that was defined by FPInnovations during this internship involved the use of Truck Scheduler to construct a weekly transportation plan for a homogeneous fleet of trucks. The full weekly plan is communicated in advance to the drivers. In this context, the plan must be followed as much as possible after a disruption. Chapter 3 presents a heuristic method that produces new transportation plans with few deviations from the initial one. The method uses different recourses, such as adding a truck or extending mills opening hours to find feasible solutions in a short computational time. Chapter 4 presents a second application context, where only one trip at a time is communicated to the drivers. This allows to reoptimize the transportation plan after a disruption. The fleet is homogeneous and demand is stated on a daily basis. We present a mathematical programming model with flow constraints and a time-space network as input. This network changes whenever unforeseen events occur, as it will incorporate the consequences of the events. In this context, the model is solved for the current day every time unforeseen events occur and modifications to the time-space network are performed. Chapter 5 presents a third application context where the fleet of trucks is heterogeneous and the demand is stated on a weekly basis. We use a mathematical model where the columns correspond to truck routes. Similarly to Chapter 4, we use a time-space network to represent the disruption consequences and to generate feasible routes. The latter must, though, not deviate a lot from the routes in the initial transportation plan. We use column generation to solve the model as the number of feasible routes can be large. Real data provided by FPInnovations and generated unforeseen events through a discrete-event simulation procedure are used to assess the three approaches.

The author of this thesis developed the solution approaches proposed in Chapters 3, 4, and 5. He wrote the code, implemented and tested the models and the algorithms, and analysed output data. Michel Gendreau and Bernard Gendron supervised the thesis, helped with the design of the experiments, discussed the results and commented on the three manuscripts at all stages. The mathematical model proposed in Chapter 4 is adapted from [38], which is part of the Ph.D. thesis of Nizar El Hachemi. Dr. El Hachemi helped understanding the application context and the mathematical model. He also contributed to adaptating the model to a real-time setting. Moreover, he revised the second manuscript in Chapter 4 and participated to the discussions that led to the design of the third study in Chapter 5.

The remainder of this thesis is organized as follows. In Chapter 2, a literature review is presented. We start by reviewing the operational research contributions to the forest industry. This includes strategic, tactical and operational planning problems. We also give some examples of uncertainties that arise in forestry. A separate section is devoted to forest transportation, which is the subject of this thesis. We review also the dynamic vehicle routing problem, which gives a general framework for the real-time forest transportation problem. Chapters 3 through 5 present the three articles that have been produced over the course of these doctoral studies. Chapter 3 presents *Managing unforeseen events in forestry transportation*, which has been accepted for publication in *J-FOR, The Journal of Science and Technology for Forest Products and Processes*. Chapter 4 presents *Real-time management of transportation disruptions in forestry*, which has been submitted for publication to *Computers & Operations Research*. Chapter 5 presents *A heuristic branch-and-price algorithm to solve real-time transportation problems in forestry*, which has been submitted for publication to the *European Journal of Operational Research*. Finally, Chapter 6 summarizes the contributions of this thesis and discusses potential future research directions.

CHAPTER 2

LITERATURE REVIEW

Operational research (OR) has played an important role in improving the forest resources management practices since the 1960s [121]. Moreover, the last 50 years have witnessed an increase in practical requirements that must be taken into account during the optimization process. The literature is rich in examples of challenges faced by OR practitioners in the forest industry. One example of these is real-time transportation, which is the subject of this thesis. This problem is closely related to dynamic vehicle routing problems (DVRP) that can be found in other industrial sectors. In this chapter, we review the literature on OR in forestry and, more in detail, in log transportation and DVRP. We conclude this chapter by stating the contributions of this thesis.

2.1 Operational research in forestry

Optimization techniques are used in a wide range of applications within the forest industry. This includes long-, medium- and short-term planning problems. The long-term strategic planning may involve a time horizon of more than one hundred years in the case of forest management, while the short-term planning for trucks dispatching involves only a few seconds. In a more formal way, [107] defines the time horizon for strategic, tactical, operational and real-time planning problems at more than five years, between six months and five years, between one day and six months and less than one day periods, respectively. Note, however, that this definition can vary depending on the country and the context of application. This section reviews OR in forestry, but we devote a separate section to OR in forest transportation, which is the subject of this thesis. The present review does not cover all the planning problems that arise in the forest industry. We refer the interested reader to the two recent books [31, 122], where various subjects, models and methods can be found.

2.1.1 Models and solution methods

An optimization model formulates a real problem using a set of decision variables, constraints, and an objective function. It ensures the best use of available resources, while minimizing or maximizing the objective function. The objective may also include different goals that must be balanced. The mathematical model is solved using an appropriate solution method. Requirements on the quality and the time of the solution are usually part of the description of the real problem. For the real-time problem that we study in this thesis, there is a need to solve the model in a small computational time. However, tactical and strategic problems can take up to several hours to find a solution. [107] provides a list of typical forest planning problems along with the time available to solve them.

The main methods used to solve planning problems in forestry are derived from mixed integer programming (MIP) and linear programming (LP) models. LP models are often used in strategic forest management, while MIP models are used for operational and real-time log-truck scheduling, annual harvesting, and equipment scheduling. MIP models may be hard to solve using a commercial solver, but the properties of the problem can be used to improve the solution approach. Branch-and-bound and cutting plane algorithms, or their combination (branch-and-cut) are often used to solve MIP models. The number of variables in these models is generally too large to be handled explicitly. Column generation is then used to add iteratively only the variables that have the potential to improve the objective function. The problem is separated into a master problem and a subproblem. The master problem has the same structure as the model to be solved, but includes only a subset of the variables. The subproblem is used to identify the new variable to be included in the master problem. This method has received a lot of attention in recent years and has been used for solving the log-truck scheduling problem (LTSP) [97, 104, 105].

Dynamic programming involves breaking the problem down into a set of stages. At each stage, there are a number of states corresponding to the possible conditions at the start of the stage. Given a state, a decision must be made, which yields a new state at the following stage. The goal of the method is to find the optimal policy, which consists of the set of optimal decisions to take at each stage and each state. Dynamic programming is often used to solve the subproblem of the LTSP in a column generation method. The columns in the master problem correspond to feasible routes for each log-truck. Using a time-space network, a route corresponds to a path in the network. Dynamic programming is used to find the shortest path in such networks [104, 105]. It is also the main method used to solve the bucking problem, which consists in cutting trees into different log types [107].

Well-designed heuristics are able to quickly find high quality solutions. They use logic rules to guide the search for these solutions, but there is no proof of their quality. Therefore, a solution is considered of high quality if it is at least as good as a manual solution derived by an experienced planner [107]. Heuristics are used in many planning problems in the forest industry, especially for operational and real-time problems where fast solutions are needed.

Simulation is used to evaluate the solutions found by solving the mathematical models. Given the stochastic nature of many parameters included in forest planning problems and the unforeseen events that can be revealed during the execution of the plans, simulation can be used to identify the bottlenecks associated with the execution of the solutions. The model can be modified and solved again after the simulation. This technique was recently used to assess the performance of a transportation plan considering uncertainty in trucks arrival time at a mill [78]. In stochastic programming, simulation can be used to generate a set of scenarios that are used as input to the optimization model. This method assumes that uncertainties can be described by probability distribution functions that allow to generate these scenarios. The goal is to find a feasible solution that maximizes or minimizes the expectation of the objective function given a set of scenarios. This is the case for example in forest management, where simulation is used to model the evolution of forests over time and to choose the potential silvicultural interventions to reach the desired forest state. Then, a mathematical model is solved to decide which silvicultural intervention must be applied to each sector [126].

The evaluation of the solutions must also be done by the planner. Since the number of decision variables can be large, a graphical interface must be used to format the solution in an easily readable way. Many decision support systems were developed in the forest industry [9, 69, 115, 124]. More details on decision support systems and OR methods used in the forest industry are given in the next subsections, which cover all planning levels (strategic, tactical and operational), as well as a subsection presenting how uncertainties are handled.

2.1.2 Strategic planning

Several surveys traced the development, as well as the challenges faced by OR in forestry. [10] reviews the OR contributions from the late 1960s to the mid 1980s, while [109] is a recent survey presenting also 33 current open problems in forestry. Some of the applications are very specific, whereas others are more general. The latter involve several decision problems. An example of a general problem is strategic forest management planning. This problem aims at maximizing the net worth of forest resources, while preserving the surrounding social, ecologic and economic systems. This involves, among other things, the wildlife, habitat and scenic beauty conservation [11].

To preserve the wildlife habitat, some contiguous areas of a certain age must be preserved [95] and several constraints are imposed on the maximal area to be harvested. Also, if a block is harvested, the adjacent ones must be preserved for a certain number of years. On the other hand, minimal area limitations are also necessary to ensure cost-efficiency. From a modelling perspective, this gives rise to a large set of constraints, and binary variables are generally used to decide which blocks to harvest [89]. The main OR contributions try to minimize the number of constraints, while ensuring that the linear relaxation solutions remains close to the initial integer program solutions [54, 55, 83, 84, 90, 118]. Besides the economic benefits of harvesting, it is also used to reduce tree diseases and fire risks. Fire is also used as a natural mean to preserve the forest health. In [79, 80], the authors review OR approaches to forest fire management problems. This involves also air tankers location and crew scheduling problems [56].

Strategic planning is considered both at governmental and industrial levels. The US Forest Service developed Timber RAM (Resource Allocation Method) [92], which is a strategic LP model used to plan long-term harvesting and reforestation. Later, it developed FORPLAN [69] to take into account the new socio-economic concerns. In Sweden, Heureka [124] contains an optimization module and a simulator to evaluate the management plans. Similarly, FOLPI [51] was developed by the New Zealand forestry ministry. These countries have different regulations and approaches to solve the strategic problems. The different requirements give rise to different constraints from a modelling perspective [60], which impacts also the quality of the solutions [59].

Strategic decisions include also road construction, the opening and location of mills, transportation equipment, information technology, etc. Synchronizing all these activities is known as supply chain management. For example, road building/upgrading investment decisions must take into account some of the supply chain components, such as the customer demands, the cost variations and the inventory management policies [96]. The goal is to improve the efficiency of the flow of wood from planting and harvesting (sourcing) to the distribution of the final product to end consumers (sales) [22, 24].

Issues that arise in strategic planning are generally hard to measure, often conflicting and hard to manage simultaneously. Finding an optimal solution is generally hard to achieve and the challenge is rather to find feasible solutions [82]. However, such solutions must be optimized, since they have long-term impacts on the quality and quantity of the available wood. Decomposition techniques and hierarchical planning approaches are generally used for this purpose [119]. The first step is generally to make decisions on the volume and then on the areas to harvest or plant. Strategic decisions impact the tactical ones, which in turn impose restrictions on the operational planning. Simulation models allow evaluating the impact of these strategic decisions [15].

An important problem in the hierarchical planning is the interaction of the decisions at different planning levels. Generally, the models are solved separately and they are linked using heuristic methods. This is done through aggregation and disaggregation of decisions between the planning levels. [81] reviews the principles of hierarchical planning and presents some practical experiences and feedbacks from both the industry and the government with regards to the hierarchical forestry planning systems that they use.

2.1.3 Tactical and operational planning

The definition of the tactical and operational planning levels is very dependent on the context of application. They are also defined by the contracts between the different parties. Therefore, they are reviewed in a single subsection.

At the medium-term planning level, tactical problems typically integrate harvesting with transportation [109]. However, tactical harvesting and road building decisions were originally taken separately. [66] analyzed the advantages of combining the two problems into a single model. Different formulations and resolution methods can be found in the literature. These problems are generally modelled as mixed integer programs (MIP). [4] uses Lagrangian relaxation, while [45] uses constraint aggregation to solve such models. Also, many decision support systems (DSS) such as FORPLAN [69], SNAP [112], PLANS [115] and RoadOpt [68] include models to make choices of roads to build or upgrade. Locating harvesting machinery is another problem that is linked with harvesting and transportation. PLANEX [44] uses a heuristic method based on a geographic information system (GIS) and economic parameters to decide of machinery installations, road building and wood transportation in Chile. [18] included also crew scheduling with harvesting decisions.

At the short-term planning level, operational problems involve harvesting, bucking and transportation [42]. The wood logs are generally defined by their species, quality, shape, diameter and length. Short-term harvesting problems decide which areas to harvest each week. They decide also which patterns to use to define the demanded logs. OPTICORT [43] uses column generation to define the areas, the volume, the products to harvest, the machinery to use and the assignment of these products to the customers. [70] use a tabu search to define the bucking patterns, given certain market prices. For the same problem, [20] uses dynamic programming and simulation, while [41] uses a Dantzig-Wolfe decomposition procedure and solve the subproblems using dynamic programming.

Horizontal collaboration also gained a lot of interest in the last few years. OR is more and more used to evaluate this kind of collaborations. In [12], the authors propose an approach to find an agreement between multiple firms on a common harvesting plan. In [49], the cost allocation problem in a collaboration at the log transportation level is addressed, while [73] proposes an evaluation methodology for a collaboration between a paper mill and a wholesaler.

2.1.4 Uncertainties in forestry

Several sources of uncertainty exist in the forest industry [109]. The wood markets include uncertainties about the prices and demand volumes. The forest areas involve uncertainties about the trees growth and quality, the global and per species volumes estimates and diseases and fire risks. The wood production includes inaccuracies about the harvesting and transportation plans. Other uncertainties include technology developments and regulations changes. These inaccuracies are generally handled through including extra travel times for transportation estimates, extra inventories for demand levels, extra dollars for costs estimates, etc. In the 80s, chance constrained programming was mainly used to solve strategic forest management problems under timber growth uncertainties [63, 123]. Other stochastic models have been lately introduced to cover other types of uncertainties in the forest industry.

In [1], the authors consider a tactical harvesting and road construction problem with market uncertainty. The stochastic model that they developed tries to find the best solution that is feasible under all the generated timber prices variations scenarios. In addition to the big number of scenarios that could be generated, the nonanticipativity constraints make the problem resolution harder. These constraints state that if two different scenarios are identical up to a certain time interval, the decision variables values must also be identical up to that interval. The authors use a branching scheme where these constraints are implicitly satisfied.

The same problem was considered in [117], but under uncertainty in timber growth. The authors use the progressive hedging method where the problem is decomposed and the model is solved for each scenario separately in an iterative fashion. This allows to remove the nonanticpativity constraints. The proposed method penalizes the deviations from the average value in all scenarios of the variables that must satisfy these constraints. The iterative fashion allows the convergence to a solution that respects the nonanticiptavity constraints.

In [23], the authors present a Scandinavian case study where there are uncertainties about the demand. The original approach was to keep a safety stock to face the demand fluctuations. The authors propose, instead, a robust optimization approach to eliminate these stocks. The approach decomposes the problem into two separate problems. In the first problem, they find a feasible solution. The second gives the worst-case scenario given this solution. This is used as a valid inequality in the first model. The process is repeated until the first problem produces a solution satisfying the worst-case scenario.

In [113], the authors study the problem of selecting natural areas to be preserved for the protection of species over time. In practice, these decisions take place sequentially as funds become available. Site availability is dynamic, since some parcels can be converted from open space to developed uses if protection is delayed. To solve the sequential site selection problem under uncertainties in budget and site availability, the authors use an integer programming model that maximizes the expected number of species represented by the set of selected sites. The uncertainty is characterized by a set of scenarios. Since the number of the scenarios increases exponentially with the number of sites, the authors choose randomly a subset of scenarios to be included in the model.

There is still a lot of work to be done in formulating and solving stochastic problems in the forest industry. In [109], at least 10 of the 33 open challenges that were identified for OR in forestry involve uncertainties. Large-scale problems in forest management and value chain management are more challenging, since they include many stochastic processes. It is hard to quantify and to represent the underlying multiple sources of uncertainties. It is also hard to construct mathematical models that explicitly consider uncertainty, especially for multi-objective and multi-stage planning problems.

There are significant computational challenges to solve these complex problems. The number of scenarios, in a scenario-based model, can be large and the problem cannot be solved directly. Ad-hoc methods can be used to reduce the size of the pool of scenarios. The problem can also be decomposed into several subproblems that are easy to solve. Stochastic optimization assumes that the uncertainty has a probabilistic description, which can be hard to define. Depending on the complexity of this description, the model may become hard to solve within a reasonable computational time. In robust optimization, the uncertainty is only known to belong to some uncertainty set and there is no requirement to have probability distributions. The goal is to find the optimal solution that is feasible in the worst-case scenario. To avoid conservative and costly solutions, care must be taken in the construction of the uncertainty set, which is challenging for large scale forest planning problems.

2.2 Dynamic vehicle routing problems

The real-time transportation problem that we study in this thesis is closely related to the dynamic vehicle routing problem that has received a lot of interest during the last few years. Indeed, the search for efficiency at all the business levels is a key element to maintain the competitiveness of a modern company. This includes the efficient distribution of goods and services. Also, the customer needs are evolving and the capability to fulfill these requirements impacts the competitiveness of the companies. The late technological advances allow to track the fleet of vehicles in real-time, which gives opportunities to react efficiently to changes in demand and travel times. Numerous decision systems were developed for that purpose. They range from scheduling maintenance workforce [74] for dynamically revealed service requests to city logistics systems [127] for changing travel and service times. In this section, we start by presenting the static vehicle routing problem before discussing the dynamic version.

2.2.1 Vehicle routing problem

The vehicle routing problem (VRP) [114] aims at designing a set of minimal cost routes for a homogeneous fleet of vehicles in order to satisfy the demand of a given set of customers. Some variations of this problem may have a different objective function such as minimizing the number of vehicles used [75]. In the capacitated VRP (CVRP), the vehicles have finite capacity, while the heterogeneous VRP (HVRP) considers vehicles with different capacities and configurations. Another variation imposes time windows within which a customer can be served. This is known as the VRP with time windows (VRPTW) [67]. The VRPTW can also be extended to include both pickup and delivery operations. In the pickup and delivery problem with time windows (PDPTW) [36], the customer demands consist of a pickup at an origin site followed by a delivery at a destination site within given time windows. The dial-a-ride and dial-a-flight problems are examples of PDPTW applications. Many distinctions can be found based on whether there is a single or multiple vehicles, if people or commodities are transported, etc.

The customers are generally visited only one time in the problems presented above. However, in some applications such as forestry, several visits to the pickup and delivery sites are necessary to satisfy the demand. Also, the possibility of serving a customer by more than one vehicle is used to potentially reduce the total costs. This can be achieved, for example, if a certain number of vehicles are not fully loaded. The spare space can be used to deliver a given demand in a split way instead of using extra vehicles. In the literature, the researchers studied split pickup problems [72], split delivery problems [58], and split pickup and delivery problems [87].

The amount of published material on VRP has exponentially increased since its first introduction. This is due to the diversity of its variants and of the optimization approaches that have been proposed to tackle them. This includes both exact and approximate approaches. The next subsection presents the dynamic version of the VRP along with some solution methods. We refer the interested reader to the books [53, 114] for a complete review of the VRP variants, solution methods, challenges and applications.

2.2.2 Dynamic vehicle routing problem

Other variations of the VRP are based on the nature and quality of the data it considers. If all the data is known beforehand, the VRP can be solved statically. But if the data is evolving over time, the VRP must be solved dynamically. This is known as dynamic VRP (DVRP). Also, the data may be assumed to be fixed and accurate. The problem is then considered as deterministic. On the other hand, data can be expressed as random variables with certain probability distributions. This is known as stochastic VRP (SVRP). Table (2.I) [100] summarizes these problems.

	Information quality	
Information evolution	Deterministic input	Stochastic input
Input known beforehand	Static and deterministic	Static and stochastic
Input changes over time	Dynamic and deterministic	Dynamic and stochastic

Table 2.I – Taxonomy of vehicle routing problems by information evolution and quality [100]

The problem that we study in this thesis falls into the category of dynamic and deterministic problems. The latter are also called real-time or online VRPs. In these problems, a part of the input is revealed dynamically and the routes are modified accordingly. In the problem we study, the dynamic input corresponds to the unforeseen events. The transportation plan is re-optimized every time a disruption happens. High quality solutions must be found in a short computational time. Also, one must reduce the deviations from the initial plan. The dynamic and deterministic log-truck scheduling problem is able to find a trade-off between solution quality, computational time and perturbations for the original plan. The main drawback of this approach is that it does not anticipate the disruptions. A stochastic model would be able to produce solutions that are less vulnerable to disruptions. Unfortunately, there exists no study on the probability distributions of such events in the forest industry. Conducting such a study is challenging given that it is hard to get historical data. This is, however, an avenue for future research since many forest companies are investing in log-trucks with onboard computers, positioning systems, and communication technologies that can be used to collect accurate data on disruptions.

The uncertainty sources that are mostly addressed in the literature on DVRP are new customer requests [100]. Some applications involve stochastic travel times [127] and less importantly vehicle breakdowns [76]. Moreover, there are problems where there is no need to define complete routes with different customers visits. This is the case for emergency vehicles [62]. Taxi bookings can also arrive dynamically, but some requests may be known in advance in the case of an early booking [21]. These applications consider generally one source of uncertainty at a time. In the forest industry, in contrast, the events that impact the transportation can be of different natures: trucks, loader breakdowns, road closures, delays, etc. In this thesis, we consider all these events simultaneously.

The approaches used to solve DVRP include many techniques both to generate and to evaluate new routes. For example, in express courier applications, it is possible to reject some new requests because of infeasibility or cost inefficiency [52]. It is also possible to divert a vehicle from its destination to serve a nearby customer [64]. This is generally allowed by the supporting urban transportation network, which enables reaching the same destination by taking different itineraries. The forest roads do not generally allow this kind of recourses.

Solution approaches for dynamic and deterministic VRP are generally based on those that were developed for the static VRP. They rely on the extensive research that was done in this field and their adaptation to dynamic contexts proved to be efficient. [101] developed a dynamic programming approach to solve a single vehicle dial-a-ride problem where they consider both service time minimization and customer satisfaction maximization. The problem aims at finding the best route each time a request is revealed. The use of dynamic programming, however, necessitates an important computational effort, as the size of the instances increases significantly. This is known as the curse of dimensionality. To cope with the latter, Approximate Dynamic Programming approximates the value function and avoids the evaluation of all possible states.

In [26], the authors use a column generation approach for a DVRP with time windows where the demands arrive dynamically. The approach uses the existing columns to generate the new ones after each request. For a DVRP with pickup and delivery where requests for transportation tasks between an origin and a destination node arrive dynamically, [125] developed five policies based on either offline reoptimizations or heuristic methods.

The solution approaches include also metaheuristics. [13] addresses a patient transportation problem where requests for transportation between locations in a hospital campus arrive dynamically. The authors use an insertion algorithm to find a feasible solution that is improved by a tabu search heuristic. [52] addresses a dynamic courier service problem using also a tabu search heuristic that is implemented on a parallel platform. Genetic algorithms are used in [116] and [61] where both requests and travel times variations are considered. [50] and [88] developed approaches based on ant colony systems.

2.3 Transportation in forestry

The problem that we study in this thesis is a dynamic and deterministic VRP with time windows. Time windows are mainly the opening times of the mills and the operating hours of the loaders at the forest sites. However, these time windows are generally not as tight as in other industries. This is also a pickup and delivery problem with split pickups and deliveries, since several visits are necessary to both the mills and the forest sites in order to satisfy the demand. Additional constraints include the synchronization with loaders for both loading and unloading operations. We consider also both the homogeneous and the heterogeneous cases. This section reviews the contributions of the OR community to forest transportation in general including the real-time problem.

2.3.1 Strategic and tactical transportation

Strategic transportation planning is mainly concerned with road building and investments in heavy transportation modes such as trains [9]. Tactical planning involves upgrading roads and managing the fleet utilization [107]. If a heavy transportation mode is used, the tactical decisions may involve the upgrading of these modes, such as adding wagons to trains or changing their schedules. Strategic and tactical transportation planning problem are generally integrated with harvesting problems. At the strategic level, the harvesting and road building decisions are mutually influential. At the tactical level, harvesting decisions involve the assignment of the harvested wood to the customers. These decisions are critical for the transportation costs since they influence directly the backhauling opportunities [25].

FlowOpt [48] is a decision system used for strategic and tactical transportation problems in Sweden. At the strategic level, the aim is to decide whether to use trains, trucks or their combination. The decisions concern also the train capacity and the terminal location. At the tactical level, decisions concern the areas to harvest, the allocation of wood to customers and backhauling opportunities. The system considers also collaboration between companies to exchange wood in order to enable backhauling. This approach generated cost savings of up to 12.8% in some case studies. The transportation planing problem is formulated as a linear programming problem. The model is solved using column generation. The subproblem is a constrained shortest path problem, which is used to generate iteratively the best possible backhauling routes.

FPInterface is a simulation tool developed by FPInnovations in Canada. It contains a transportation planning method named MaxTour [9]. The latter generates backhauling tours given the volumes of supply, their destinations and different configurations of trucks. The method is based on the Clarke and Wright savings algorithm [27].

2.3.2 Operational transportation

Operational transportation decisions deal generally with trucks scheduling. Depending on the quality of the roads, transportation can be done in one or two phases. If the harvested areas are difficult to access, the first phase consists of forwarding the logs from the harvested areas to sites that are close to roads. The second phase is the transportation to the mills. This phase involves the most important saving opportunities and has then received more attention. This problem is known as the log-truck scheduling problem (LTSP).

Many approaches were developed to solve the LTSP. [104] studies a daily LTSP. The problem is formulated using an integer programming model with binary variables representing truck routes. These variables represent the columns of the model. The integrality constraints are relaxed and the model is solved using column generation where the routes are generated dynamically. This is done by solving a shortest path problem in a time-space network. An integer solution is found using a branch-and-bound algorithm. A similar approach is used in [98]. However, the authors generate a priori a set of routes following certain rules. Also, the authors use a heuristic branch-and-price method to find an integer solution.

In [38], the authors use a two-phase approach to solve the weekly LTSP. The first phase involves an integer linear program that determines the destinations of full truckloads. The demand is then expressed as a set of trips to operate throughout the week and the weekly LTSP is decomposed into seven daily LTSPs. The daily LTSP is formulated as an integer programming model using arcs from a time-space network. Besides minimizing transportation costs, the authors consider also minimizing the loader working time at forest sites. To derive integer solutions, a heuristic branching strategy favouring solutions with lower loader working times is used. This two-phase approach is also used in [40] . However, the daily LTSP is solved using two approaches. The first one uses a constraint-based local search method where constraints are used to control the local search. The second one integrates the first approach with a constraint programming model.

In [57], the authors assume a context where all the transportation requests are a priori known. The loaded travel times are then fixed and the goal is to minimize empty driven trips. The authors adapted the unified tabu search algorithm, initially proposed for a general VRP [28], to solve a daily LTSP. The same context is studied in [39], but the authors combine a constraint programming model with an integer programming model to solve the problem.

In [85], a simulated annealing method is used to solve a daily LTSP. Simulation was used to evaluate the quality of the first found solution with regards to different metrics such as empty travel and waiting time. Then, heuristic rules were used to modify this solution based on its performance. This produces new routes, which are considered during the simulated annealing process. A new solution is then generated, simulated and evaluated again. If this solution has a better performance than the previous one, it is kept with a certain probability. The process is repeated until there is no improvement.

Many decision support systems were also developed. ASICAM is a DSS that was developed in Chile [120]. It is used on cases with a fleet reaching 220 trucks. The system uses a simulation process based on heuristic rules to design daily transportation plans. The reported results show an improvement of the quality of the schedules and cost savings between 15% and 35%.

EPO is a system that integrates strategic, tactical and operational planning problems [77]. The system was developed and used in Finland on cases with a fleet of 250 trucks but there were different dispatching centers, each managing up to 20 trucks. The approach starts first by allocating the wood from forest sites to mills. The second phase aims at designing weekly transportation plans. This is done using a combination of an optimization method and some heuristics in addition to some prior knowledge. The third phase involves manual post-processing based on the dispatcher's experience.

RuttOpt was developed in Sweden [5]. The DSS produces transportation plans for up to five days for a fleet size ranging from 10 to 110 trucks. The proposed approach is based on a two-phase algorithm. The first one solves a linear programming model to generate a feasible flow between the forest sites and the mills for each truck. The second phase solves the routing problem using a tabu search method. The reported cost savings range from 5% to 30 %.

2.3.3 Real-time transportation

Real-time transportation deals with truck dispatching where typically only one load at a time is assigned to idle trucks. In this thesis, however, we study a real-time problem where we generate daily or weekly plans every time the original weekly transportation plan is disrupted. Note that this problem involves the dispatching problem as well, but tries to take advantage of the available information about future trips.

The LTSP literature addresses mostly the static LTSP. In contrast, not much research has been done for solving the real-time LTSP. A New Zealand forestry company has developed a decision support system called CADIS (for Computer Aided Dispatch) for dispatching the log-trucks in real-time in a context where the supply and the demand for forest products are continuously changing during the day [110]. Given more accurate data about the supply and the demand, the dispatcher communicates one trip at a time to the drivers as they complete their current trip. This system was tested with a fleet of 120 trucks. However, it was used for only a short period because of the dissolution of the company due to financial issues. Moreover, because of confidentiality agreements, only the general scheme of how this system works was provided in [110]. However, [108] gives a feedback about the implementation challenges and the lessons learned from this application. The author reports that CADIS was particularly useful when the overall supply was low.

In [110], the problem is modelled using binary variables to decide which routes to select. These route variables represent the columns of the mathematical model. The constraints are demand satisfaction and not exceeding both supply and the number of available trucks. The solution approach starts by looking for the best single trip for the trucks. The trucks are ordered by the expected completion time of their current trip. These single trips are found using a greedy algorithm. The dispatcher can use these trips for urgent requests. Then, a second greedy heuristic is used to derive full routes. Again, the first trips of these routes are made available to the dispatchers. The routes are also used in the linear relaxation of the mathematical model, which is solved using column generation. To find integer solutions, a branch-and-bound algorithm with heuristic branching rules is used. The first trips are again communicated to the dispatcher. This approach provides the dispatchers with feasible trips quickly, which are improved when enough time is available between successive trip requests from the drivers.

Few works considered a priori stochastic information about the input of a dispatching problem. In [86], the authors use a simulation approach to solve a dispatching problem for a case study involving 30 trucks. Although the travel speed is supposed equal for all trucks, some assumptions were made about the distribution of truck delays both during transportation and waiting at mills. Stochastic information was also available about loading times, the quantities and the nature of the product being loaded. For example, the probability of a product to be loaded is equal to the percentage of that product available at the forest site. The dispatcher has to assign a new pickup operation to the trucks each time they quit the wood mills. Three methods were developed for this purpose. The first one considers a random assignment where the probability of a truck being sent to a forest site is the same among all the sites. The second one is a static assignment where each truck is only allowed to visit an a priori fixed forest site. The authors refer to the third method as the "informed" assignment. In this method, an estimate of the time when the supply will be exhausted at a forest site is calculated. The truck is then sent to the forest site such that the difference between its exhaustion time and the truck arrival time is minimized. However, if a loader at a forest site remains idle for a long time, the truck is sent there. These methods were evaluated using simulation. Although the authors expected that the "informed" method would outperform the two other methods, the fixed assignment method was able to deliver slightly more wood. This is explained by the fact that the fixed method assigned wood flows between mills and forest sites that are close. The "informed" method, instead, does not contain enough information about the transportation distances and costs.

Some commercial DSSs include also a dispatching module. FLO is a system developed in the US and used in applications involving a fleet of 50 trucks on average [9]. The system comes with different communication and reporting technologies. The loader operators at the forest site have to enter the quantities of available wood and to define the product to be loaded and its destination. They repeat this process until the end of their working day. An optimization module matches this data with the availability of the trucks and their positions and selects a truck to pickup the available load. The truck position is tracked using global information and positioning systems (GIS and GPS). Once the truck reaches the mill, it is unloaded. If the truck is still within operating hours, it can be considered for a second pickup by the optimization module. Otherwise, it is sent to its home base. The optimization module is, however, undocumented and it is then hard to assess its performance.

Forestruck is an information system that was developed in Chile. It considers the whole supply chain from a management perspective starting at the sourcing (wood production) and ending at the sales (wood delivery) [9]. The system contains a module for truck dispatching, but there is no indication on whether an optimization model is used or not, nor on its performance.

2.4 Contributions of this thesis

We demonstrate in this thesis that real-time re-planning in reaction to disruptions can be done effectively and efficiently for realistic instances of log-truck scheduling problems. This problem is different from the dispatching problem where only one trip at a time must be planned, without the need to generate a complete daily plan. In this thesis, we generate new daily and weekly plans after each disruption. We present a list of frequent disruptions in forest transportation. Unlike in the urban context, where the literature deals generally with one problem at a time (often demand changes), we propose an approach that is valid for all types of disruptions. We relate also the real-time LTSP to other known routing problems and show the differences such as volumes to be transported, quality of the forest roads networks, and synchronization constraints. Transportation management practices differ from one forest company to another, which defines different application contexts. We study three contexts of application driving different solution methods. We start by defining core events, which can be used to describe other disruptions. In the first application, we use a heuristic to construct a new transportation plan with very limited deviations from the initial one. Given the conservative culture of the Canadian forest industry, we suggest different solution options rather than imposing one optimal solution. In the second and third applications, a time-space network is used to capture the impacts of the disruptions and represent the trips and activities that are still feasible in the physical network. Then, an arc-based mathematical model is used in the second application, where the transportation plan is completely re-optimized for the remaining of the current day. In the third application, we use a route-based model where a column in the model corresponds to a feasible route. The goal is to reoptimize a weekly transportation plan with few deviations from planned truck trips. The model is solved using column generation and a branch-and-price algorithm with a heuristic branching rule. During the column generation process, we use a *compatibility degree* to capture the trade-off between minimizing cost and deviations from planned truck trips. We demonstrate that MIP technology can be exploited in this context since effective solutions can be found in a few minutes. We evaluate the three solution methods using real instances from the Canadian forest industry and a simulation procedure that we developed to generate the disruptions.

CHAPTER 3

ARTICLE 1: MANAGING UNFORESEEN EVENTS IN FORESTRY TRANSPORTATION

Chapter notes

This chapter was accepted for publication in *J-FOR, The Journal of Science and Technology for Forest Products and Processes* on February 26, 2016. Preliminary work was presented at the following conferences:

- VCO Workshop, Montreal, Canada, February 27, 2015
- 56th CORS Annual Conference, Ottawa, Canada, June 26-28, 2014
- VCO Summer School 2014, Halifax, Canada, June 11-13, 2014
- Optimization Days 2014, Montreal, Canada, May 5-7, 2014
- 55th CORS Annual Conference, Vancouver, Canada, May 27-29, 2013
- VCO Summer School 2013, Montreal, May 17, 2013
- 1st FIBRE Conference, Cornwall, Canada, May 13 -16, 2013
- Optimization Days 2013, Montreal, Canada, May 6-8, 2013

MANAGING UNFORESEEN EVENTS IN FORESTRY TRANSPORTATION

Amine Amrouss [1,2,*], Michel Gendreau [1,3], Bernard Gendron [1,2]

[1] *Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation (CIRRELT), Canada*

[2] *Department of Computer Science and Operations Research, Université de Montréal, P.O. Box 6128, Station Centre-Ville, Montreal, Canada H3C 3J7*

[3] *Department of Mathematics and Industrial Engineering, Polytechnique Montréal, C.P. 6079, succursale Centre-ville, Montreal, Canada H3C 3A7*

When wood is transported from forest areas to plants, several unforeseen events may occur and disrupt planned trips. In this paper, we present heuristic methods for solving real-time transportation problems in which the initial plan must be re-optimized to deal with unforeseen events. Here, the nature of the events that must be dealt with differs radically from what can be found in the routing literature, since one must deal not only with changes in the demand (e.g., the arrival of a new request), but also with changes in the topology of the transportation network (e.g., road closures). Since the consequences of such events on transportation are cancellations and/or delaying of some trips, the proposed approach deals first with single cancellations and single delayed trips and builds on these simple events to deal with more complex ones. The heuristic methods we propose provide the dispatchers with alternative plans in a few seconds.

Keywords: Forestry, real-time, transportation.

3.1 Introduction

According to Natural Resources Canada [91], Canada is, by any measure, a forest nation. With 348 million hectares of forests, it is the second largest exporter of raw forest products after the USA. In 2013, the contribution of the forest industry to Canada's GDP stood at $19.8 billion. However, the recent years have been difficult for the forest industry. It was heavily impacted by the global economic downturn, the collapse of the US housing market and the development of electronic media to the detriment of newspapers. Meanwhile, Canadian companies have managed to reduce their dependence on the US market by opening up to Asian markets.

In order to stay competitive, Canadian companies have to adopt innovation at the center of their competition strategies. In this context, the use of optimization models and methods arises naturally as a way to reduce costs, respond efficiently to customer needs and reduce negative environmental impacts of forest operations. In particular, Canadian firms are using optimization models and methods to improve the efficiency of different components of the forest supply chain, including the transportation of forest products [30, 107]. In this paper, we focus particularly on the optimization of the real-time timber transportation from the forest sites to the wood mills. During the transportation operations, several unforeseen events may occur and disrupt the original plan. Recent technological advances allow capturing such events in real-time, providing opportunities to improve the performance of the existing dispatching systems that generally rely on the experience of the dispatcher. We propose methods capable of processing information in real-time and reacting to various unforeseen events that may occur during the execution of the transportation plan.

The remainder of this paper is organized as follows. We first describe the challenges related to the planning and control of forest transportation in real-time. The next section reviews the literature related to real-time transportation in forestry. We then describe the proposed methods to react to real-time disruptions. The results of our approach are presented in the following section. Finally, the last section concludes this work.

3.2 Problem description

In this paper, we are interested in the real-time transportation operations of timber from the forest sites to the wood mills, i.e., the first half of the forest supply chain. The other half of the supply chain, linking the wood mills to the consumers of finished wood products, is beyond the scope of this work.

Every wood mill and every forest site is equipped with a loader. The mills are usually open only for a specific period of the day. There are different forest products to transport. A product is defined by its species, its quality, its length and diameter. The truck fleet is homogeneous and each truck is associated with a base, which is usually a wood mill. It must begin and end its day at its base. The mills demand for forest products and their availability are generally given for a period of one week but may be divided by days. The volumes can be expressed in terms of full truckloads since the truck fleet is homogeneous.

When transporting wood from the forest sites to the mills according to a predefined plan, many unexpected events may disrupt the plan. When such events are known only during a trip, the truck that accomplishes this trip should be provided with an alternative plan. In the absence of relevant information on the state of the supply chain, the truck driver may choose an alternative trip that is unnecessarily long or worse, which becomes infeasible following another unforeseen event. It is therefore essential to provide drivers with real-time information, especially alternative itinerary suggestions when a planned route is impracticable. Also, such events can disrupt the entire supply chain. For example, a forest site becoming inaccessible may impact several trucks that are planned to visit this site. The dispatcher is then required to proceed to major changes in the initial transportation plan. In that case, the dispatcher must provide a new plan in real-time to the drivers while taking into account their current positions and whether or not their trucks are already loaded. Little time is available to solve such heavily constrained problem, hence the need for the development of a real-time decision support system.

Table 3.I presents a list of the unforeseen events that often disrupt transportation plans. We have classified these unforeseen events into three categories. The first covers the disturbances that are likely to appear at the forest sites. The second concerns those affecting trucks and road networks. The third category shows problems that occur at the wood mills. Note that this separation is not exclusive in the sense that, for example, a forest fire makes not only some forest sites inaccessible but also a set of forest roads. Also, an unexpected event in a supply chain component typically generates disturbances for the other components. For example, the slower pace of a loader at a forest site causes delay for the truck driver and then disturbs the timber reception plan at the wood mill.

Forest sites	Transportation	Wood mills
-Truck in the wrong site	-Road closure	-Loader not available
-Loader not available	-Too many trucks inside the same cycle	-Increase in demand
-Changing the loader's grapples	-Mandatory stop for weighing at check-points	- Queue for weighing at the mills
-Inventory inaccuracies	-Deficient forest road configurations (like sharp turns)	-Changing the loader's grapples
-Forest fires	-Road degradation, presence of a grader	-Decrease in storage capacity
-Pile of a product not available (behind a pile of another product)	-Truck breakdown, misdirected truck, unforeseen stop of the driver	-Decrease in demand
	-Poor trucks synchronization (some drive fast and some slow)	
	-Weather conditions: heavy rain, thawing soils, poor visibility	
	-Traffic jam or seasonal traffic increase (opening of the fishing or hunting seasons)	

Table 3.I – List of common unforeseen events

This project is in partnership with FPInnovations, a non-profit forest research center dedicated to the improvement of the Canadian forest industry through innovation. FPInnovations provided us with the case study presented in this paper. We assume that weekly transportation plans are developed to operate a fleet of homogeneous trucks in a supply chain with multiple products, multiple sources and multiple destinations. Drivers are given a priority ranking based on specific criteria established by transportation companies, such as seniority, performance or encouraging natives in some regions. A higher priority ensures the driver a higher amount of workload. Working hours are limited by regulatory constraints on the maximum number of consecutive driving hours and minimum rest hours between successive shifts.

We assume a context where every effort must be made to respect the agreements with truckers on the granted working hours for each week. The rationale behind this goal is to retain truckers for availability in the forest industry where there is competition for truckers' availability with other industries such as oil and mining. In this context, one must observe as much as possible the original plan agreed with the drivers, which limits possible changes in the plans when unforeseen events occur. In this paper, we present methods that deal with this lack of flexibility and produce consistent alternative plans in reaction to real-time disruptions. This context and the proposed methods differ significantly with what can be found in the literature on similar problems in forestry and other sectors. The next section reviews this literature.

3.3 Literature review

The construction of a transportation plan that satisfies the demand at the lowest cost, while respecting availability and other operational constraints (such as synchronizing trucks with loaders and not exceeding their capacities) is known in the literature as the log-truck scheduling problem (LTSP). The LTSP can be seen as a vehicle routing problem with pickup and delivery [14, 53]. The majority of case studies addressed in the literature relevant to this problem require that only a single vehicle serves each client only one time. In forestry, this is not realistic, as a single vehicle cannot satisfy the entire demand of a wood mill. Several visits are needed for each mill. This kind of problems is known as the vehicle routing problem with split deliveries [33, 94]. In addition, some constraints on the delivery time windows must be respected. Variants of this problem have been addressed [114]. But the approaches presented do not take into account the constraints of synchronization with the loaders that are crucial to the problem presented in this work. These constraints appear in the LTSP, introduced in the context of optimizing forest transportation plans [37, 40, 105].

Optimization models and operational research methods are increasingly used to solve planning problems in the forest industry [30, 107]. Among the most ambitious forest transportation optimization projects reported in the literature we cite the ASICAM project in Chile [120] and the EPO project in Finland [77]. The literature also contains contributions on routing and scheduling of logging trucks [9].

Recent technological advances have allowed the development of Intelligent Transportation Systems. A review of these systems and the contribution of operational research in their development can be found in [29]. They are already in use for the dispatching of vehicles in real-time in applications such as milk collection [35] or mail services [8]. There are also applications where the dispatching is relatively easier such as the cases of emergency vehicles and taxis dispatching in urban environments where it is not necessary, and usually not even possible, to assign more than the next trip. The most frequent disturbances in these applications are the emergence of new demands during operations. However, the nature of the problems addressed in the context of our work is fundamentally different from the problems found in the literature, since one must not only deal with changes in demand (e.g., the arrival of a new request), but also with changes in the topology of the transportation network (e.g., road closure). The other major difference between real-time transportation problems in forestry and problems in urban environments is at the level of the supporting road network. In fact, access to a forest site depends generally on a single forest road. It is usually impossible to find alternative ways to reach the same forest site, in contrast to urban areas, where the driver can access the same destination via an alternative route. This allows the dispatcher, in the urban context, to divert a vehicle from its initial destination to respond to a new request [17, 64, 102, 103].

The literature on forest transportation problems in real-time is scarce. The only work that we are aware of is the development of a decision support system called CADIS (for Computer Aided Dispatch) for a New Zealand forestry company that no longer exists [110]. For privacy reasons, few details were given about the system but it is reported that the results were 'of high quality' [108]. There are also other marketed decision support systems for forest transportation that may include modules for truck dispatching, but these modules are typically managed manually or using undocumented methods. It is reported [9] that, despite the lack of data in this area, FLO (for Logistics and Forest Optimization) would probably be one of the most advanced tools in the forestry industry. The American company Trimble, which offers the installation of onboard computers and communication means, is selling FLO. No indication is given on the possible presence of models or optimization methods to react to disruptions. Note also that in these systems, the dispatching of a vehicle is done only after receiving a call from the driver at the end of its delivery. Planning future trips can be made in advance, but they are not communicated to the drivers. In the event of unforeseen circumstances, they must be updated. However, some forestry companies may have a different mode of operation. Truckers require receiving their working schedule for the whole week. In the event of unforeseen circumstances, the flexibility decreases in this case since we must try our best to follow the original plan that was agreed with the drivers. In the next section, we present methods that can be used in this particular context.

3.4 Methodology

The starting point of the proposed approach is a predefined weekly transportation plan that can be derived manually or through an optimization method. This transportation plan is composed of a set of routes for each truck driver. Note that some routes in this set may be identical but different truck drivers operate them. As mentioned above, there is a priority ranking of the truck drivers and, therefore, we need to distinguish the routes by truck drivers. A route is then defined as a sequence of empty and full trips performed by a particular truck driver. A route also includes loading and unloading operations, as well as waiting and resting between successive shifts. We discretize the time into a set of equal intervals and express the duration of the transportation operations as a multiple of these intervals. The loading and unloading duration are approximately equal and we use their duration as a discretization step. We express the demand and supply in full truckloads since the fleet is homogeneous. To model the problem of selecting the best route for each truck while respecting the operational constraints, let us first define the input parameters of the model:

F = Set of forest sites,

M = Set of mills,

P = Set of wood products,

T = Set of trucks,

R = Set of feasible routes,

I = Set of time intervals,

c_r = Cost of route r,

c = Unmet demand penalty,

d_{mp} = Demand of product p at mill m,

s_{fp} = Quantity of product p available at forest site f,

n_{fmpr} = Quantity of product p transported from forest f to mill m on route r,

n_{fi} = Number of loaders available over time interval i at forest site f,

n_{mi} = Number of loaders available over time interval i at mill m,

L_{rfi} = 1 if route r includes a loading operation at forest site f over time interval i,
0 otherwise,

U_{rmi} = 1 if route r includes an unloading operation at mill m over time interval i,
0 otherwise.

The decision variables of the model are defined as follows:

x_{fmp} = Quantity of product p delivered from forest f to mill m,

δ_{mp} = Quantity of unsatisfied demand of product p at mill m,

y_r = 1 if route r is selected,
0 otherwise.

The mathematical model is inspired from the work [105] that deals with tactical (annual) transportation planning problems. The model was adapted for an operational (weekly) problem as follows:

$$
\begin{aligned}
\text{Min} \sum_{r \in R} c_r y_r &+ \sum_{m \in M} \sum_{p \in P} c\delta_{mp} && (3.1)\\
\sum_{f \in F} x_{fmp} + \delta_{mp} &= d_{mp}, \ \forall m \in M, p \in P && (3.2)\\
\sum_{m \in M} x_{fmp} &\leq s_{fp}, \ \forall f \in F, p \in P && (3.3)\\
\sum_{r \in R} n_{fmpr} y_r &\geq x_{fmp}, \ \forall m \in M, f \in F, p \in P && (3.4)\\
\sum_{r \in R} L_{rfi} y_r &\leq n_{fi}, \ \forall f \in F, i \in I && (3.5)\\
\sum_{r \in R} U_{rmi} y_r &\leq n_{mi}, \ \forall m \in M, i \in I && (3.6)\\
\sum_{r \in R} y_r &\leq |T| && (3.7)\\
y_r &\in \{0,1\}, \ \forall r \in R && (3.8)\\
x_{fmp} &\in R^+, \ \forall f \in F, m \in M, p \in P && (3.9)\\
\delta_{mp} &\in R^+, \ \forall m \in M, p \in P && (3.10)
\end{aligned}
$$

The objective function 3.1 minimizes both the costs associated with the routes and the penalty incurred if the demand is not satisfied. The cost of a route is the sum of the costs of all the activities that constitute this route. This cost depends also on the priority ranking of the truck driver. The penalty cost is chosen large enough to ensure that the demand is satisfied whenever it is possible. Constraints 3.2 and 3.3 aim at satisfying the demand while not exceeding the available supply. Constraints 3.4 ensure that the products are delivered only if the corresponding routes are selected. Constraints 3.5 and 3.6 ensure that the number of trucks using a loader over any time interval is less than or equal to the number of available loaders at this interval. Constraint 3.7 limits the number of selected routes to the number of available trucks. Constraints 3.8 ensure that the routes selection variables are binary and constraints 3.9 and 3.10 ensure the non-negativity of the flow variables and unmet demand variables respectively.

The number of routes that can be produced is generally large. The work [105] uses column generation to solve this model iteratively by considering only the subset of routes that have the potential to improve the objective function. After solving this model, the results are formatted to define weekly transportation plans for each truck driver and working schedules for each loader operator at wood mills and at forest sites. This allows defining the agreements with truckers on the workload that is granted for each week. In this particular context, one must try to respect the original plan agreed with the truckers, since the forest industry is competing with oil and mine companies to attract those drivers.

When an unforeseen event is revealed, we do not solve the model from scratch. We use heuristic methods to define new routes that satisfy the model constraints based on the actual solution. The consequences of the occurrence of an unforeseen event are, ultimately, a combination of canceled and/or delayed trips. This is why we started by developing recourse strategies for these two elementary disruptions with the aim to build on these simple events to develop strategies for more complex disruptions.

3.4.1 Canceled trip

As mentioned earlier, we want to limit the changes in the original transportation plan. To do so, we try to reschedule the canceled trip within the same week in such a way that a loader is free at the truck arrival time both at the forest site and at the mill. In this way, none of the other trips will be impacted or changed.

Figure 3.1 represents an excerpt from a transportation plan obtained through solving the mathematical model above. Each line represents the schedule of a truck. It is composed of a set of shifts. Each shift contains a set of trips that may be interrupted by some waiting. Let us assume that the first trip of the second truck is canceled. We can reschedule it at the beginning or at the end of any of the shifts of the same truck or the other trucks as long as the regulations on maximum driving hours and minimum rest time are met. To avoid impacting the other trucks plans, the truck chosen to operate this trip may have to wait for the loader to be free before being loaded or unloaded. One must try to minimize this waiting time as much as possible. Figure 3.2 shows one of the possible solutions where this trip was assigned to the first truck at the end of its second shift.

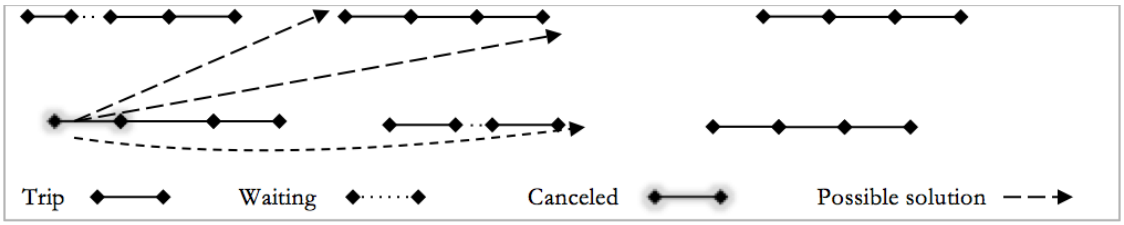

Figure 3.1 – Initial solution

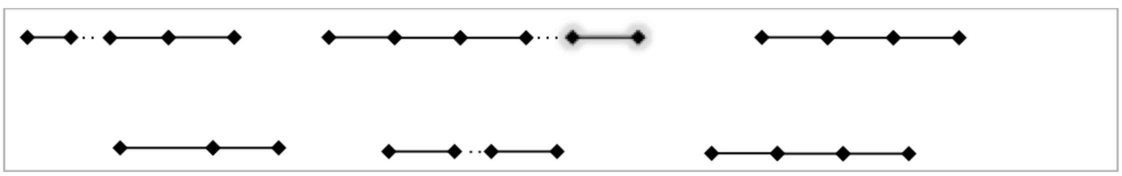

Figure 3.2 – New solution

To describe the algorithm that we developed to find all possible solutions, we use the following notation:

$Start_s^t$	=	Start time of shift s of truck t,
End_s^t	=	End time of shift s of truck t,
Max	=	Maximum consecutive driving hours,
Min	=	Minimum rest hours between two shifts,
L_{ij}	=	0 if the loader is free at site (forest site or mill) j over time interval i, Max otherwise (this prevents choosing time intervals initially assigned to other trucks),
del_s^t	=	Last delivery point of shift s of truck t,
b_t	=	Home base of truck t,
$d_{i_1 i_2}$	=	Distance between two sites i_1 and i_2,
f^*	=	Impacted forest site,
m^*	=	Impacted wood mill,
α_s^t	=	Waiting time at impacted forest site,
β_s^t	=	Waiting time at impacted mill.

The following algorithm allows identifying all the solutions that assign the canceled trip to a truck at the end of one of its shifts while respecting the constraints on driving and rest time:

for every truck t

  for every shift s of truck t

    for $\alpha_s^t \in \{0, 1, \ldots\}$

      $\theta_1 \leftarrow End_s^t + d_{del_s^t f^*} + \alpha_s^t$

      (the truck arrives at the impacted forest site and waits α_s^t time units)

      if $(\theta_1 - Start_s^t < Max$ **and** $Start_{s+1}^t - \theta_1 \geq Min)$

        for $\beta_s^t \in \{0, 1, \ldots\}$

          $\theta_2 \leftarrow \theta_1 + L_{f^* \theta_1} + d_{f^* m^*} + \beta_s^t$

          (the truck is loaded, arrives at the mill and waits β_s^t time units)

          $\theta_3 \leftarrow \theta_2 + L_{m^* \theta_2} + d_{m^* b_t}$

          (the truck is unloaded and goes back to its base)

          if $(\theta_3 - Start_s^t \leq Max$ **and** $Start_{s+1}^t - \theta_3 \geq Min)$

            store the solution

            break

          end if

        end for

      end if

    end for

  end for

end for

For solutions assigning the canceled trip to trucks before the start of one of their shifts, we use the same algorithm but instead of starting from the end of a shift, we start from the start of a shift and go backward to find a solution.

Since in the original transportation plan, the wood allocation was made from the impacted site to the impacted mill (optimally), we want to keep the same allocation. Accordingly, we ensure that this trip will be done. Knowing the distance that separates them, all that remains to do is to find a truck that would be able to reach the forest site in a time interval when the loader at this site is free and that allows reaching the mill in a time interval when the loader at the mill is also free. First, we verify if one of the trucks used in the current week could make the trip at the end of one of its shifts or prior to the start of one of its shifts. If the tool is able to find several solutions using only initially planned trucks, the user has the choice between three decision criteria for determining the trucker who will perform this trip. The user may choose to allocate this trip to the trucker who lost it if it is part of the found solutions. The second decision criterion proposed is the priority of the drivers. The third criterion is the cost of the solution. We study the impact of these criteria on the solution in the results section.

If no solution is found using the existing trucks, other strategies are available to the user. The first is to extend the opening hours of the wood mill in order to receive the delivery. Similarly, such an extension might show several potential solutions and the user should select a decision criterion as above. Another option is to check whether the transportation company could provide an additional truck to perform the trip. If so, the tool provides a preferential ranking of the bases or starting points if several additional trucks were available. It is also possible to combine the addition of a truck to the extension of working hours.

Some events may prevent keeping the initial wood allocation and the product being transported in the impacted trip must be delivered from another forest site. In this case, we make a ranking of the other sites where there is a surplus of this product according to the distance separating them from the delivery wood mill. We replace the impacted site by the top ranked forest site and repeat the process described above. After the selection of the final solution, we update the initial plans for loaders and trucks.

This user interaction and the multitude of the available options may seem heavy but it is intended. Indeed, the experience accumulated by FPInnovations staff showed them that their partners prefer to have control over their decision process. Certainly, the aim of this work is to support the users for making decisions, but they want to be able to choose the solution that best suits their needs among a small set of choices, rather than having to deal with one imposed 'optimal' solution.

3.4.2 Delayed trip

The second elementary event that we treat is the delay in the completion of a trip, which can be caused by various events such as the degradation of forest roads, lower visibility or traffic jams. When such a delay is observed, the truck arrives at a wood mill or at a forest site in a different time slot than the one it was allocated in the initial transport plan. It may, therefore, occupy a time interval that is initially granted to another truck. The latter is then obliged to wait for the loader to become free and will have then some delay relative to its initial plan. Any truck that joins the queue will also undergo delay. Similarly, each of these trucks will automatically have delay for their future trips, too. Thus, there is a propagation of the delay throughout the logistics system. In the case of a trip that is planned for the end of the working day of a loader, this propagation of delay could lead to deliveries arrivals after the closing time of the corresponding mill. This is why we capture this propagation, anticipate late arrivals (after closing time) and react in real-time.

First, we update the transportation plan to reflect this propagation of delay using the following algorithm where we define TP as the set of the impacted trips. In the beginning, this set includes the current (delayed) trip in addition to all the following trips of this truck that are planned for the current shift. The algorithm works as follows:

1. If $TP \neq \emptyset$, sort the trips according to their arrival times at a loader and go to 2; else, stop.
2. If the first truck arrives when a loader is free, update the loader and this truck plan and go to 1; else, go to 3.
3. If the truck arrives over a time interval initially planned for another truck, add one time slot to this truck time arrival for all the current shift trips, add them to TP and go to 1; else, go to 4.

4. If the truck arrives after business hours, store it as a late arrival and go to 1.

If our tool predicts a delivery after business hours, it offers the user the option to extend the business hours to receive this delivery. In the case such extension is impossible, this trip is considered as canceled. So we end up in the case of the first elementary unexpected event (canceled trip) and use the corresponding strategies.

3.4.3 Loader breakdown

When a loader breakdown occurs, we assume that we have an estimate of the failure duration. Having this information, the user specifies the maximum accepted waiting time of a truck at this loader. If, for example, a truck were planned to arrive at this loader a few minutes before the estimated time for completion of a repair, it would be better to accept delay for this trip instead of canceling it. Note again that this may not be the “optimal" solution, but this particular context favors simple adjustments instead of the complete re-optimization of the transportation plan.

Once the waiting time threshold is defined, we separate all trips heading to this loader into delayed trips and canceled trips (to reschedule). We start by treating delayed trips as they can lead, by propagation, to the cancellations of some trips, as explained above. These are added to the list of canceled trips and are processed by the corresponding strategies. The processing order of these trips is that of the priorities of the truck drivers that operate that trip. This is justified by the fact that the earlier a trip is treated, the greater the possibility to reschedule it in the same week. This ensures maximum volume of work agreed with the priority drivers. Again, the solution may not be optimal in terms of costs, but it is easier to be adopted by drivers and dispatchers. This method can be summarized as follows:

1. Go to the impacted loader plan and extract the trips related to the breakdown duration. The loader plan consists of a set of tuples ‘truck, product, origin/destination of the product’ for each time interval. This gives information about the trips.

2. Given the waiting time threshold, separate these trips into trips to cancel (TC) and trips to incur delay (TD).
3. Use the delay method to update the plan. Add trips arriving after business hours to TC.
4. Order TC by truck driver priority.
5. For every trip in TC, use the canceled trips method.
6. Update the plan.

Other disturbances can also be represented as loader failures. Forest fire or a forest road closure undermine the access to one or more forest sites and thus to the products there. The failure of loader at a forest site leads exactly to the same result. Similarly, several unexpected events resemble in terms of their consequences and the strategies presented in this section remain valid for their treatment. For this reason, we illustrate our approach with loader breakdowns only.

3.5 Results

FPInnovations provided us with a weekly planning problem based on real data. This problem is solved using an optimization method to obtain the transportation plan. In the week we consider, the number of available trucks is 62 but the solution obtained by the optimization method uses only 50 trucks to satisfy the demand. We assume that the user can have recourse to some of the unused trucks to operate some trips if the initial transportation plan is disrupted due to an unforeseen event. As mentioned above, the fleet of trucks is homogeneous but there is a priority ranking of the drivers. There are 4 wood mills, 59 forest sites and 12 different products. The total demand is 542 full truckloads for the week we consider. The mills business hours vary from 6am to 7pm from Monday to Friday. However the drivers can start their shifts before 6am and end their shifts after 7pm. Initially, there are no restrictions on the schedules of the loaders at forest sites but the optimization method, by minimizing the transportation costs, minimizes also the operating hours of the loaders. In the literature, one can find models that consider explicitly the loaders scheduling, like [40]. The loading and unloading times are estimated at 20 minutes for this case study. Therefore, we use 20 minutes steps to discretize the planning horizon. The trip duration between forest sites and wood mills ranges from 1.5 to 4.5 hours and depends on whether the trucks are loaded or not. The approximate driving cost is around $100 per hour and the waiting cost is about $75 per hour. The difference between loaded and empty driving costs is captured in the duration of these trips. The maximal driving time per shift is 14 hours and the minimal rest time between shifts is 10 hours.

We start by testing how the tool reacts to single trip cancellations and delays. When a trip must be rescheduled, the user is provided with different decision criteria. In the current practice, there is no standard recourse used to react to unforeseen events. Therefore, we cannot predict the user behavior regarding the choice of the options offered by the tool. Moreover, this behavior and the available options vary depending on the usage context. For example, the efficiency of adding a new truck depends on the time when the truck becomes available, as well as on the price charged by the truck owner. Hence, several scenarios are developed to examine how the different options impact the resulting solutions and to achieve a proof-of-concept of the ideas presented in this paper. To reschedule a trip with respect to the regulation on driving and rest time, the following scenarios are defined:

- Scenario 1: We try to assign the trip to one of the existing trucks. If different solutions are found, the trip is assigned to the least costly truck. This is a classical scenario where the main objective is to minimize the costs. We use this scenario as a reference to compare the results of the other scenarios.
- Scenario 2: We try to assign the trip to one of the existing trucks. If different solutions are found, the trip is assigned to the truck that has the highest priority.
- Scenario 3: We try to assign the trip to the impacted truck (that lost the trip).
- Scenario 4: We try to assign the trip to one of the existing trucks. If no solution is found, the mill opening time is extended to allow the truck with the highest priority to operate the trip.

- Scenario 5: We try to assign the trip to one of the existing trucks. If no solution is found, the mill opening time is extended to allow the least costly truck to operate the trip.
- Scenario 6: We try to assign the trip to the impacted truck. If no solution is found, the mill opening time is extended to allow this truck to operate the trip.
- Scenario 7: We use an additional truck to operate the trip.
- Scenario 8: We extend the mill opening time and use an additional truck to operate the trip.

We cancel one of the planned trips and measure the performance of the tool under each of the 8 scenarios. We repeat the test for every single trip of the weekly transportation plan that we consider. This allows testing the tool under a wide scope of experimental conditions ranging from early cancellations that are generally easier to reschedule to late cancellations happening at the end of the week. Table 3.II presents a summary of the results comparing the effectiveness of the tool under the different scenarios. We first compare the average transportation costs generated by the solutions. Note that the extension of the mill opening time and the usage of additional trucks generate extra costs that are not included in the results, since they depend on how the transportation is managed and who owns the trucks (for instance, penalties may be incurred if a trip is canceled for some drivers). Only the routing costs are used in Table 3.II. The first scenario is used as a reference with regards to the cost. The third column of Table 3.II gives the percentage of canceled trips that the tool was able to reschedule within the same week under the different scenarios. The last column shows the changes in the fleet of trucks usage rate.

It can be seen that scenarios 2 and 3 yield solutions with the highest costs. This could be easily predicted since the cost of solutions found for the highest priority truck or for the impacted truck is at best equal to the least costly solution. However, the average cost in scenario 3 is slightly lower than the one in scenario 2. This can be explained by the percentage of solutions found under these scenarios: a solution using the impacted truck (scenario 3) is found in only 2.58% of the instances, while under

Scenario	Cost	Rescheduling	Truck usage
1	0.00%	51.85%	0.55%
2	35.54%	51.85%	0.74%
3	34.86%	2.58%	0.73%
4	-5.77%	91.14%	0.51%
5	5.23%	91.14%	0.58%
6	-1.18%	10.52%	0.54%
7	-14.70%	88.75%	-0.47%
8	-15.29%	97.60%	-0.47%

Table 3.II – Results for one canceled trip

scenario 2, a solution is found in 51.58% of the instances. The latter includes some solutions with high costs that affect the overall average cost. Scenarios 4, 5 and 6 allow the extension of the mill opening time. This improves the performance of the tool with regards to its ability to reschedule the trips. The number of solutions using existing trucks increased from 51.58% to 91.14%. Although the extension of the mill opening time generally allows finding less costly solutions, it generates also new solutions that could not be found under the first three scenarios. These solutions can have high costs that impact negatively the average cost (this can be observed under scenario 5).

Adding a new truck in scenario 7 yields a rescheduling capacity of 88.75%, which is lower than in scenarios 4 and 5, where the mill opening time is extended. However, it finds better solutions with regards to the costs. This reduction is even more important when the addition of a truck is combined with the extension of the mill opening time (scenario 8). But one must keep in mind that this additional flexibility comes at a certain cost. Taking this extra cost into account, the user can decide whether or not to use these options. It should also be noted that the rescheduling capacity of the tool never attains 100%. This is due to the trips that are planned at end of final day of the week, since there is not enough time left to reschedule these trips within the same week.

The last column shows that the overall usage of the trucks is improved for all the scenarios except the scenarios where a new truck is added. These results could be easily predicted but one must note that these solutions can be found only if the reaction is done in real-time. Otherwise, these opportunities are lost. Our strategy of allocating trips at the end or before the start of a shift allows the use of under-utilized trucks, increasing their productive hours and their annual tonnage. This key indicator could help convince the forest companies of the importance of using such tools.

To measure the impact of the tool on delayed trips, we developed a random generator that selects an impacted trip and delay duration. In all the conducted experiments, our tool anticipated at least one truck arrival after the closing time of the mill. In the current practice, those trucks would have to wait until the morning to be unloaded, which disturbs the plan of the following day. The proposed approach allows avoiding unnecessary trips to inaccessible points. These trips are processed as canceled trips and since the cancelations of all the trips of the studied transportation plan were considered in the first set of tests, we do not reproduce the results here.

We explained in the methodology section that many events have the same consequences on the transportation plan as a loader breakdown. Therefore, we use the loader breakdown example to measure the performance of the tool for more complicated disruptions. One involved loader is chosen randomly and the impacted trips are chosen depending on the breakdown occurrence time and its duration. The latter are also randomly generated. The first half of the trips is canceled. The other half is delayed. Recall that delay could generate extra canceled trips. We define three different test settings that we refer to as:

- Light: The breakdown duration ranges between one hour and two hours.
- Medium: The breakdown duration ranges between three and five hours.
- Hard: The breakdown duration ranges from six hours up to thirteen hours, which is the number of daily working hours. In fact, it attains 13 hours only if the loader is down exactly when the loader starts the day. In average, the duration remains between six and seven hours.

Note that when the loader is down for one hour, three trips are impacted on average. In fact, the loading/unloading time is approximately equal to twenty minutes. Therefore, three trucks will not be able to load/unload their shipments during a one hour breakdown.

The occurrence time of the breakdowns is chosen to be in the beginning or in the middle of the week. We do not choose breakdowns for the last day of the week since we know that the impacted trips cannot be rescheduled within the same week. For a real application, the tool can easily handle a rolling horizon where we try to reschedule the trips within the following seven days rather than in the same week. Note also that we ensure that the same random numbers are used for all the scenarios to have a fair comparison.

Since we have to reschedule a set of trips rather than only one trip as in the first tests, scenarios 3 and 6 are modified. When a solution using the impacted truck is not available, we use the highest priority truck instead. We also added a new scenario:

- Scenario 9: We try to assign the trip to the initial trucks. If the trip could not be rescheduled, we use an extra truck.

Table 3.III presents the results for the different instances. The average cost under scenario 1 is used as a reference. The costs used in Table 3.III are the costs per trip. We do so because the number of trips rescheduled varies in each scenario. The results using the impacted trucks are approximately equal to those using the trucks with higher priority (scenario 2 versus scenario 3, or scenario 5 versus scenario 6). This is explained by the fact that when no solution is found for the impacted truck, the trip is assigned to the highest priority truck. Scenarios 7 and 8 give exactly the same results. This means that the extension of the mill opening time does not improve the solutions obtained when using extra trucks. The absence of impacted trips late in the week may explain that phenomenon. Using extra trucks only for unassigned trips (scenario 9) generates additional costs that attain about 74% in the 'Hard' setting. This allows rescheduling all the trips but at a higher cost. As shown by the results of scenarios 7 and 8, one may

conclude that it is always better to assign all the trips to extra trucks. However, these costs do not include the extra fees for adding a new truck nor the penalties that are incurred if the agreement made with the drivers on the workload volume is not respected.

The extension of the mill opening time produces interesting results. Comparing scenarios 1 and 4, the extension of the mill opening time increases the cost. The difference can be explained by the fact that the extension of the mill opening time generates new solutions that can be costly but allows improving the rescheduling performance. This raises the question of finding the right tradeoff between satisfying the demand and minimizing the costs. Also, comparing scenarios 2 and 3 to scenarios 5 and 6, respectively, shows that even if the extension of the mill opening time increases the number of rescheduled trips it also improves the cost, except in the 'Hard' setting. Therefore, in scenario 4, we can conclude that the improvement of the solutions found in scenario 1 was not enough to offset the high costs of the new scheduled trips.

As for the rescheduling capacity, we note that it is better than in the single trip cancellation tests (Table 3.II). This is explained by the fact that very late trips were discarded from these tests. We note also that increasing the duration of the breakdown decreases the rescheduling capacity for the first six scenarios. In the last scenarios, there were enough trucks to operate the impacted trips. There was also enough time to reschedule the trips within the same week.

Compared to current practices, the tool will allow increasing the overall efficiency of a fleet through the quicker responses resulting in less lost time and the reduction of truck costs over the long term. However, the results show that only 56% to 62% of the trips could be rescheduled under the first three scenarios, which do not allow the extension of the mill opening time and the use of additional trucks. The density of the loader plans, which contain only 15% of free time intervals, can explain this phenomenon. We use two additional case studies to confirm this explanation. In these case studies that we denote Case 2 and Case 3, we consider the same supply chain but we decrease the

	Light		Medium		Hard	
	Cost	Rescheduling	Cost	Rescheduling	Cost	Rescheduling
Scenario 1	0.00%	61.67%	0.00%	57.33%	0.00%	56.67%
Scenario 2	88.32%	61.67%	81.73%	57.33%	73.08%	56.67%
Scenario 3	88.71%	61.67%	81.98%	57.33%	73.24%	56.67%
Scenario 4	5.57%	89.35%	21.89%	86.67%	28.75%	80.29%
Scenario 5	52.48%	89.35%	71.73%	86.67%	80.68%	80.29%
Scenario 6	51.18%	89.35%	70.82%	86.67%	80.00%	80.29%
Scenario 7	-2.78%	100.00%	12.15%	100.00%	18.37%	100.00%
Scenario 8	-2.78%	100.00%	12.15%	100.00%	18.37%	100.00%
Scenario 9	16.12%	100.00%	38.60%	100.00%	73.95%	100.00%

Table 3.III – Results for one loader breakdown

demand to lower the density of the loader plans. The demand is decreased by 15% in Case 2 and by 25% in Case 3. The free time intervals now constitute 28% and 37% of the new loader plans in Case 2 and Case 3, respectively. Case 1 refers to the initial case study analyzed above. We use the 'Medium' test setting to compare these three case studies. For the ease of comparison, we reproduce also the results of Case 1 in Table 3.IV.

	Case 1		Case 2		Case 3	
	Cost	Rescheduling	Cost	Rescheduling	Cost	Rescheduling
Scenario 1	0.00%	57.33%	0.00%	94.56%	0.00%	100.00%
Scenario 2	81.73%	57.33%	98.65%	94.56%	86.27%	100.00%
Scenario 3	81.98%	57.33%	95.38%	94.56%	46.11%	100.00%
Scenario 4	21.89%	86.67%	-4.38%	100.00%	0.00%	100.00%
Scenario 5	71.73%	86.67%	32.45%	100.00%	44.26%	100.00%
Scenario 6	70.82%	86.67%	18.04%	100.00%	25.62%	100.00%
Scenario 7	12.15%	100.00%	-6.99%	100.00%	0.00%	100.00%
Scenario 8	12.15%	100.00%	-6.99%	100.00%	0.00%	100.00%
Scenario 9	38.60%	100.00%	3.27%	100.00%	0.00%	100.00%

Table 3.IV – Comparison of case studies

The rescheduling capacity in Case 2 and Case 3 has increased as predicted. Also, scenario 2 is now quite different compared to scenario 3 (the same can be observed when comparing scenario 5 and scenario 6). This is explained by the fact that the tool is able to find more solutions for the impacted trucks in Case 2 and Case 3. Also, the highest priority trucks seem to generate higher costs. The remaining results have similar characteristics than those obtained in Case 1.

One must keep in mind that the results presented in this section are averaged over a set of instances. In a real application, each case must be examined individually. For example, one might conclude from the results presented in this section that adding an extra truck always yields the best solution with regards to the routing costs. However, for a single trip, this would not be the case if the home base of the available additional truck were far from the impacted sites. This is another reason for having many options available to the users.

The tool aims to assist the dispatchers in making the best combinations of recourse strategies on the basis of cost, rescheduling capacity, and the nature of the unforeseen event. However, one key element of real-time decision support systems is the computational time. The advantage of the presented heuristic methods is that they are solved in a few seconds. The drawback of these methods is that they do not always produce the optimal solution with regards to the cost. But the context that we study in this paper favors adjustments of the transportation plan rather than complete re-optimization. Current work includes developing mathematical models and exact methods to deal with contexts where the defined agreements with the truckers allow the re-optimization of the disrupted transportation plan. This is the case, for example, when only one trip at a time is communicated to the drivers, instead of a weekly plan.

3.6 Conclusion

The goal of the tool presented in this paper is to re-optimize a disturbed transportation plan, in response to some unforeseen events that may happen while performing the trips between the forest sites and the wood mills. The input to the tool is a weekly plan (obtained manually or after using an optimization method) and some unforeseen events. The testing prototype includes a generator of unforeseen events and implements heuristic methods used to reschedule the trips depending on the type of unforeseen event. The results show that the tool is able to reschedule, for most of the tested instances, at least half of the trips that could be lost if no reaction was performed in real-time. This improves the overall usage rate of underutilized trucks over the long term. The instances are derived from a real application and tested under generated unforeseen events. The developed approach allows the production of alternative solutions in little computation time. The costs of the solutions depend on the decision criteria that are considered. These include minimizing costs, increasing the priority trucks workload or re-assigning the lost trips to the impacted trucks. Recourse strategies try to reschedule the trips using the existing trucks, adding extra trucks or extending the mill opening time. The last two options generate extra costs that vary depending on the configuration of the supply chain. The analysis of the case studies, however, excludes these costs and compares the different solutions only on the basis of the routing costs. The results showed that these two options allow rescheduling more trips. This avoids incurring penalty costs for unsatisfied demand. The user should then find a trade-off between satisfying the demand and minimizing the average cost per truckload.

The context of utilization that we consider in this paper is rather rigid in that it is necessary to provide truck drivers with weekly plans and adhere to these plans as much as possible. Better solutions can be found if the dispatcher is allowed to make more changes to the original plan. We are currently developing approaches to deal with different contexts such as when only one trip at a time is communicated to drivers. In these contexts, the dispatcher waits until the driver finishes its trip before revealing the next trip. This context is more flexible and provides more recourse possibilities. Other contexts offer more flexibility by including trucks equipped with onboard loaders. These trucks are especially useful when a loader breakdown happens. They can also be used in forest sites where only small volumes of wood are available, thus avoiding the use of a loader in these sites.

The unforeseen events take different characteristics at each occurrence. Some events are even interrelated. For example, a forest road degradation could lead to its closure with a given probability. The aim of this paper is to present an approach to react to transportation disruptions. Collecting and validating the relevant data and designing the statistical processes that describe these interrelations is beyond its scope. This study is a topic for future research that fits into the current developments of Big Data in the forest sector. Indeed, more and more log trucks are equipped with onboard computers and communication technologies that can be used for this purpose.

CHAPTER 4

ARTICLE 2: REAL-TIME MANAGEMENT OF TRANSPORTATION DISRUPTIONS IN FORESTRY

Chapter notes

This chapter has been submitted for publication in *Computers & Operations Research*. It was also published as a technical report in CIRRELT [2]. The paper ranked 3rd in the 2016 David Martell Student Paper Prize in Forestry. Preliminary work was presented at the following conferences:

- 9th Triennial Symposium on Transportation Analysis (TRISTAN IX), Oranjestad, Aruba, June 13-17, 2016
- 58th CORS Annual Conference, Banff, Canada, May 30 - June 1, 2016
- SSAFR 2015, Uppsala, Sweden, August 19-21, 2015
- CORS/INFORMS International Conference, Montreal, Canada, June 14-17, 2015
- 3rd Annual FIBRE Conference, Montreal, Canada, May 11-13, 2015
- VCO Workshop, Montreal, Canada, February 27, 2015

Please note that this paper uses a notation slightly different than the previous one.

Real-time management of transportation disruptions in forestry

Amine Amrouss [1,2,*], Nizar El Hachemi [1], Michel Gendreau [1,3], Bernard Gendron [1,2]

[1] *Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation (CIRRELT), Canada*

[2] *Department of Computer Science and Operations Research, Université de Montréal, P.O. Box 6128, Station Centre-Ville, Montreal, Canada H3C 3J7*

[3] *Department of Mathematics and Industrial Engineering, Polytechnique Montréal, C.P. 6079, succursale Centre-ville, Montreal, Canada H3C 3A7*

In this paper, we present a mathematical programming model based on a time-space network representation for solving real-time transportation problems in forestry. We cover a wide range of unforeseen events that may disrupt the planned transportation operations (e.g., delays, changes in the demand and changes in the topology of the transportation network). Although each of these events has different impacts on the initial transportation plan, one key characteristic of the proposed model is that it remains valid for dealing with all the unforeseen events, regardless of their nature. Indeed, the impacts of such events are reflected in a time-space network and in the input parameters rather than in the model itself. The empirical evaluation of the proposed approach is based on data provided by Canadian forestry companies and tested under generated disruption scenarios. The test sets have been successfully solved to optimality in very short computational times and demonstrate the potential improvement of transportation operations incurred by this approach.

Keywords: Real-time, transportation, forestry, mathematical programming.

4.1 Introduction

Optimization models and operations research (OR) methods have been used in the forest industry since the 1960s [121]. Recent reviews on how these models and methods are used to solve planning problems in forestry can be found in [16, 109]. These planning problems cover a wide range of activities such as sylviculture, harvesting, road building, production and transportation, which present to this day several challenges to OR practitioners [82, 109], as the forest industry attempts to improve its competitiveness and reduce its environmental impact. In particular, improving transportation planning in forestry has been the object of recent research of highly practical relevance, since transportation costs are estimated at more than one-third of wood procurement costs [109]. Minimizing transportation costs therefore represents a key element to improve the competitiveness of forest companies.

Recently, a number of OR models and methods have been developed to solve the *log-truck scheduling problem* (LTSP) [40, 57, 99, 105, 120], which consists in deriving schedules for trucks to transport different wood products between forest sites and wood mills. In addition, several decision support systems, such as the ASICAM project in Chile [120] and the EPO project in Finland [77], were developed to ease transportation planning. A review of transportation planning systems in the forest industry and the contribution of OR in their development can be found in [9]. Note that few decision support systems are available to forest companies (compared to other industrial sectors [24]), as many forest companies still rely on experienced dispatchers to manually derive their transportation plans.

Whether the transportation plans are obtained through an optimization method or manually, their implementation in practice is vulnerable to unforeseen events. For example, in Canada, spring thawing soils and summer rains degrade the forest roads condition and prevent the trucks from accomplishing their trips within the planned time. The late arrival of these trucks may also create queues for loading and unloading operations. In this case, the disruption consequences may stream through the whole supply chain and many trips could become infeasible. There is then a need to re-optimize the transportation plan as early as possible to minimize the impact of such disruptions. Real-time rescheduling of log-trucks has not been subject to much attention in the literature, in spite of the growing body of literature on similar problems in other industrial sectors, with the advent of intelligent transportation systems [29]. To the best of our knowledge, CADIS (for Computer Aided Dispatch) is the only documented decision support system for real-time dispatching in forestry [110]. The authors reported very few details about this system because of non-disclosure agreements with the New Zealand company that used it. The system produced encouraging results [108], although it was used only for a short period, as the company ceased its activities because of financial issues. Other commercial decision support systems [9] may include real-time dispatching modules, but they are generally manually managed. The recent work [109] defines real-time transportation management as one of 33 open problems in the forest industry for OR practitioners.

The most frequent source of uncertainty related to transportation planning problems in other industrial sectors is the arrival of new requests (e.g., new customers or change in the demand) [93, 111]. In forest transportation planning problems, one must deal with unforeseen events of a different nature such as changes in the topology of the transportation network (e.g., road closure). In this paper, we propose a mathematical programming model that remains valid for every unforeseen event that may occur during forest transportation operations, regardless of its nature. The model is based on a time-space network representation of the forest supply chain where the impacts of the unforeseen events are represented.

The remainder of this paper is organized as follows. Section 4.2 describes the problem, starting with a generic description of the LTSP. Section 4.3 presents the proposed approach to re-optimize the transportation plan in real-time in response to an unforeseen event. The description of the test sets and the results of our approach are presented in

Section 4.4. Section 4.5 concludes this work.

4.2 Problem description

We begin this section with a generic description of the LTSP, whose solution produces a transportation plan that consists of a sequence of empty and loaded trips in addition to loading and unloading operations. Note that our approach remains valid whether such a plan is derived manually or by using optimization methods, but the LTSP provides a conceptual framework for the subsequent development of our model for real-time rescheduling of log-trucks.

We assume a homogeneous fleet of trucks. Each truck is associated with a base, usually a wood mill, where it must begin and end its shift. The planner must assign a route to each truck over a planning horizon of one week. A route is composed of a set of trips in addition to waiting, loading and unloading operations. We define as R, V, M, and F the sets of routes, trucks, mills and forest sites, respectively. R_v is the subset of routes linked to truck $v \in V$. Each route $r \in R$ has a cost c_r. This cost includes productive (loaded trips, loading and unloading) and unproductive activities (empty trips and waiting). The basic LTSP aims at minimizing the total cost while satisfying the demand D_m at each mill $m \in M$ given a certain amount of available wood products S_f at each forest site $f \in F$. The problem can be formulated as follows [107]:

$$\text{Min} \quad \sum_{r \in R} c_r y_r \tag{4.1}$$

$$\sum_{r \in R} b_{mr} y_r = D_m, \ \forall m \in M \tag{4.2}$$

$$\sum_{r \in R} a_{fr} y_r \leq S_f, \ \forall f \in F \tag{4.3}$$

$$\sum_{r \in R_v} y_r = 1, \ \forall v \in V \tag{4.4}$$

$$y_r \in \{0,1\} \ \forall r \in R \tag{4.5}$$

The variables y_r are 1 if route r is chosen and 0 otherwise. The parameters a_{fr} (b_{mr}) represent the total amount of products picked up at forest site f (delivered at mill m) if route r is selected. The objective function (4.1) minimizes the total cost. Constraints (4.2) and (4.3) ensure demand satisfaction while not exceeding the supply. Constraints (4.4) ensure that at each truck is assigned a route.

Note that the transportation cost includes a fixed cost for using a truck and a variable cost proportional to the distance that depends on whether the truck is empty or loaded. The trucks have to travel empty from the mills to the forest sites. Thus, a truck that operates only trips between the same mill and the same forest site loses half of its transportation capacity. Instead, once at a mill, one must try to allocate the wood products from the closest forest sites to the mills in the opposite direction. This is known in the literature as backhauling and we refer the interested reader to [25] for more details about decision support systems using backhauling in the forest industry.

Loading and unloading operations are performed by loaders at forest sites and mills. These loaders are usually operated only for a specific period of the day. Moreover, the number of loaders available at a mill or a forest site may vary during the day. To avoid creating queues at the loaders and thus reduce the cost of unproductive activities, another objective that must be met by the dispatcher is the synchronization of the trucks with the loaders given accurate information about the available loaders. These constraints appear in the recent works on the LTSP [40, 105] and are considered in our work.

In the context of real-time rescheduling of log-trucks, we assume that truck drivers receive one trip at a time, the dispatcher waiting for each truck driver to finish its current trip before revealing its next destination. This mode of transportation planning management gives more flexibility to re-optimize the routes, since it avoids drivers resistance to change.

While re-optimizing the transportation plans following the occurrence of an unforeseen event, the dispatcher must avoid diverting a truck from its destination unless the unforeseen event prevents the completion of the current trip. This improves the consistency of the proposed schedules and facilitates their real-life implementation. Moreover, in a real-time context, the amount of time available to the dispatcher to derive alternative transportations plans is limited.

The nature of the unforeseen events that arise in the forest industry is distinct from what can be found in the literature on similar problems found in other industrial sectors. We have drawn up a list of the most frequent unforeseen events. The list includes unforeseen events that are likely to appear at the forest sites, those involving trucks and road networks, and the events that occur at the mills. To develop effective recourse strategies when facing such events, one must focus on the impacts they have on the transportation network rather than on the events themselves. The next section describes the proposed approach to implement these recourse strategies.

4.3 Proposed approach

Our approach to real-time rescheduling log-trucks is built on a time-space network representation, which is used in the definition of our mathematical programming model. The time-space network represents the evolution of the forest supply chain over time. This representation varies depending on the nature of the unforeseen events that are revealed over time. The space and time dimensions of the network allow to track the trucks in real-time and to capture the impacts of the unforeseen events on the transportation network (e.g., by removing the arcs that become inaccessible). The distances between two locations in the transportation network are expressed as a time measure. This helps to capture the impact of some unforeseen events. In the case of a road degradation or a traffic jam, for example, the trip duration may become longer, while the geographical distance remains the same. The mathematical programming model takes this time-space network as an input and is solved using a commercial solver.

4.3.1 Time-space network

When an unforeseen event is revealed, one must collect real-time information about the state of the transportation network elements. We refer to the state of a truck, for example, as the information about its position, its destination and the product it is transporting if it is loaded. Moreover, if the truck is directly impacted by the unforeseen event as in the case of a truck breakdown, we assume that we have additional information about the estimated characteristics of the corresponding event, such as an estimate of the truck repair duration. The collection and validation of these estimates is beyond the scope of this work, but the current development of onboard computers, geo-location and communication technologies, in addition to the development of big data algorithms, make the collection of good quality estimates of the disruptions characteristics more affordable and easier.

The state of the transportation network can be seen as an instant picture of this network that we represent as a time-space network. The space dimension of the network contains the set of wood mills and forest sites in addition to their linking roads. For the time dimension of the network, we divide the planning horizon into a set of intervals. The necessary time for loading and unloading operations is approximately equal and the driven distances are quite large in the context we consider in this paper. Therefore, we use the loading duration as a time-step for discretizing the planning horizon. The time-space network representation (Figure 4.1) contains four types of vertices :

- A *source vertex for each truck* representing its current location (or its base if it has not yet started its shift) when the unforeseen event is revealed. These individual truck vertices are different from what can be found in a conventional time-space network. We need to introduce them to track the truck positions in real-time. Note also that the trucks that finish their shift before the occurrence of the disruption are not represented in the network.
- A *sink vertex for each truck*. It corresponds to its base and represents the shift end for the truck.

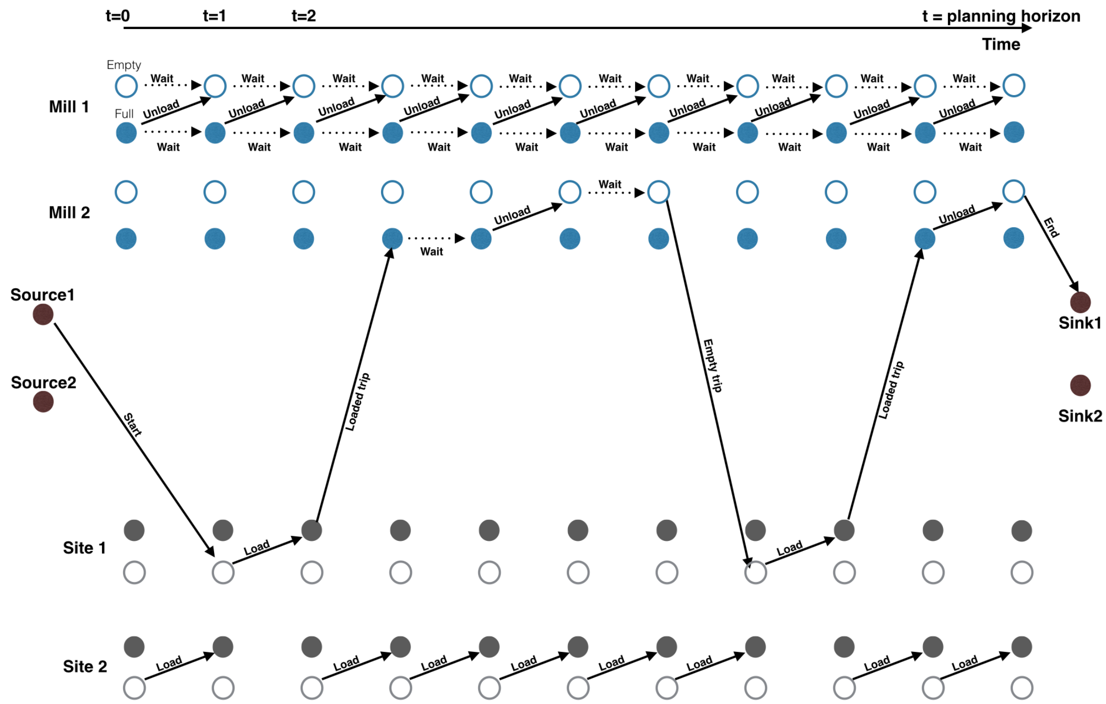

Figure 4.1 – Time-space network

- *Forest site vertices.* Each vertex is replicated for each time interval of the discretized planning horizon. This allows to capture real-time information about the forest sites. This includes the current supply of each product and the number of loaders available at the correspondent interval. These vertices are duplicated to represent whether the truck is full or empty.
- *Mill vertices.* They are similarly replicated. The vertex state contains information about the current demand for each product and the number of loaders available at the corresponding interval.

The replication of the vertices is done horizontally in Figure 4.1. Each pair of lines represents either a mill or a forest site evolving over time. For reasons of clarity, only a subset of the arcs is represented in Figure 4.1 and their length does not represent the real distances. The arcs kept for the first truck give an example of a small sequence of trips. There are seven types of capacitated arcs in the time-space network:

- Start arcs connecting source vertices to empty forest site vertices, if the corresponding truck is empty, and to full mill vertices, otherwise. Their capacity is one truck.
- End arcs connecting empty mill vertices that correspond to a truck base to this truck sink vertex. Their capacity is one truck.
- Loaded driven arcs connecting a full forest site vertex to a full mill vertex demanding at least one of the available products at this forest site. Their capacity is equal to the number of available trucks.
- Empty driven arcs connecting an empty mill vertex to an empty forest site vertex supplying at least one requested product. Their capacity is equal to the number of available trucks.
- Waiting arcs connecting two successive mill vertices. Note that, as the number of mills is usually smaller than the number of forest sites and to reduce the symmetry, we prefer that the trucks wait at mills instead of forest sites. Their capacity is equal to the number of available trucks.
- Loading arcs connecting two successive empty and full forest site vertices. Their capacity is equal to the number of available loaders.
- Unloading arcs connecting two successive full and empty mill vertices. Their capacity is equal to the number of available loaders.

It should be noted that the length of the arcs represents the duration of the corresponding operation. Therefore, these arcs exist only between vertices at intervals separated by at least this duration. Moreover, the vertices and arcs constituting this time-space network vary over time and depend on the nature of the revealed unforeseen events. We describe how these transformations are done in the following subsection.

4.3.2 Dealing with disruptions

At the occurrence of an unforeseen event, we first collect the necessary information about the trips that were executed before the disruption in order to update the remaining demand and the number of trucks still in operation. We also collect the relevant information about the trucks, their positions and if they are loaded or empty. Having this information in addition to the estimates of the unforeseen event impacts, a new time-space network is produced. All the vertices and arcs that start before the occurrence of the event are removed from the initial time-space network. One exception is the truck start vertices. Outgoing arcs from these start vertices are updated according to the nature of the unforeseen event and to the corresponding truck positions.

The recourse strategies when an unforeseen event is revealed depend on its impact on the transportation network rather than on the event itself. Different unforeseen events can have the same impact on the transportation network. For example, in the case of the presence of a single loader at a forest site or at a mill, the breakdown of this loader can be seen as the corresponding site closure, assuming that the loaders are not allowed to move between different sites and that the trucks do not include onboard loaders. The following describes the disruptions categories based on their impact on the network, in addition to the corresponding recourse strategies.

Closures

This category contains the closures of forest sites, wood mills and roads. Also, there is generally one single forest road to access a forest site in contrast with urban context where the same point may be reached by different paths. Therefore, the closure of such road can also be considered as a forest site closure. A mill closure means that no product can be delivered to this mill during the closure. This can be caused, for example, by a decrease in the storage capacity or by the breakdown of the loader associated with this mill.

In the event of such disruptions at a mill or at a forest site, we remove the loading or unloading arcs at the corresponding vertices in addition to outgoing driven arcs for all time intervals that lie within the estimated duration of the disruption. We keep the waiting arcs at the mills. For trucks planned to arrive at the closed vertices before the operations start back, their start vertices are connected to the other mills or forest sites depending on whether they are loaded or not. The remaining truck start vertices are connected to their current destination at the time the disruption is revealed. The rest of the network is unchanged. If the disruption occurs on a road linking a mill to a forest site, we remove the corresponding arcs in the network for all the time intervals that lie within the closure duration.

Delays

Delays can be caused by a variety of unforeseen events. This includes bad weather conditions (poor visibility, thawing soils, heavy rains), degradation of forest roads, traffic jams, opening of hunting or fishing season and so on. Delays can be observed at a single truck level. This is the case, for example, when the truck is undergoing some mechanical issues and thus slowing down. In contrast, when a forest road is damaged, for instance, all the trucks taking this road will be impacted.

When a truck is delayed, we link its start vertex to its current destination vertex but at an interval that takes into account both the remaining distance and the estimation of the delay. For delays observed between two vertices, we move the arcs to take into account the delay estimation. We do so for all the arcs that lie within the estimation of the duration necessary to return to normal operations.

A truck breakdown can also be seen as a delayed truck. We assume that we have an estimate of the necessary time to repair this truck. If the repair time does not exceed the planning horizon, the arrival time of the truck to its next destination is delayed by the repair duration. Otherwise, we just remove the truck from the network.

Demand and supply variations

Mill breakdowns may lead to a decrease in its storage capacity. The demand of some products must therefore be adjusted downwards. Also, we may have an increase in the demand for some products. If the mill is not already connected to forest sites where the product is available, we add empty and loaded driven arcs between the mill and these forest sites. We also adjust the demand parameter in the input data. Similarly, if, during the day, we have more accurate data about the supply, its parameter is updated in the input data.

Loader breakdowns

We assume to have an estimate of the necessary repair time and we update the number of available loaders during this period. When an arc is modified in the network, its cost is also updated according the nature of the disruption. Once the new time-space network is obtained, it is combined with the new cost matrix, the remaining demand, and the number of available trucks and loaders. These constitute the input parameters of the mathematical programming model.

4.3.3 Mathematical programming model

A two-phase approach for solving a weekly LTSP is introduced in [38]. The authors solve, in the first phase, a tactical MIP to assign forest supply to mills. In the second phase, they solve seven daily LTSPs where the demand is expressed as a set of trips between forest sites and mills obtained from the assignment phase. As the resulting transportation plans are vulnerable to unforeseen events, the following mathematical model presents the results of adapting this work to a real-time context. For example, as the demand and supply may vary over time, we reintroduce supply constraints and disaggregate the demand by products in the daily LTSP. The demand and supply are expressed in full truckloads since the fleet is homogeneous and the supply is quite large in the case studies we consider.

Some unforeseen events can have severe impacts on the supply chain and prevent the demand satisfaction. A penalty cost for each unmet demand is incurred. The penalty cost is chosen large enough to ensure demand satisfaction whenever it is possible.

As the input data and the time-space network evolve over time, depending on the nature of the revealed unforeseen events, one must index all the model parameters and variables by the event category and by their occurrence time. However, for the sake of clarity and ease of reading, we omit these indices. Hereafter, we list the parameters and the variables of the model, and then introduce the model itself.

Parameters

F	:	set of forest sites,
M	:	set of mills,
V	:	set of trucks,
P	:	set of wood products,
I	:	set of time intervals,
N	:	set of vertices,
A	:	set of arcs,
$A^+(n)$	:	set of outgoing arcs from vertex n,
$A^-(n)$	:	set of incoming arcs into vertex n,
A_{fmp}^{loaded}	:	set of loaded driven arcs from forest site f to mill m transporting wood product p,
A^{WLE}	:	set of waiting, loaded and empty driven arcs,
$Start_v$	:	start vertex for truck v,
End_v	:	end vertex for truck v,
A_{mi}^U	:	unloading arc at mill m at time interval i,
A_{fi}^L	:	loading arc at forest site f at time interval i,
c_a	:	cost associated with arc a,
c	:	penalty cost of unmet demand,
u_a	:	capacity of arc a,
d_{mp}	:	demand of product p at mill m,
s_{fp}	:	supply of product p at forest site f,
l_{mi}	:	number of available loaders at mill m at time interval i,
l_{fi}	:	number of available loaders at forest site f at time interval i.

Variables

x_a : number of trucks that follow arc a,

δ_{mp} : unmet demand of product p at mill m.

Model

$$
\begin{aligned}
\text{Min} \sum_{a\in A} c_a x_a &+ \sum_{m\in M}\sum_{p\in P} c\,\delta_{mp} && (4.6)\\
\sum_{a\in A^{+}(Start_v)} x_a &= \sum_{a\in A^{-}(End_v)} x_a,\ \forall v\in V && (4.7)\\
\sum_{a\in A^{+}(n)} x_a &= \sum_{a\in A^{-}(n)} x_a,\ \forall n\in N\setminus \cup_{v\in V}\{(Start_v, End_v)\} && (4.8)\\
\sum_{f\in F}\sum_{a\in A^{loaded}_{fmp}} x_a + \delta_{mp} &= d_{mp},\ \forall m\in M, \forall p\in P && (4.9)\\
\sum_{m\in M}\sum_{a\in A^{loaded}_{fmp}} x_a &\leq s_{fp},\ \forall f\in F, \forall p\in P && (4.10)\\
\sum_{a\in A^{+}(Start_v)} x_a &\leq 1,\ \forall v\in V && (4.11)\\
x_a &\in \{0,1\},\ \forall a\in A^{+}(Start_v)\cup A^{-}(End_v) && (4.12)\\
x_a &\in \{0,\dots,l_{mi}\},\ \forall m\in M, \forall i\in I, \forall a\in A^{U}_{mi} && (4.13)\\
x_a &\in \{0,\dots,l_{fi}\},\ \forall f\in F, \forall i\in I, \forall a\in A^{L}_{fi} && (4.14)\\
x_a &\in \{0,\dots,u_a\},\ \forall a\in A^{WLE} && (4.15)\\
\delta_{mp} &\in \{0,\dots,d_{mp}\},\ \forall m\in M, \forall p\in P && (4.16)
\end{aligned}
$$

The objective function (4.6) minimizes the total cost, including waiting, loading and unloading, and loaded and empty driven trips. The total cost includes also the penalty costs of the unmet demand. Constraints (4.7) ensure that every used truck goes back to its base. Constraints (4.8) are flow conservation constraints for each mill and forest site vertex. Constraints (4.9) and (4.10) guarantee the satisfaction of the remaining demand while not exceeding the supply. Constraints (4.11) ensure that each truck uses, at most, one start arc. Constraints (4.12) ensure the unicity of the capacity of start and end arcs. Constraints (4.13) and (4.14) ensure that each loader only serves one truck at a time. Constraints (4.15) limit the capacity of waiting, loaded and empty driven arcs to the number of available trucks. Finally, constraints (4.16) ensure the non-negativity of the unmet demand and limits its value to the actual demand.

We assume that we have a weekly transportation plan as the starting point. The transportation operations follow this schedule until an unforeseen event is revealed. The time-space network and the input parameters are updated according the nature of the unforeseen event, then we solve the model for the current day. The new transportation plan is used until another unforeseen event is revealed and the same operation is repeated until the end of the planning horizon.

4.4 Computational results

FPInnovations, a non-profit forest research centre dedicated to the improvement of the Canadian forest industry through innovation, provided us with six case studies from Canadian forest companies. All these case studies represent weekly planning problems. Moreover, we developed a disruptions generator that produces several “weeks" of unforeseen events. A week of unforeseen events is a set of disruptions scattered over one week. The goal is to assess the proposed approach performance on different forest supply chain configurations under different disruption scenarios. The main performance indicators considered in this paper are demand satisfaction, transportation cost and computational time.

4.4.1 Unforeseen events

Unforeseen events have different impacts on the transportation network. For testing purposes, these events and their impacts are randomly generated. We developed a discrete-event model that produces a succession of events that happen at different discrete times. Note that different events are allowed to happen at the same time. The aim of this simulation is to generate unforeseen events that may happen during a full week. Therefore, after running the simulation model several times, we obtain different types of weeks with regard to the severity of the impacts. A hard week, for example, may be considered as a spring week with thawing soils, traffic jams and increasing risk of accidents because of the opening of the fishing season.

Some assumptions regarding the probability distributions of the disruptions and their impacts were made. To represent the impacts of these events, one needs to have an estimate of the expected time of the return to normal operations. It is common for the impacts to last for a shorter time and only a smaller amount of the impacts lasts for a longer time. We use then an exponential distribution to generate the disruptions duration. Note that the impacts of some unforeseen events are not measured in time units such as changes in the demand but the same observation could be applied to the demand variation volumes. As for the disruptions occurrence time, we assume that they can occur at any time in the week. Therefore, we use the uniform distribution to generate their occurrence time. We make also some assumptions about the maximum number of events that can happen simultaneously. This is done for each single unforeseen event category presented in Section 4.3.2 and also for the total number of all the event categories. During the events generation, if an unforeseen event is generated and the maximum number of simultaneous disruptions is attained, this event is rejected. Consequently, we need to keep track of the start and the end of the unforeseen events and to maintain a list of the current events. To generate the sequence of disruptions, we represent each disruption category by a special data type in our program that memorizes the occurrence time and duration of the disruption. For each disruption, we consider two types of simulation events : *Start* and *End*. The role of these events is to update the state of the simulation given that a disruption starts or ends. This includes generating the necessary random variables and scheduling future events as follows:

Event 1 *Start*

if the maximum number of simultaneous events is not attained

Generate the current disruption random duration d

Schedule the end of the event in d time units

else

Reject the event

end if

Generate a random occurrence time t

Schedule the future disruption at time t

Update the number of current events and the statistics.

Event 2 *End*

Update the number of current events and the statistics.

To start the simulation, we schedule a dummy first *Start* event at the beginning of the planning horizon. We also schedule an end-of-simulation event at the planning horizon end to stop the simulation and extract the statistics. This simulation was done using SSJ, a framework for Stochastic Simulation in Java [71] .

4.4.2 Case studies

The collaboration with FPInnovations allowed us to obtain realistic data about the forest supply chain and to validate the proposed methods. We were provided with six weekly planning problems. We assume that these problems are initially solved using an optimization method rather than manually by a dispatcher. For testing purposes, we use the method described in [38] to derive a weekly transportation plan. In these case studies, the number of initially available trucks is provided. However, the optimization method may pick only a subset of these trucks to transport the wood products. Table 4.I describes the six case studies that we denote C1 through C6. For each case study, we provide the number of wood mills ($|M|$), forest sites ($|F|$), wood products ($|P|$), the total demand (D) in full truckloads, the number of initially available trucks ($|V|$) and the number of trucks used in the resulting transportation plan ($|V_u|$). The approximate driving cost (c^D) is around 100$ per hour in average and the average waiting cost (c^W) is about 75$ per hour. The difference between loaded and empty driving costs is captured in the duration of these trips. The trip duration between forest sites and wood mills ranges from 1 to 6 hours in the 5 case studies. The loading and unloading times (t^{LU}) depend on the used equipment and the nature of the wood products. They are estimated at 20 or 30 minutes for these case studies. Therefore, we use 20 or 30 minutes steps to discretize the planning horizon.

To assess our approach, we performed complete information tests on the case studies and compared the results to our real-time re-optimization approach. We refer to complete information tests as settings where we assume we know all the unforeseen events in advance and we run the optimization method on the case studies taking into account these disruptions. In contrast, as we progress through the planning horizon and each time an unforeseen event is revealed, our real-time re-optimization approach produces a new transportation plan. This plan is used until the next disruption. Although the complete information setting is expected to outperform our approach because it takes into account all the disruptions in advance, we are nevertheless able to demonstrate the effectiveness of our real-time approach, as we show next.

4.4.3 Experimental results

We implemented the algorithms in C++, and used Gurobi 6.0 with default settings to solve the mathematical programming model. All experimentation was done on an Intel

Case study	$\|M\|$	$\|F\|$	$\|P\|$	D	$\|V\|$	$\|V_u\|$	$\|c^D\|$	$\|c^W\|$	$\|t^{LU}\|$
C1	5	6	3	618	26	11	90	75	30
C2	5	6	3	400	13	8	90	75	30
C3	1	5	1	200	37	7	110	100	20
C4	1	5	1	215	10	8	90	75	20
C5	1	5	1	215	8	8	90	75	20
C6	4	59	12	273	40	11	90	60	20

Table 4.I – Description of case studies

Core i7, 2.2GHz processor with 16 GB of memory. We used the disruptions generator to derive several "weeks" of unforeseen events. We then picked the 10th, 50th, 75th and 90th percentiles of these weeks. The lowest percentile, for instance, consists of a week with events happening at the beginning of the day and having the lowest impacts among the generated weeks. In contrast, the highest percentile means that the events occur close to the end of the days and have hard impacts. We also combined weeks with early occurrences and hard impacts, and vice-versa. Note that a different set of weeks is generated for each case study. The first part of Table 4.II describes 8 weeks (W) that we picked for each case study. For each week, we provide the number of additional demand (*DM*) in full truckloads, the number of loader breakdowns (*LO*), the number of closures (*CL*) and the number of delays (*DL*). Some weeks may have the same number of disruptions but their occurrence times are different, which explains the differences in performance.

For each of these weeks, we first transform the weekly time-space network according to the generated events. We then solve the problem for the whole week. This is the complete information test. The second part of Table 4.II presents the results of these tests. All the instances were solved to optimality. We first report the number of additional trucks (AT) used in the optimal solution compared to the initial transportation plan without any disruption. The usage of an additional truck implies a fixed cost so the model tries to minimize the number of used trucks. This allows to use the under-utilized trucks rather than using additional trucks. However, the model prioritizes the demand satisfaction since a higher penalty is incurred in the event of default. We report the unmet demand (UD) under these disruptions. In fact, in some cases, even if the disruptions are known in advance, nothing can be done to satisfy all the demand within the planning time. This includes, for example, the case where a product is available at only a set of forest sites that are closed by an unforeseen event or the case where the unloading equipment at a mill is broken for a long time. The results for case study $C6$ show an example of this behaviour.

Disruptions					Complete information		Real-time		Deviation	
W	DM	LO	CL	DL	AT	UD	AT	UD	Co	De
C1										
1	5	3	2	13	0	0.48%	1	0.48%	0.00%	0.00%
2	20	6	6	17	0	1.72%	0	1.72%	0.00%	0.00%
3	22	7	7	20	0	2.03%	0	2.19%	-0.15%	0.16%
4	31	9	7	23	0	0.92%	1	2.47%	-1.06%	1.54%
5	6	3	3	13	0	0.00%	0	0.00%	0.00%	0.00%
6	21	6	3	17	1	0.47%	1	0.47%	0.25%	0.00%
7	14	7	5	20	0	1.58%	0	1.58%	0.00%	0.00%
8	26	9	9	23	0	1.86%	1	2.48%	-0.33%	0.62%
C2										
1	5	1	1	10	0	0.74%	0	0.74%	0.00%	0.00%
2	12	5	3	15	0	2.18%	0	2.18%	0.00%	0.00%
3	19	5	4	15	0	2.39%	1	2.63%	0.20%	0.24%
4	15	8	6	18	0	0.00%	1	0.00%	0.00%	0.00%
5	7	2	2	10	0	0.00%	0	0.25%	-0.27%	0.25%
6	10	4	4	15	0	0.00%	0	0.00%	0.13%	0.00%
7	14	6	4	19	0	0.48%	0	0.48%	0.00%	0.00%
8	19	8	7	20	0	0.95%	1	1.19%	0.13%	0.24%
C3										
1	7	2	1	8	1	0.00%	2	0.00%	1.05%	0.00%
2	14	5	5	13	0	0.00%	1	0.00%	0.52%	0.00%
3	17	6	5	16	1	0.00%	3	0.92%	-0.40%	0.92%
4	19	8	6	19	2	0.00%	2	0.00%	0.49%	0.00%
5	9	2	2	9	3	0.00%	4	0.96%	-1.93%	0.96%
6	11	5	6	14	0	0.00%	1	0.00%	0.00%	0.00%
7	17	5	6	16	1	0.00%	2	0.00%	0.50%	0.00%
8	23	9	7	20	2	0.00%	8	2.69%	-3.08%	2.69%

Disruptions					Complete information		Real-time		Deviation	
W	DM	LO	CL	DL	AT	UD	AT	UD	Co	De
C4										
1	7	2	2	10	0	0.00%	1	0.00%	0.00%	0.00%
2	14	5	5	15	0	0.00%	1	1.75%	-1.86%	1.75%
3	24	6	6	17	0	0.00%	2	1.67%	-1.74%	1.67%
4	19	8	7	20	0	0.00%	1	0.85%	-0.91%	0.85%
5	9	2	2	10	0	0.00%	2	0.00%	0.00%	0.00%
6	11	5	5	15	0	0.00%	2	0.00%	0.00%	0.00%
7	17	6	6	17	0	0.00%	0	0.00%	0.00%	0.00%
8	19	8	7	20	0	0.00%	1	2.99%	-3.01%	2.99%
C5										
1	7	2	2	10	0	0.00%	0	0.00%	0.00%	0.00%
2	14	5	5	15	0	0.00%	0	2.62%	-2.80%	2.62%
3	24	6	6	17	0	0.00%	0	2.51%	-2.63%	2.51%
4	19	8	7	20	0	0.00%	0	2.56%	-2.73%	2.56%
5	9	2	2	10	0	0.00%	0	1.34%	-1.43%	1.34%
6	11	5	5	15	0	0.00%	0	0.88%	-0.95%	0.88%
7	17	6	6	17	0	0.00%	0	0.43%	-0.46%	0.43%
8	19	8	7	20	0	0.00%	0	3.42%	-3.64%	3.42%
C6										
1	5	2	2	8	0	0.72%	0	0.72%	0.00%	0.00%
2	12	5	4	16	0	1.05%	6	1.05%	0.74%	0.00%
3	19	7	4	17	0	4.79%	0	4.79%	0.00%	0.00%
4	15	9	5	20	0	5.21%	0	5.21%	0.00%	0.00%
5	7	2	1	10	0	2.50%	0	2.50%	0.00%	0.00%
6	10	5	3	15	0	3.53%	0	3.53%	0.00%	0.00%
7	14	8	6	18	0	4.88%	0	4.88%	0.00%	0.00%
8	19	10	7	21	0	5.14%	0	5.14%	0.06%	0.00%

Table 4.II – Results on case studies

The third part of Table 4.II presents the results of the proposed real-time approach where the model is solved every time an unforeseen event is revealed. The model is solved for a planning horizon starting at the event occurrence time and ending at the the current day end. For case studies $C1$ through $C5$, an optimal solution was found within 1 minute. Case $C6$ is larger and was solved to optimality within 10 seconds to 5 minutes depending on the nature of the events. We report the number of additional trucks used by our approach compared to the initial transportation plan and the unmet demand. The fourth part of Table 4.II represents the deviation in transportation cost (Co) and unmet demand (De) compared to the complete information test. This cost does not include both the fixed cost for using trucks and the unmet demand penalty. Negative values of cost deviation do not mean that the real-time approach does better than the complete information approach. It only means that the real-time model was unable to satisfy as much demand as in the complete information setting. This happens generally when the request of additional volumes is revealed close to the end of the day. Knowing in advance this information, the complete information approach manages to satisfy the demand. In contrast, the real-time approach does not have enough time to satisfy this late revealed demand. The unmet demand deviation is computed as the difference between the two approaches resulting unmet demand divided by the total demand. This includes both the initial demand and the new requests revealed during the week.

Although the complete information benefits from an information advantage, the real-time approach offers the same performance in about 50% of the cases. Only, one must note that in some cases, even though the unmet demand and cost deviation are equal for both approaches, the number of used trucks might be unequal. If a truck undergoes a breakdown or a lot of delay, the first approach, knowing this information in advance, picks another truck instead beforehand. In contrast, the real-time approach uses this truck until these events are revealed and decides then to add an additional truck as a replacement. The routes produced by the two approaches are the same, but they are not operated by the same trucks.

The case study $C5$ is the same as $C4$ under the same disruptions scenarios. The only difference is that no additional truck is allowed in $C5$. The results show that the real-time approach yields an average difference between the two cases of 0.81% for the unmet demand and -0.89% for transportation cost. Since the main goal is to satisfy the demand, adding a truck is the best option for this context. Also, for these two case studies, one may notice that the deviations in costs are approximately proportional to the unmet demand deviations. This is due to the configuration of these case studies. In fact, we have one product and one mill and the distances between the forest sites and the mill are approximately similar. Therefore, the cost of transportation is approximately proportional to the number of demand that is satisfied.

For the cases where the performances of the two approaches are different, the results are still of a very good quality since the cost deviation is under 3.64% and the demand deviation in under 3.42%. This includes three extreme cases: $C3$, $C4$ and $C5$ under week 8. In these cases, around 5% of the total initial demand was added at the end of the days. This naturally explains the performance difference between the two approaches. Another extreme case is $C6$. The results for $C6$ show an example where difficulties are met to satisfy the demand. With an equal performance with regards to demand satisfaction, the complete information approach outperforms our approach by 0.06 and 0.75% for two scenarios and has the same results for the 6 remaining scenarios. This shows that the proposed approach results deviate slightly from the ideal setting where all the information about the disruptions is known in advance, and therefore demonstrates effectiveness of the real-time approach.

4.5 Conclusion

We have introduced a new approach to re-optimize the log-truck transportation plans in real-time when an unforeseen event is revealed. This approach uses a time-space network to represent the evolution of the transportation network over time and the changes it undergoes following a disruption. The allowed trips and loading and unloading operations are used as an input for the mathematical model. The latter is solved to obtain a new transportation plan. Ease of deployment of this new plan is taken into account through ensuring the continuity of trips that are in progress when the disruption is revealed unless they are directly impacted by the disruption. A simulation model was developed to generate the unforeseen events for real applications provided by FPInnovations. Compared to a complete information scenario where disruptions are assumed to be known in advance, the proposed approach produces very good results. Also, the mathematical model was solved in a few seconds and is thus well suited for a real-time context.

Future work involves using a heterogeneous fleet of trucks. The presence of trucks with a loader onboard may give more recourse strategies especially when facing loader breakdowns at forest sites or mills. The approach proposed in this paper could be adapted to this context. The time-space network could be used to represent the disruptions impacts on the forest supply chain. However, since the trucks may have different capacities and loading constraints, one must duplicate the arcs for each truck class. This will increase the size of the problem. In this context, column generation could be used for solving this problem.

CHAPTER 5

ARTICLE 3 : A HEURISTIC BRANCH-AND-PRICE ALGORITHM TO SOLVE REAL-TIME TRANSPORTATION PROBLEMS IN FORESTRY

Chapter notes

This chapter has been submitted for publication in *European Journal of Operational Research*. Preliminary work was presented at:

- 9th Triennial Symposium on Transportation Analysis (TRISTAN IX), Oranjestad, Aruba, June 13-17, 2016
- 58th CORS Annual Conference, Banff, Canada, May 30 - June 1, 2016
- VCO Webinar, Montreal, Canada, October 14, 2015
- VCO Workshop, Montreal, Canada, February 27, 2015

Please note that this paper uses a notation slightly different than the previous one.

A HEURISTIC BRANCH-AND-PRICE ALGORITHM TO SOLVE REAL-TIME TRANSPORTATION PROBLEMS IN FORESTRY

Amine Amrouss [1,2,*], Nizar El Hachemi [1], Michel Gendreau [1,3], Bernard Gendron [1,2]

[1] *Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation (CIRRELT), Canada*

[2] *Department of Computer Science and Operations Research, Université de Montréal, P.O. Box 6128, Station Centre-Ville, Montreal, Canada H3C 3J7*

[3] *Department of Mathematics and Industrial Engineering, Polytechnique Montréal, C.P. 6079, succursale Centre-ville, Montreal, Canada H3C 3A7*

In the forest industry, the transportation of wood from forest sites to mills is generally done according to a predefined transportation plan. However, several unforeseen events may occur and disrupt the planned trips. The dispatcher is then required to derive an alternative transportation plan in real-time. In this paper, we present an approach to derive new routes for the impacted trucks in a short computational time. We use a time-space network to represent the impact of unforeseen events on the forest supply chain and to generate routes that take into account truck drivers' constraints. The linking constraints, such as the synchronization with the loaders and the satisfaction of the demand, are managed in a master problem. The latter is solved using column generation and a heuristic branch-and-price algorithm is developed to find integer solutions. Several disruption scenarios were generated to test the proposed approach on case studies from the Canadian forest industry. Numerical results are presented.

Keywords: Real-time, transportation, forestry, mathematical programming, column generation.

5.1 Introduction

According to the U.N. Food and Agriculture Organization (FAO) [46], different management plans were developed for about 52% of the world forest area in 2010. This includes both plans intended to wildlife conservation and to wood production. The operational research contributions to derive such plans date back to the 1960s [121] and cover a broad spectrum of activities such as planting trees, managing forest fires, harvesting, transportation, etc. In this paper, we present a real-time planning problem in forestry where we propose alternative truck schedules to dispatchers in response to disruptions in timber transportation. In this problem, the disruptions are caused by unforeseen events such as delays, breakdowns or demand variations. They perturb the planned transportation operations and require quick modifications to minimize their impact. This problem was identified as one of 33 open challenges for operational research in forestry [109].

Transportation plans in forestry are derived for strategic, tactical, operational and real-time levels [107]. For the strategic level, the planning is done for a period of more than 5 years and concerns mainly road building and upgrading. For the tactical level, the planning deals with the wood flow assignment between forest sites and mills. This wood assignment is tightly linked to the harvest planning. This is why tactical transportation planning is typically included in harvest planning. The planning horizon is generally between 6 months and 5 years. For the operational level, the aim is to define detailed plans for the fleet of trucks that minimize the total cost for a period ranging from 1 day to 6 months. This involves minimizing the number of empty driven trips, which depend on the configuration of the supply chain. In fact, there must be a flow of wood going in opposite geographical directions to enable backhauling. We refer the reader to [25] for more details on backhauling in forestry. Deriving truck schedules in forestry is known in the literature as the log-truck scheduling problem (LTSP). The real-time application is also a LTSP where the transportation plans are disrupted by unforeseen events that must be dealt with in a short amount of time.

The LTSP is a vehicle routing problem (VRP) with pickup and delivery, time windows and split deliveries [14, 53]. We mean by split deliveries that several visits are necessary to satisfy the wood mills demand. This variation of VRP was introduced in [34] and reviewed in [6, 7]. In forestry, trucks must also be synchronized with loaders at forest sites and mills for loading and unloading operations [19, 40, 105]. This allows avoiding queues at the loaders and thus reducing the cost of unproductive activities. Different types of synchronization constraints, such as movement synchronization, transshipment and so on, were addressed in the VRP literature [32].

For the real-time level, there are two types of VRP applications. The first one deals with truck dispatching where the main aim is to assign a trip to idle trucks one at a time [110]. There is no need to derive the entire daily or weekly transportation plan. This type of approach is adapted to applications such as cabs and ambulances dispatching where there is insufficient information about future requests. In contrast, in forestry one must take advantage of the available data on forest supply chain including demand and supply. The second type of applications, which we consider in this work, aims at generating new weekly plans in real-time. It starts by solving the LTSP for one week and operating the trucks according to the resulting plan. When an unforeseen event is revealed and given more accurate data on the forest supply chain, a new transportation plan is generated for the rest of the week.

The recourse strategies used for solving real-time transportation problems in forestry depend on the context of application. In [3], weekly transportation and loaders plans are derived a priori. When an unforeseen event occurs, some trips are delayed, while other trips are cancelled. First, the transportation and loaders plans are updated according to the delays. Then, the cancelled trips are rescheduled in such a way that trucks arrive at loaders when the latter are free. This avoids impacting the schedules of the other trucks. Different options are used to obtain such solutions, including opening hours extensions and using additional trucks. The goal of this application is to minimize changes in the initial transportation plan, since agreements are made with truck drivers according to this plan. In contrast, in [2], the authors study a context where drivers are provided with one trip at a time and, even if weekly transportation plans are available, the demand is stated on a daily basis. Therefore, the problem is solved only for the rest of the day after each disruption. The transportation plan is completely reoptimized for the current day in contrast to [3] where the plan is only adjusted. In both works, the fleet of trucks is homogeneous. The use of a heterogeneous fleet, as in this paper, adds to the complexity of the problem. We also assume a different context, where we reoptimize the weekly transportation plan, but by limiting the scope of changes made to the initial plan.

In the current practice, modifications to the transportation plans are usually done manually. However, given the large number of constraints that must be met, this manual approach does not always produce an efficient solution. There are commercial decision support systems that include real-time log-truck dispatching modules, but they are not documented and there is no indication whether an optimization model is used or not. The reader is referred to [9], an extensive survey of the systems used in timber transportation for different planning horizons. In this paper, we propose a heuristic branch-and-price approach to provide dispatchers with an alternative transportation plan in real-time. This plan aims at minimizing costs, but also changes to the original transportation plan.

This paper is organized as follows. We introduce some notation and describe the real-time transportation problem in forestry in Section 5.2. In Section 5.3, we describe the proposed optimization approach to react to unforeseen events. Then, we present computational experiments that demonstrate the efficiency of the approach in Section 5.4. Conclusions are reported in Section 5.5.

5.2 Problem description

We assume that the wood is transported from forest sites to mills according to a predetermined weekly transportation plan. We define F, M, and T as the sets of forest sites, wood mills and heterogeneous trucks, respectively. The transportation plan is composed of a set of weekly routes that are assigned to each truck $t \in T$ from a set of feasible routes R_t. The trucks differ in their capacities and their types. Some trucks have a loader on board and therefore do not have to synchronize with loaders at forest sites and mills. Their number is limited because the onboard loader decreases significantly their capacity and thus increases their cost. However, they remain useful in forest sites with small amounts of supply as they avoid installing a loader there. We refer to the other trucks as regular trucks. There is also a ranking of truck drivers that allows the highest priority drivers to get the largest workloads. Some high priority drivers may also have preferences concerning the mills and forest sites to visit. Each truck is associated with a mill that is considered as its base. The truck must start and end its shift at its base. Also, the maximum driving duration of a shift and the minimum rest periods between successive shifts must comply with regulatory limits.

A route $r \in R = \bigcup_{t \in T} R_t$ consists of a set of working shifts that are separated by rest periods. A shift is a sequence of empty and full driven trips, loading and unloading operations and waiting between operations. The time is discretized into a set of equal intervals I. Although the loading and unloading durations may depend on the truck capacity and on the transported product, we assume that they are approximately equal. This duration is used as the discretization step of I. The durations of transportation operations are expressed as a multiple of these intervals. Each operation has a different cost per unit of time and the cost of a route $r \in R$ is the sum of the unit costs of the operations associated with r multiplied by their durations.

To define a feasible route, we consider a time-space network $G = (N, A)$ (Figure 5.1), where N is the node set and A is the arc set. The node set contains eight types of nodes. Each truck has a source and a sink that represent the start and the end of its week, respectively. Each forest node is represented by two types of nodes : empty and full nodes. An empty node represents the fact that the truck has just arrived at this forest site. A full node means that the truck is now loaded. These two types of nodes are replicated at every time interval $i \in I$. Each mill node is also replicated similarly, but two other types of nodes are added : empty and full rest nodes. They represent the fact that the truck driver is resting between shifts. Note that the truck can access the rest nodes even when it is loaded. However, this rest node must correspond to its base. In this case, the truck waits for the mills to open to be unloaded the following day.

The arc set contains twelve types of arcs. Start arcs connect a source node to a shift node. Loading arcs connect an empty forest node to a full forest node. As for loaded driven arcs, they connect a full forest site to a full mill node. These arcs are duplicated for each product and for each truck since they have different capacities, which also depend on the type of product. Waiting arcs connect successive full or empty mill nodes. We choose to make the trucks wait at the mills rather than at the forest sites to avoid symmetry, because the fuel stations are generally close to the mills and because the number of mills is generally less than the number of forest sites. Unloading arcs connect a full mill node to an empty mill node. Empty driven arcs connect an empty mill node to an empty forest site node, while end arcs connect an empty mill node to the sink node. End of shift arcs can be loaded and then connect a full mill node to a full rest node, or unloaded and then connecting empty mill nodes to empty rest nodes. Start of new shift arcs can similarly be loaded or unloaded. Their length is at least equal to the minimum rest time. Finally, rest arcs connect successive empty or full rest nodes.

A feasible route for a truck can be expressed as a path in this network starting at a source node and ending at a sink node. Given the set of feasible routes for each truck, one must decide which route must be assigned to which truck. In fact, there are some constraints that link the routes of these trucks. The demand must be satisfied while the supply must not be exceeded. The demand and supply for a set of commodities K are

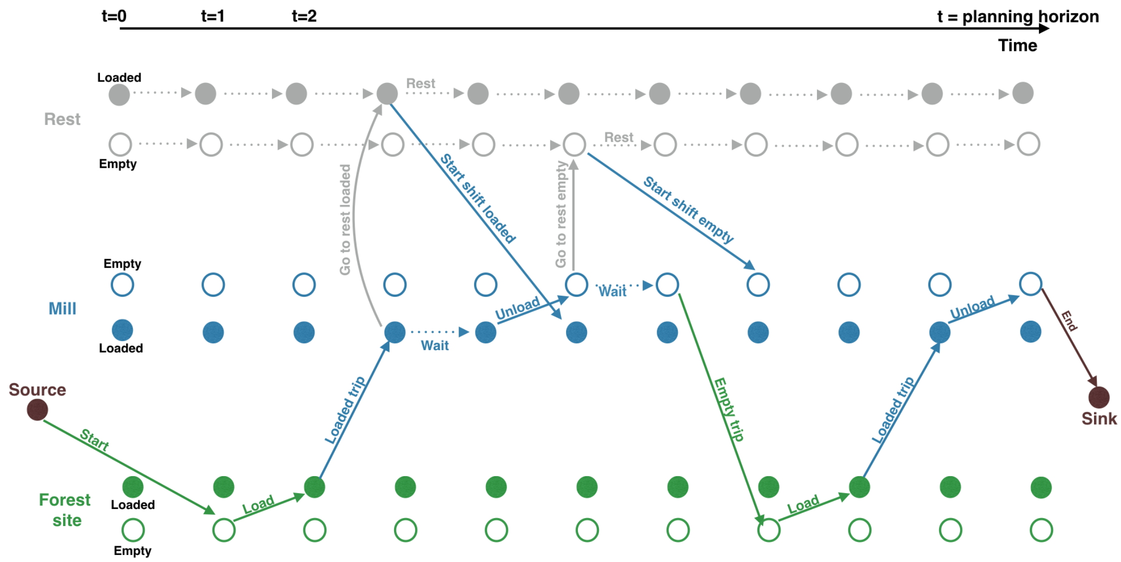

Figure 5.1 – Time-space network

often expressed in terms of full truckloads in the literature, but we use cubic meters instead, since the fleet is heterogeneous. However, we assume that the available supply is large enough to fully load each truck at each trip. Let d_{mk} and s_{fk} be the demand of commodity k at mill m and the quantity of commodity k available at forest site f, respectively. If the demand could not be wholly satisfied during the given planning horizon, a penalty is incurred.

Another linking constraint is the fact that the number of trucks being loaded or unloaded at a node at a certain time interval must be lower than the number of available loaders. Note that we do not count trucks with onboard loaders. For each truck $t \in T$, we define the binary constant l_t that takes value 1 if truck t needs a loader. Let also n_{fi} and n_{mi} be the number of loaders available over time interval i at forest site f and mill m, respectively.

For each route $r \in R$, we define binary constants L_{rfi} (U_{rmi}) that take value 1 if route r includes a loading (unloading) operation at forest site f (at mill m) over time interval

i. Also, q_{fmkr} is defined as the quantity of commodity k that is transported from forest site f to mill m on route r. The first set of decision variables δ_{mk} represents the quantity of unsatisfied demand of commodity k at mill m, while p denotes the unit penalty incurred in this case. This penalty is chosen large enough to ensure demand satisfaction whenever possible. Finally, we define the binary variables y_r that take value 1 if route r is selected, c_r denoting the cost of that route.

In [105], a mathematical model is proposed for tactical transportation problems. The authors consider also other aspects such as inventory management, wood freshness and balancing schedules in terms of workload between different weeks in the planning horizon. These constraints are not relevant to the problem that we study in this paper. The following model, denoted (P), is an adaptation of the one proposed in [105] that allows to define the weekly schedules for each truck :

$$\begin{aligned}
\text{Min} \sum_{r \in R} c_r y_r &+ \sum_{m \in M} \sum_{k \in K} p \delta_{mk} && (5.1) \\
\sum_{f \in F} \sum_{r \in R} q_{fmkr} y_r + \delta_{mk} &= d_{mk},\ \forall m \in M, k \in K && (5.2) \\
\sum_{m \in M} \sum_{r \in R} q_{fmkr} y_r &\leq s_{fk},\ \forall f \in F, k \in K && (5.3) \\
\sum_{r \in R_t} y_r &\leq 1,\ \forall t \in T && (5.4) \\
\sum_{t \in T} l_t \sum_{r \in R} U_{rmi} y_r &\leq n_{mi},\ \forall m \in M, i \in I && (5.5) \\
\sum_{t \in T} l_t \sum_{r \in R} L_{rfi} y_r &\leq n_{fi},\ \forall f \in F, i \in I && (5.6) \\
y_r &\in \{0,1\},\ \forall r \in R && (5.7) \\
\delta_{mk} &\in R^+,\ \forall m \in M, k \in K && (5.8)
\end{aligned}$$

The objective function (5.1) minimizes the sum of routing costs and unsatisfied demand penalties. Constraints (5.2) and (5.3) aim at satisfying the demand in compliance with the available supply. Constraints (5.4) ensure that a truck is assigned at most one single route. Note that not all the available trucks have to be used. Constraints (5.5) and

(5.6) guarantee that each available loader serves only one regular truck in a given time interval.

Column generation methods [98, 104, 105] are used to solve this problem, producing a daily or weekly transportation plan. This plan can be perturbed by many unforeseen events that impact the forest supply chain including mills, forest sites, transportation network, and trucks. The consequences of these events can be a road closure, a mill or a forest site closure, a loader breakdown, a truck breakdown, a delay, or a variation in demand. Note that these consequences can be caused by different events. For example, a delay can be caused by a forest road degradation, traffic jams, bad weather, etc. These events may prevent the completion of the planned operations and the dispatcher is then required to propose alternative routes to the truck drivers in a limited time. In the following section, we propose an approach to manage the disruptions and to generate new cost-efficient routes in a short computational time.

5.3 Optimization approach

We propose an approach to derive a new transportation plan each time a disruption occurs. First, we build a time-space network composed of the nodes and the arcs that are still available after the disruption. We also set the new values of the parameters of the mathematical model (P). The integer programming model is solved using a branching tree. At each node, the linear relaxation of the model is solved by dynamically generating columns with the potential to improve current solutions. Then, some rules are used to fix a fractional variable to a binary value, which yields a new node in the branching tree. The process is repeated until a stopping criterion is reached.

5.3.1 Reacting to an unforeseen event

Several unforeseen events can occur while wood is being transported from forest sites to mills. They can impact the mills, the forest sites, the trucks or the transportation network. A list of the most frequent disruptions in the forest industry can be found in

[3]. Our approach is based on a mathematical model that uses a time-space network (Figure 5.1) to represent the impact of the unforeseen events on the forest supply chain. If a road is closed, the arcs representing this road are removed for all the closure duration. If a truck has one time unit of delay, this can be represented by adding an arc linking the destination node to the same destination, but in the next time interval. This arc can be used only by the impacted truck. In fact, many events have similar impacts on the network from the modelling perspective. For example, bad weather conditions and traffic jams both cause delay. Also, certain forest sites are only accessible by one forest road. Site or road closure have the same impact in this case. Therefore, one should focus on the impact rather than the event itself. In [2], these impacts are classified into four groups : closures, delays, loader breakdowns, and demand and supply variations. Also, modifications made to the time-space network are presented for each group. Although the time-space network (Figure 5.1) used in this work is enhanced to take into account rest periods, these modifications remain relevant.

Once the modifications are done on the time-space network, we are able to generate new feasible routes given the current state of the supply chain. These routes are then used as input to the mathematical model (P). Other parameters of (P) are also used to capture the impact of unforeseen events. For instance, the numbers of available loaders n_{fi} and n_{mi} could be decreased for all the intervals when a loader is under repair. We remove the corresponding arcs in the time-space network, too. Doing so, we avoid generating routes that are not feasible for the model. To summarize, only the inputs of (P) are changed, not the mathematical model itself. The latter remains valid to generate a new transportation plan after each disruption occurrence.

5.3.2 Column generation

The LTSP as formulated above is solved using column generation. A column in the model corresponds to a feasible route and is therefore easily interpretable. The set of feasible routes is generally large and the problem would become very difficult to solve if all of these routes were generated for all trucks. Therefore, the column generation approach starts with a restricted master problem where only a set of routes that form a feasible solution is considered. In this problem, the integrality constraints (5.7) are also relaxed and the columns that have the potential to improve the solution are added iteratively. These promising columns have negative reduced costs and are obtained after solving a set of subproblems. The management of the pool of columns can be done in various ways. In the case of the LTSP, [98] generate a priori, according to numerous rules, only a subset of feasible routes, which are then stored in a pool of columns. The pricing procedure is used to verify the optimality of the current solution or to add new routes from the pool to the set of active columns in the model. In [104, 105], routes with negative reduced costs are dynamically generated using a time-space network. In [104], a single column ending at each demand point is generated at each iteration. In [105], the 200 columns with the most negative reduced costs are added at each iteration. Finally, the solution of the linear relaxation is generally fractional and branching is commonly used to obtain integer solutions.

For the real-time problem that we study in this paper, we solve our model every time a disruption occurs. The planning horizon starts at the disruption occurrence time and ends at the current week end. The unforeseen events may occur in the beginning of the week and thus the number of feasible routes may remain large. However, we assume that the initial weekly transportation plan is obtained through an optimization method. Therefore, the initial routes are generally of good quality and some promising columns for the new problem could be obtained from the initial plan by removing the transportation tasks that were already performed. As for the routes that are impacted by the disruption, we try to minimize the number of deviations from the initial plan. Let n_{fmkr} be the number of trips where commodity k is transported from forest site f to mill m on route r. For an impacted truck, we define r_{init} and r_{new} as the initial route and the newly generated route, respectively. We define the *compatibility degree* as the maximum difference allowed between the number of trips to each forest site and to each mill on r_{init} and r_{new}. The commodity being transported on these trips is also taken into account.

We then impose an upper bound on the quantity:

$$\sum_{f\in F}\sum_{m\in M}\sum_{k\in K}|n_{fmkr_{init}} - n_{fmkr_{new}}|$$

Once a disruption occurs, our approach starts by solving the restricted master problem, where we first relax the integrality constraints (5.7) of the initial model (P) and consider only the subset of routes that are not impacted by the disruption. Note that some disruptions might impact all routes. When this happens, the routes that consist of sending all trucks to their sink nodes are considered. The unsatisfied demand variables ensure that a feasible solution can be found for this problem. To generate new columns, we first update A and N, the sets of available arcs and nodes of the time-space network according to the unforeseen event impact. This network is used to generate feasible routes that respect the compatibility degree restriction. Note that there are many other restrictions on the feasible routes. Some trucks may not be allowed to use all the roads or to visit some nodes. The closing and opening times of the wood mills can be different. All these constraints are implicitly managed during the construction of the time-space network and the generation of routes.

When an unforeseen event is revealed, the start node for a truck represents its actual position. If the event does not prevent the truck from attaining its current destination, only arcs linking its start node to its sink node or its destination nodes at different interval times are kept. Otherwise, it is linked to all the other possible destinations. To find a feasible route, we draw a path from the start node to the sink node that respects the compatibility degree restriction.

To calculate the reduced cost of the generated routes, we use the dual variables associated with the constraints of the linear relaxation of the problem. Let α_{mk}, β_{fk}, θ_t, γ_{mi} and π_{fi} be the dual variables associated with constraints (5.2), (5.3), (5.4), (5.5) and (5.6), respectively. Recall also that a shift node in the time-space network represents a physical forest site or mill in a certain time interval. We define $i(n)$ and $s(n)$ as the time interval and physical site of node n, respectively. Also the loaded driven arcs are duplicated for each commodity and truck. Let k_{od} be the commodity transported on arc (o,d) and $C_{tk_{od}}$ the carrying capacity of this commodity of truck t. The cost of an arc (o,d) linking an origin node o to a destination node d is c_{od}. From the time-space network, we define for a certain route $r \in R$ the following sets:

$$
\begin{array}{lcl}
A_r & : & \text{set of arcs of route } r. \\
A_r^{loaded} & : & \text{set of loaded driven arcs of route } r. \\
A_r^{loading} & : & \text{set of loading arcs of route } r. \\
A_r^{unloading} & : & \text{set of unloading arcs of route } r. \\
A_r^{remaining} & = & A_r \setminus (A_r^{loaded} \cup A_r^{loading} \cup A_r^{unloading}) \\
 & : & \text{set of remaining arcs of route } r.
\end{array}
$$

We define binary variables x_{od} that take value 1 if arc (o,d) is selected in the path. The reduced cost of a feasible route is the sum of the reduced costs of the arcs of its underlying path minus the dual variable θ_t that represents the reduced cost of using a truck $t \in T$ and is calculated as follows:

$$
\begin{aligned}
& \sum_{(o,d)\in A_r^{remaining}} c_{od} x_{od} + \sum_{(o,d)\in A_r^{loaded}} (c_{od} - C_{tk_{od}}(\alpha_{s(d)k_{od}} - \beta_{s(o)k_{od}})) x_{od} \\
+ & \sum_{(o,d)\in A_r^{unloading}} (c_{od} - \gamma_{s(o)i(o)}) x_{od} + \sum_{(o,d)\in A_r^{loading}} (c_{od} - \pi_{s(o)i(o)}) x_{od} - \theta_t
\end{aligned}
$$

Finding routes with negative reduced costs can be expressed as a shortest path problem with resource constraints [65] in the time-space network. The resources include commodities, loaders availabilities, time for minimal and maximal shift and break durations, etc. For the compatibility degree restriction, we also keep track of the number of trips and the commodities being transported to extend the path from one node to another. We solve this problem using a label-setting algorithm as in [104, 105].

This algorithm allows to find the shortest path (with the smallest negative reduced cost). However, other metrics such as the path length with regards to time are also updated and used to control the extension from a node to another during the algorithm iterations. For example, if the minimal time length is not yet attained, arcs going to sink nodes are not considered. To find paths that respect the compatibility degree restriction, we extend the algorithm to capture the number of differences with the initial route.

The extended algorithm works as follows. As in any shortest path algorithm, we use a label $\{prev[n], rc[n]\}$ for each node n of the time-space network. This label stores information about the previous node ($prev[n]$) that ensures the lowest reduced cost ($rc[n]$) to attain node n at a certain iteration of the algorithm. Let rc_{od} be the reduced cost of an arc (o,d). We also include a new set of variables $cd_{fmk}[n]$ in each node label to measure the number of trips from forest site f to mill m carrying commodity k on the current path to node n. These variables are initialized at $n_{fmkr_{init}}$ for the source node and at ∞ for the other nodes. We define also constants $cd_{fmk}[od]$ to represent the contribution of arc (o,d) in satisfying the compatibility degree restriction. We set $cd_{fmk}[od] = -1$ for loaded arcs from site f to mill m carrying commodity k and $cd_{fmk}[od] = 0$ for the other arcs. We refer to the compatibility degree as CD while $\mathbb{1}(x \in X)$ is the indicator function, which takes value 1 if $x \in X$ and 0 otherwise. The algorithm is then extended as follows:

---

for each arc (o,d) following the chronological order of origin nodes o

  if $\sum_{f\in F, m\in M, k\in K} \mathbb{1}[cd_{fmk}[o] + cd_{fmk}[od] < 0](cd_{fmk}[o] + cd_{fmk}[od]) \geq -CD$

    if $\sum_{f\in F, m\in M, k\in K} cd_{fmk}[d] > \sum_{f\in F, m\in M, k\in K} (cd_{fmk}[o] + cd_{fmk}[od])$

      if $rc[d] > rc[o] + rc_{od}$ **:**

        $rc[d] = rc[o] + rc_{od}$ and $prev[d] = o$

        $cd_{fmk}[d] = cd_{fmk}[o] + cd_{fmk}[od]$

      end if

    end if

  end if

end for

We use this algorithm to derive paths with different lengths with regard to shift duration. For each $l \in \{$minimum shift duration, ..., maximum shift duration$\}$, we find the shortest path with regard to the reduced cost that has a duration l. This is achieved by duplicating the labels at each node for each duration l. These labels contain then the compatibility variables, the previous node and the reduced cost of the shortest path such that the duration of the current shift is l. The subproblem is solved for each impacted truck. Routes with negative reduced costs are added to the master problem that we solve again. This yields a new solution, as well as new dual variable values. We repeat this process until no negative-cost routes can be found. This means that an optimal solution is found for the linear relaxation. Note that, since not all the feasible routes were considered, the solution is optimal only for the subset of routes that respect the compatibility degree restriction. This process can also be stopped before obtaining the optimal solution if a time limit is reached.

The routes reduced costs are based on a dual information, and we consider the compatibility degree restrictions as a primal information, since they are based on the initial transportation plan. Another primal information that we consider is the fact that a truck that is used in the initial transportation plan must be assigned a route. Therefore, if no negative reduced costs routes are generated for this truck, we generate positive reduced cost routes that respect the compatibility degree restriction only once during the column generation process. If still no routes can be found, the compatibility degree restriction is removed.

5.3.3 Heuristic branch-and-price

To derive integer solutions from the relaxed problem, we use a heuristic branch-and-price algorithm. We integrate our column generation approach into a branch-and-bound tree where we use a heuristic branching strategy inspired by [106]. At the root node, the original relaxed restricted master problem is solved using column generation. Once an optimal solution is found or a stopping criterion such as a time limit is attained, we generally obtain a fractional solution. We focus first on route variables that are linked to the highest priority truck. This increases the likelihood of obtaining routes with the largest workloads that could become infeasible if other trucks routes were fixed before. We branch on the highest priority truck routes variables by fixing the variable with the largest fractional value to 1. We solve the restricted master problem with this additional constraint. If the problem becomes infeasible, we fix this variable to 0. Otherwise, since a route is fixed, some loaders become unavailable to the other trucks at certain time intervals. In this case, the loading and unloading arcs are removed from the time-space network for the corresponding time intervals. Also, the supply may be exhausted and the demand may be satisfied. We similarly remove the corresponding arcs and nodes from the network. This yields a reduction in the number of feasible routes generated in subsequent iterations. Also, some of the routes being considered in the current restricted master problem become infeasible because of the linking constraints, and are then removed from the model. The column generation is then applied to the resulting model at the child node. This process is repeated until all trucks have a fixed route or a time limit criterion is attained. In the latter case, we enforce the integrality constraint on the resulting reduced model and use the branch-and-bound procedure from a state-of-the-art software to solve it (we use Gurobi 6.5).

5.3.4 Summary

To summarize, we solve the model to derive a weekly plan using the initial input data and operate the trucks accordingly. When an unforeseen event is revealed, we :

- update the time-space network according to its consequences, thus defining the new feasible routes,
- update the model parameters (the remaining demand and supply and the number of available trucks and loaders),
- solve the model using a heuristic branch-and-price algorithm based on the new time-space network and model parameters.

5.4 Computational results

The proposed approach was tested on two case studies provided by FPInnovations, a research centre with the aim to improve the Canadian forest industry through innovation. In the first case study (Case 1), there are 4 wood mills, 60 forest sites, 12 commodities and 62 trucks. The total demand volume is 20,187 m^3, while the total supply is 71,023 m^3. Although the total supply is larger than the demand, the available supply differs from one site to another. In the second case study (Case 2), there are 5 wood mills, 6 forest sites and 3 commodities with a 16,800 m^3 demand, while the total supply is 3 times larger than the demand. The number of available trucks is 21. In the two cases, there is one loader per mill and per forest site, while the loading and unloading times are estimated at approximately 20 minutes in Case 1 and 30 minutes in Case 2. The average transportation cost is estimated at \$100 per hour and the waiting cost is estimated at \$71.5 per hour.

Several scenarios of occurrence of unforeseen events were developed to test and evaluate the performance of the approach presented in the previous section. We used a discrete-event simulation model, as in [2], to generate a set of unforeseen events scattered throughout the week. These events can occur in the beginning of the week and last for a short time. This gives more room to deal with their consequences in contrast to events happening at the end of the week. We generated sets with different severity levels. In Table 5.I, we use a subset of disruptions that we generated for Case 1. The proposed approach is used to generate new routes each time an unforeseen event happens. The tests are done in a continuous mode where we start with the weekly transportation plan, solve the model when the first event is revealed, continue with the obtained solution until the second disruption is revealed, resolve the model, and so on. We imposed a time limit of 5 minutes on the approach, which means that only a fraction of this time is used to generate routes with negative reduced costs for each impacted truck.

We report six events in chronological order in Table 5.I. After each event, we compare the resulting routes to the initial transportation plan before any disruption. The first three metrics measure the average difference between the proportion of empty trips, loaded driven trips and waiting time over all the trucks in the initial and new solution. The results show that the configuration of the transportation plan remains approximately the same. The slight variations can be explained by the nature of the events. The first truck delay impacts only one truck and the model picked a solution that makes it wait instead of allowing it to be loaded or unloaded at time intervals that were assigned to other trucks. This prevents the propagation of delay to these trucks. In contrast, a loader breakdown caused the cancelation of some deliveries where the trucks had to wait before being unloaded in the previous transportation plan. This explains the decrease in waiting time. The last truck delay increases again the waiting time compared to the second last solution.

Event	Empty	Loaded	Waiting	Cost	Demand	Time (s)
One-hour truck delay	-0.1%	-0.1%	0.2%	0.08%	0.00%	95
2 new demands	0.1%	0.1%	-0.2%	0.13%	0.00%	93
One hour traffic jam (incurred delay : 30 min)	0.2%	0.2%	-0.4%	0.50%	0.00%	107
2-hour loader breakdown	0.3%	0.3%	-0.6%	-0.38%	0.78%	104
6-hour truck breakdown	0.3%	0.3%	-0.6%	-0.82%	1.31%	102
2-hour truck delay	0.1%	0.1%	-0.2%	-1.07%	1.53%	96

Table 5.I – Results after six disruptions

The third and fourth metrics represent the transportation cost and the proportion of demand that could not be met. The first three events caused an increase in transportation cost, but the model was still able to satisfy the demand. For the other three events, in turn, the model was unable to satisfy a portion of the demand, which resulted in decreased transportation cost, but there may be a penalty to pay. The unsatisfied demand can be explained by the severity of the events. For a mill that is open from 6 am to 6 pm, the loader breakdown represents one sixth of its daily operational time and results in less satisfied demand. Also, these events are happening late in the week. For example, the last truck delay has more impact than the first one. Finally, the last metric represents the resolution time for the model each time an event is revealed. All the solutions were obtained within less than two minutes.

The results presented in Table 5.I were obtained using a compatibility degree of 4. To measure the impact of the latter on the solutions, we use 4 sets of unforeseen events on each case. The description of these sets is in Table 5.II where we use 4 categories of events. The first refers to loaders breakdowns, but since there is one loader per mill and per forest site, their closure can also be represented as a loader breakdown with regards to regular trucks. Also, the impact of a truck breakdown is that the truck is not available for a certain amount of time. This can be seen as a big delay in its operations. The second category contains then truck breakdowns and delays. The third one represents the number of new requests for different products and they are expressed in terms of an average full truckload. The last category contains the number of events that impact roads. Unlike the truck breakdowns, these events impact not only an individual truck but all the trucks using this road.

	Loader breakdown	Truck breakdown	New requests	Road disruptions
Case 1a	1	6	5	1
Case 1b	5	9	12	11
Case 1c	6	16	19	7
Case 1d	8	15	15	12
Case 2a	1	6	5	1
Case 2b	6	13	20	10
Case 2c	7	18	22	9
Case 2d	9	20	31	12

Table 5.II – Description of unforeseen events

We used compatibility degrees ranging from 0 to 6 on the resulting case studies. For each case, we compute the average cost per cubic meter of transported products, the proportion of demand that was unsatisfied and the time required to obtain a solution after each disruption. The average of these values over all the solutions are reported for each compatibility degree in Table 5.III.

Degree	Case 1a			Case 2a		
	Cost ($/m3)	Demand	Time (s)	Cost ($/m3)	Demand	Time (s)
0	15.56	1.84%	114	10.46	0.32%	21
1	15.69	1.50%	110	9.70	0.00%	22
2	15.59	0.87%	104	9.76	0.00%	20
3	15.86	1.10%	104	9.87	0.00%	23
4	15.86	0.60%	99	9.87	0.00%	23
5	15.83	0.80%	97	10.89	0.00%	17
6	15.96	0.63%	111	10.87	0.00%	18
	Case 1b			Case 2b		
	Cost ($/m3)	Demand	Time (s)	Cost ($/m3)	Demand	Time (s)
0	15.49	3.81%	107	9.88	0.51%	18
1	15.61	3.01%	116	10.03	0.49%	19
2	15.65	2.19%	99	10.04	0.49%	19
3	15.72	2.21%	106	10.05	0.49%	19
4	15.82	1.35%	97	10.06	0.49%	19
5	15.81	1.42%	104	10.06	0.49%	19
6	15.90	0.59%	94	10.06	0.49%	19
	Case 1c			Case 2c		
	Cost ($/m3)	Demand	Time (s)	Cost ($/m3)	Demand	Time (s)
0	15.54	2.90%	107	10.34	0.55%	26
1	15.62	1.62%	106	10.01	0.12%	28
2	15.65	2.18%	112	9.96	0.10%	22
3	15.73	2.01%	106	9.98	0.10%	19
4	15.78	1.21%	110	10.17	0.10%	20
5	15.78	1.30%	94	10.28	0.10%	23
6	15.86	2.12%	102	10.19	0.10%	20
	Case 1d			Case 2d		
	Cost ($/m3)	Demand	Time (s)	Cost ($/m3)	Demand	Time (s)
0	15.62	1.48%	85	10.90	2.49%	23
1	15.76	1.48%	84	10.90	0.02%	18
2	15.76	0.96%	83	10.26	0.00%	22
3	15.87	0.91%	106	10.23	0.00%	21
4	15.83	1.19%	92	10.36	0.00%	13
5	15.99	0.05%	103	10.60	0.00%	19
6	16.02	0.02%	94	10.49	0.00%	16

Table 5.III – Impact of compatibility degree on results

The results for Case 1 show an increase in the average cost as we increase the compatibility degree. Also, the approach tends to satisfy the demand as we increase the degree. This tendency is clearer in Case 2. Moreover, the demand satisfaction becomes stable after degree 1 or 2 in Case 2. In contrast, once the degree is enough to satisfy the demand, any further increase of degree results in more expensive solutions. This shows that there is a need to tune the compatibility degree for each case study beforehand. The right degree can be clear as in Case 2. But, in Case 1, a tradeoff between demand satisfaction and transportation cost must be found. As for the solution time, the approach produces new transportation plans in less than 30 seconds in Case 2 and 2 minutes in Case 1. This shows the compatibility of the approach with a real-time application.

We also compared the performance of the proposed approach to the method used in [105]. The method was developed for a tactical LTSP that we adapted to a weekly problem. We compare then the proposed approach to a naive approach where a weekly planning method is used to resolve the problem after each disruption. The method uses column generation until an optimal solution is obtained for the relaxed master problem or a stopping criterion is attained. The method is used only at the root node. After that, branch-and-bound is used to find an integer solution. We use a 20-minute time limit to stop the column generation even if an optimal solution is not obtained and allow another 20 minutes to branch-and-bound. We report in Table 5.IV the results obtained by this method that we refer to as CG. We average the results obtained in Table 5.III over all the compatibility degrees and compare the results against the CG ones. We refer to the heuristic branch-and-price approach as HBP.

The heuristic branch-and-price approach outperforms clearly the CG method almost on all the case studies. Only in Case 1b, the CG approach is able to find a less costly solution but with lots of unsatisfied demand. The differences in solution times are also large.

	Cost ($/m3)		Demand		Time (s)	
	HBP	CG	HBP	CG	HBP	CG
Case 1a	15.76	15.83	1.05%	11.8%	105.57	2400
Case 1b	15.71	15.59	2.08%	12.2%	103.29	2400
Case 1c	15.71	15.90	1.90%	12.9%	105.29	2400
Case 1d	15.84	16.05	0.87%	15.8%	92.43	2400
Case 2a	10.20	11.88	0.05%	4.30%	20.57	136
Case 2b	10.02	12.41	0.49%	3.74%	18.86	1556
Case 2c	10.13	12.60	0.17%	3.04%	22.57	2400
Case 2d	10.54	13.07	0.36%	1.64%	18.86	486

Table 5.IV – Approaches comparison

5.5 Conclusion

In this paper, we introduced a heuristic branch-and-price algorithm to solve the real-time transportation problem in forestry. The approach uses the initial transportation plan to generate new routes when an unforeseen event prevents the completion of planned operations. The first phase of the approach makes changes in a time-space network and in the input parameters to take into account the impacts of the revealed unforeseen event. The second phase solves a restricted master problem where the integrality constraints are relaxed and only the routes that were not impacted are added. The dual variables produced by this problem are used to price the columns generated from the time-space network. These routes must not deviate too much from the initial routes for the impacted trucks. The routes with negative reduced costs are added to the restricted master problem. The process is repeated until no negative reduced cost routes are available or a stopping criterion is met. To obtain integer solutions, we branch on the route variables according to the priority of the trucks they are assigned to and repeat the column generation process again.

Several disruption scenarios were developed to evaluate the proposed approach on case studies provided by FPInnovations. The results show that the produced routes remain cost-efficient compared to the initial routes. Also, it shows that the number of allowed differences with regards to the initial plan should be tuned to obtain the right trade-off between demand satisfaction and transportation cost.

Driven by the recent technological developments, the cost of collecting, storing and analyzing data is continuously decreasing. This may enable the collection of historical data about the unforeseen events in the forest industry. Having more accurate data about the probability distributions of the disruptions, future work may allow the development of stochastic models to take into account a priori the possible future unforeseen events.

CHAPTER 6

CONCLUSION

Growing interest in dynamic vehicle routing problems can be observed in the industrial and academic communities. In this thesis, we propose three approaches for a similar problem in forestry. The real-time log truck scheduling problem differs from a classical DVRP in several ways though. Furthermore, there may be large differences in the definition of the objectives of the problem depending on the context of application. Some contexts favor the use of heuristic methods to adjust the initial transportation plans after a disruption, while others allow the complete reoptimization of this plan.

Chapter 3 presents a context of application where it is necessary to provide truck drivers with the weekly transportation plans in advance. When an unforeseen event is revealed, the proposed heuristics try to limit, as much as possible, the changes in the initial plan. The approach starts by decomposing the impacted trips into trips to be delayed and others to be canceled. A first heuristic updates the transportation plan according to the delays incurred by these trucks. This can produce other trips to be canceled, especially when a delayed truck has a delivery that is planned close to the closing hours of the mills. The second heuristic tries then to reschedule the canceled trips in such a way that a loader is free when the truck arrives on a site. These heuristics provide the dispatcher with different rescheduling options. We developed different scenarios that represent the dispatcher choices of one of the available options after a disruption. The results show that this approach is able to reschedule, for most of the tested instances, at least half of the trips that could be lost if no reaction was performed in real-time.

Chapter 4 assumes a context where only one trip at a time is communicated to the drivers. This context allows to avoid drivers resistance to change in the case of a complete reoptimization of the transportation plan. However, the proposed solutions must be feasible in practice. Therefore, the proposed approach forbids diverting a truck from its current destination unless it is heading to a closed road or site. Also, this kind of reoptimizations is allowed only when the size of the instance is not very large, otherwise it may necessitate a large computational time. Unlike the contexts studied in Chapters 3 and 5, the reoptimization is done only for the current day and not for the whole week. The approach uses a time-space network to capture the impacts of the disruptions and solves the resulting problem using a commercial solver. To evaluate the performance of this approach, a simulation model was used to generate the disruptions, which were applied to case studies from the Canadian forest industry. The results show that this approach performs very well in comparison to an approach where we assume that all the disruptions are known beforehand.

Chapter 5 introduces a heuristic branch-and-price algorithm to solve a weekly dynamic LTSP. We assume a context where we try to minimize the changes in the initial transportation plan, but at a lower degree than in Chapter 3. The fleet of trucks is heterogeneous unlike in Chapters 3 and 4. This gives rise to new possible recourses in the case of a loader breakdown since some trucks have onboard loaders. However, this adds to the complexity of the problem. Similarly to Chapter 4, a time-space network is used to represent the impacts of the disruptions, but we added other types of arcs to represent the rest periods of truck drivers. We use column generation to generate new routes, which must not deviate too much from the initial routes. The simulation model used in Chapter 4 was also used to generate disruption scenarios to evaluate the approach. We also compared the approach to a similar one derived from a tactical problem that we adapted to the weekly LTSP. The proposed approach produces cost-efficient transportation plan in a small computational time.

Given the current state of the transportation planning environment in forestry, the use of the proposed methods in an industrial setting can be challenging. To the best of our knowledge, transportation planning optimization is somewhat marginal in the Canadian forest industry. The main goal of a log-truck dispatcher is to satisfy the demand.

Backhauling opportunities and queuing minimization by synchronizing the trucks with loaders are not usually taken into account. However, many companies recognize the need of optimization tools to increase the efficiency of their trucks. FPInnovations is leading negotiations with many forest companies to use Truck Scheduler. The developments of this thesis are compatible with this tool and with any optimization tool that produces a transportation plan.

In the current practice, disruptions are primarily managed via radio communication or over the phone (when possible) through the contractor, a dispatcher or directly with the driver depending on the organization of each company. The main goal is to ensure that enough demand is satisfied to ensure the proper functioning of the mills. There is no standard recourse to deal with each type of disruptions. Also, the information about a disruption is not available in real-time. However, many forest companies are currently investing in onboard tracking and monitoring systems. This will allow to collect the necessary data about disruptions in real-time. Still, a lot of work must be done to integrate the proposed models and methods with these systems. Since the decisions are made in a quick manner, one must also include procedures to check for data errors and inconsistent solutions.

The organization of log transportation varies between companies. In some cases, each driver is given a number of loads to transport and organizes the work independently from the other drivers. They are therefore paid per load rather than per hour. Using transportation plans where the trucks are assigned time intervals to be loaded and unloaded will require a change in the current practice. Some companies make detailed transportation plans that the trucks should follow. Their collaboration will result in more backhauling opportunities and lower transportation costs. However, information concerning supply and demand may be sensitive and could not be shared. A solution could be to engage a neutral transportation coordinator, but they would still have to figure out how savings should be allocated to each company. There exist also large transportation companies that do the transportation for several forest companies.

The methods proposed in this thesis would be more fruitful if the forest companies collaborate to manage their transportation.

Conducting the experiments presented in this thesis was challenging. The unforeseen events take different characteristics at each occurrence. For instance, a truck breakdown can be caused by a variety of factors. Also, its repair time will depend on various parameters such as the truck location, the breakdown cause, etc. Moreover, these events are interrelated. Forest road degradation, for example, may lead to its closure with a certain probability. Generating a disruption that is close to real practice necessitates an important work on the collection of historical data and the design of the statistical processes that would describe the different interrelations. This is a future research direction that could improve the current practices by including stochastic information during the optimization process.

In this work, we did not cover all the possible application contexts. For example, we did not consider the possibility of moving loaders between forest sites as a recourse. This can be useful if there is a new demand for a product, which is not available at the currently open sites. Of course, this demand must be large enough to justify such decisions. Also, we did not consider changing the configurations of the trucks by adding extra trailers. One could also allow the trucks to visit other forest sites in order to load these extra trailers.

For smaller instances, the expertise of the dispatchers might be of a great help. They generally construct high quality solutions with regards to operational requirements. The latter are often simplified in the mathematical models. One could use machine learning algorithms in order to "learn" how the dispatcher makes decisions. This can be done by defining a set of features for each instance. These features can consist of the positions of the trucks, the current demand and supply, the characteristics of the disruption, etc. The model is also provided with the decision taken in that case. Given a large number of instances, the model tries to recognize the pattern used to make the decisions. One can then use this pattern for larger problems. There can also be mutual learning where the dispatcher tries to learn from the solutions produced by optimization techniques.

BIBLIOGRAPHY

[1] A. Alonso-Ayuso, L. F. Escudero, M. Guignard, M. Quinteros, and A. Weintraub. Forestry management under uncertainty. *Annals of Operations Research*, 190(1):17–39, 2011.

[2] A. Amrouss, N. El Hachemi, M. Gendreau, and B. Gendron. Real-time management of transportation disruptions in forestry. Technical Report CIRRELT-2016-13, CIRRELT, 2016.

[3] A. Amrouss, M. Gendreau, and B. Gendron. Managing unforeseen events in forest transportation. *To appear in: J-FOR: Journal of Science and Technology for Forest Products and Processes*, 2016.

[4] N. Andalaft, P. Andalaft, M. Guignard, A. Magendzo, A. Wainer, and A. Weintraub. A problem of forest harvesting and road building solved through model strengthening and lagrangean relaxation. *Operations Research*, 51(4):613–628, 2003.

[5] G. Andersson, P. Flisberg, B. Lidén, and M. Rönnqvist. Ruttopt-a decision support system for routing of logging trucks. *Canadian Journal of Forest Research*, 38(7):1784–1796, 2008.

[6] C. Archetti and M. G. Speranza. The split delivery vehicle routing problem: A survey. In B. Golden, S. Raghavan, and E. Wasil, editors, *The Vehicle Routing Problem: Latest Advances and New Challenges*, pages 103–122. Springer US, Boston, MA, 2008.

[7] C. Archetti and M. G. Speranza. Vehicle routing problems with split deliveries. *International Transactions in Operational Research*, 19(1-2):3–22, 2012.

[8] A. Attanasio, J. Bregman, G. Ghiani, and E. Manni. Real-time fleet management at ecourier ltd. In *Dynamic Fleet Management*, pages 219–238. Springer, 2007.

[9] J.-F. Audy, S. D'Amours, and M. Rönnqvist. Planning methods and decision support systems in vehicle routing problems for timber transportation: A review. Technical Report CIRRELT-2012-38, CIRRELT, 2012.

[10] B. B. Bare, D. G. Briggs, J. P. Roise, and G. F. Schreuder. A survey of systems analysis models in forestry and the forest products industries. *European Journal of Operational Research*, 18(1):1–18, 1984.

[11] T. Barrett, J. Gilless, and L. Davis. Economic and fragmentation effects of clearcut restrictions. *Forest Science*, 44(4):569–577, 1998.

[12] D. Beaudoin, J.-M. Frayret, and L. LeBel. Negotiation-based distributed wood procurement planning within a multi-firm environment. *Forest Policy and Economics*, 12(2):79–93, 2010.

[13] A. Beaudry, G. Laporte, T. Melo, and S. Nickel. Dynamic transportation of patients in hospitals. *OR Spectrum*, 32(1):77–107, 2010.

[14] G. Berbeglia, J.-F. Cordeau, I. Gribkovskaia, and G. Laporte. Static pickup and delivery problems: a classification scheme and survey. *Top*, 15(1):1–31, 2007.

[15] P. Bettinger and M. Lennette. Landscape management policy simulator (lamps): version 1.1: user guide. Technical report, Corvallis, OR: Oregon State University, College of Forestry, Forest Research Laboratory, 2004.

[16] J. G. Borges, L. Diaz-Balteiro, M. E. McDill, and L. C. Rodriguez. *The management of industrial forest plantations*. Springer, 2014.

[17] R. M. Branchini, V. A. Armentano, and A. Løkketangen. Adaptive granular local search heuristic for a dynamic vehicle routing problem. *Computers & Operations Research*, 36(11):2955–2968, 2009.

[18] D. Bredström, P. Jönsson, and M. Rönnqvist. Annual planning of harvesting resources in the forest industry. *International Transactions in Operational Research*, 17(2):155–177, 2010.

[19] D. Bredström and M. Rönnqvist. Combined vehicle routing and scheduling with temporal precedence and synchronization constraints. *European Journal of Operational Research*, 191(1):19–31, 2008.

[20] D. G. Briggs. Tree value system: Description and assumptions. Technical report, Portland, OR: U.S. Department of Agriculture, Forest Service, Pacific Northwest Research Station, 1989.

[21] M. Caramia, G. F. Italiano, G. Oriolo, A. Pacifici, and A. Perugia. Routing a fleet of vehicles for dynamic combined pick-up and deliveries services. In *Operations Research Proceedings 2001*, pages 3–8. Springer, 2002.

[22] D. Carlsson, S. D'Amours, A. Martel, and M. Rönnqvist. Supply chain planning models in the pulp and paper industry. *INFOR: Information Systems and Operational Research*, 47(3):167–183, 2009.

[23] D. Carlsson, P. Flisberg, and M. Rönnqvist. Using robust optimization for distribution and inventory planning for a large pulp producer. *Computers & Operations Research*, 44:214–225, 2014.

[24] D. Carlsson and M. Rönnqvist. Supply chain management in forestry—-case studies at Södra Cell AB. *European Journal of Operational Research*, 163(3):589–616, 2005.

[25] D. Carlsson and M. Rönnqvist. Backhauling in forest transportation: models, methods, and practical usage. *Canadian Journal of Forest Research*, 37(12):2612–2623, 2007.

[26] Z.-L. Chen and H. Xu. Dynamic column generation for dynamic vehicle routing with time windows. *Transportation Science*, 40(1):74–88, 2006.

[27] G. Clarke and J. W. Wright. Scheduling of vehicles from a central depot to a number of delivery points. *Operations Research*, 12(4):568–581, 1964.

[28] J.-F. Cordeau, G. Laporte, A. Mercier, et al. A unified tabu search heuristic for vehicle routing problems with time windows. *Journal of the Operational Research Society*, 52(8):928–936, 2001.

[29] T. G. Crainic, M. Gendreau, and J.-Y. Potvin. Intelligent freight-transportation systems: Assessment and the contribution of operations research. *Transportation Research Part C: Emerging Technologies*, 17(6):541–557, 2009.

[30] S. D'Amours, M. Rönnqvist, and A. Weintraub. Using operational research for supply chain planning in the forest products industry. *INFOR: Information Systems and Operational Research*, 46(4):265–281, 2008.

[31] J. L. de Moraes Gonçalves, L. D. Silva, M. Behling, and C. A. Alvares. Management of industrial forest plantations. In G. J. Borges, L. Diaz-Balteiro, E. M. McDill, and C. L. Rodriguez, editors, *The Management of Industrial Forest Plantations: Theoretical Foundations and Applications*, pages 91–119. Springer Netherlands, Dordrecht, 2014.

[32] M. Drexl. Synchronization in vehicle routing-a survey of vrps with multiple synchronization constraints. *Transportation Science*, 46(3):297–316, 2012.

[33] M. Dror, G. Laporte, and P. Trudeau. Vehicle routing with split deliveries. *Discrete Applied Mathematics*, 50(3):239–254, 1994.

[34] M. Dror and P. Trudeau. Savings by split delivery routing. *Transportation Science*, 23(2):141–145, 1989.

[35] T. Du, F. Wang, and P.-Y. Lu. A real-time vehicle-dispatching system for consolidating milk runs. *Transportation Research Part E: Logistics and Transportation Review*, 43(5):565–577, 2007.

[36] Y. Dumas, J. Desrosiers, and F. Soumis. The pickup and delivery problem with time windows. *European Journal of Operational Research*, 54(1):7–22, 1991.

[37] N. El Hachemi. *Problème de transport avec contraintes d'horaires*. PhD thesis, École Polytechnique de Montréal, 2009.

[38] N. El Hachemi, I. El Hallaoui, M. Gendreau, and L.-M. Rousseau. Flow-based integer linear programs to solve the weekly log-truck scheduling problem. *Annals of Operations Research*, 232(1):87–97, 2015.

[39] N. El Hachemi, M. Gendreau, and L.-M. Rousseau. A hybrid constraint programming approach to the log-truck scheduling problem. *Annals of Operations Research*, 184(1):163–178, 2011.

[40] N. El Hachemi, M. Gendreau, and L.-M. Rousseau. A heuristic to solve the synchronized log-truck scheduling problem. *Computers & Operations Research*, 40(3):666–673, 2013.

[41] G. Eng, H. Daellenbach, and A. Whyte. Bucking tree-length stems optimally. *Canadian Journal of Forest Research*, 16(5):1030–1035, 1986.

[42] R. Epstein, J. Karlsson, M. Rönnqvist, and A. Weintraub. Harvest operational models in forestry. In A. Weintraub, C. Romero, T. Bjørndal, R. Epstein, and J. Miranda, editors, *Handbook Of Operations Research In Natural Resources*, pages 365–377. Springer US, Boston, MA, 2007.

[43] R. Epstein, R. Morales, J. Seron, and A. Weintraub. Use of or systems in the chilean forest industries. *Interfaces*, 29(1):7–29, 1999.

[44] R. Epstein, A. Weintraub, P. Sapunar, E. Nieto, J. B. Sessions, J. Sessions, F. Bustamante, and H. Musante. A combinatorial heuristic approach for solving real-size machinery location and road design problems in forestry planning. *Operations Research*, 54(6):1017–1027, 2006.

[45] P. Flisberg, M. Frisk, and M. Rönnqvist. Integrated harvest and logistic planning including road upgrading. *Scandinavian Journal of Forest Research*, 29(sup1):195–209, 2014.

[46] Food and Agriculture Organization of the United Nations. *Global forest resources assessment 2010: Main report*. 2010.

[47] Forest Products Association of Canada. *Industry by the Numbers*. 2014.

[48] M. Forsberg, M. Frisk, and M. Rönnqvisty. Flowopt–a decision support tool for strategic and tactical transportation planning in forestry. *International Journal of Forest Engineering*, 16(2):101–114, 2005.

[49] M. Frisk, M. Göthe-Lundgren, K. Jörnsten, and M. Rönnqvist. Cost allocation in collaborative forest transportation. *European Journal of Operational Research*, 205(2):448–458, 2010.

[50] L. M. Gambardella, A. E. Rizzoli, F. Oliverio, N. Casagrande, A. Donati, R. Montemanni, and E. Lucibello. Ant colony optimization for vehicle routing in advanced logistics systems. In *Proceedings of the International Workshop on Modelling and Applied Simulation (MAS 2003)*, pages 3–9, 2003.

[51] O. Garcıa. FOLPI, a forestry-oriented linear programming interpreter. In *Proceedings IUFRO Symposium on Forest Management Planning and Managerial Economics', University of Tokyo, Tokyo*, pages 293–305, 1984.

[52] M. Gendreau, F. Guertin, J.-Y. Potvin, and E. Taillard. Parallel tabu search for real-time vehicle routing and dispatching. *Transportation science*, 33(4):381–390, 1999.

[53] B. L. Golden, S. Raghavan, and E. A. Wasil. *The Vehicle Routing Problem: Latest Advances and New Challenges*, volume 43. Springer Science & Business Media, 2008.

[54] M. Goycoolea, A. Murray, J. P. Vielma, and A. Weintraub. Evaluating approaches for solving the area restriction model in harvest scheduling. *Forest Science*, 55(2):149–165, 2009.

[55] M. Goycoolea, A. T. Murray, F. Barahona, R. Epstein, and A. Weintraub. Harvest scheduling subject to maximum area restrictions: exploring exact approaches. *Operations Research*, 53(3):490–500, 2005.

[56] F. E. Greulich and W. G. O'Regan. Optimum use of air tankers in initial attack: selection, basing, and transfer rules. Technical Report PSW-163, Berkeley, CA: Pacific Southwest Forest and Range Experiment Station, Forest Service, U.S. Department of Agriculture, 1982.

[57] M. Gronalt and P. Hirsch. Log-truck scheduling with a tabu search strategy. In K. F. Doerner, M. Gendreau, P. Greistorfer, W. Gutjahr, R. F. Hartl, and M. Reimann, editors, *Metaheuristics: Progress in Complex Systems Optimization*, pages 65–88. Springer US, Boston, MA, 2007.

[58] D. J. Gulczynski, B. Golden, and E. Wasil. Recent developments in modeling and solving the split delivery vehicle routing problem. *Tutorials in Operations Research*, pages 170–180, 2008.

[59] E. Gunn. Some perspectives on strategic forest management models and the forest products supply chain. *INFOR: Information Systems and Operational Research*, 47(3):261–272, 2009.

[60] E. A. Gunn. Models for strategic forest management. In A. Weintraub, C. Romero, T. Bjørndal, R. Epstein, and J. Miranda, editors, *Handbook Of Operations Research In Natural Resources*, pages 317–341. Springer US, Boston, MA, 2007.

[61] A. Haghani and S. Jung. A dynamic vehicle routing problem with time-dependent travel times. *Computers & Operations Research*, 32(11):2959–2986, 2005.

[62] A. Haghani and S. Yang. Real-time emergency response fleet deployment: Concepts, systems, simulation & case studies. In *Dynamic fleet management*, pages 133–162. Springer, 2007.

[63] J. Hof. *Coactive forest management.* Academic Press, Inc., 1992.

[64] S. Ichoua, M. Gendreau, and J.-Y. Potvin. Diversion issues in real-time vehicle dispatching. *Transportation Science*, 34(4):426–438, 2000.

[65] S. Irnich and G. Desaulniers. Shortest path problems with resource constraints. In G. Desaulniers, J. Desrosiers, and M. M. Solomon, editors, *Column Generation*, pages 33–65. Springer US, Boston, MA, 2005.

[66] J. Jones, J. Hyde, and M. Meachan. Four analytical approaches for integrating land and trans portation planning on forest lands. Technical Report INT361, Odgen, Utah : United States Dept. of Agriculture, Forest Service, Intermountain Research Station, 1986.

[67] B. Kallehauge, J. Larsen, O. B. Madsen, and M. M. Solomon. Vehicle routing problem with time windows. In G. Desaulniers, J. Desrosiers, and M. M. Solomon, editors, *Column Generation*, pages 67–98. Springer US, Boston, MA, 2005.

[68] J. Karlsson, M. Rönnqvist, and M. Frisk. Roadopt: A decision support system for road upgrading in forestry. *Scandinavian Journal of Forest Research*, 21(S7):5–15, 2006.

[69] B. Kent, B. B. Bare, R. C. Field, and G. A. Bradley. Natural resource land management planning using large-scale linear programs: The USDA forest service experience with FORPLAN. *Operations Research*, 39(1):13–27, 1991.

[70] A. J. Laroze and B. J. Greber. Using tabu search to generate stand-level, rule-based bucking patterns. *Forest Science*, 43(2):157–169, 1997.

[71] P. L'Ecuyer, L. Meliani, and J. Vaucher. SSJ: a framework for stochastic simulation in Java. In *Proceedings of the 34th conference on Winter simulation: exploring new frontiers*, pages 234–242. Winter Simulation Conference, 2002.

[72] C.-G. Lee, M. A. Epelman, C. C. White, and Y. A. Bozer. A shortest path approach to the multiple-vehicle routing problem with split pick-ups. *Transportation Research Part B: Methodological*, 40(4):265–284, 2006.

[73] N. Lehoux, S. D'Amours, and A. Langevin. Methodology for assessing collaboration strategies and incentives in the pulp and paper industry. In T.-M. Choi and E. T. C. Cheng, editors, *Supply Chain Coordination under Uncertainty*, pages 625–650. Springer, Berlin, Heidelberg, 2011.

[74] D. Lesaint, C. Voudouris, and N. Azarmi. Dynamic workforce scheduling for British Telecommunications plc. *Interfaces*, 30(1):45–56, 2000.

[75] C.-L. Li, D. Simchi-Levi, and M. Desrochers. On the distance constrained vehicle routing problem. *Operations Research*, 40(4):790–799, 1992.

[76] J.-Q. Li, P. B. Mirchandani, and D. Borenstein. A lagrangian heuristic for the real-time vehicle rescheduling problem. *Transportation Research Part E: Logistics and Transportation Review*, 45(3):419–433, 2009.

[77] S. Linnainmaa, J. Savola, and O. Jokinen. EPO: A knowledge based system for wood procurement management. In *Proceedings of the 7th Innovative Applications of Artificial Intelligence Conference*, pages 107–113, 1995.

[78] A. F. Marques, J. P. de Sousa, and M. Rönnqvist. Combining optimization and simulation tools for short-term planning of forest operations. *Scandinavian Journal of Forest Research*, 29(sup1):166–177, 2014.

[79] D. L. Martell. A review of operational research studies in forest fire management. *Canadian Journal of Forest Research*, 12(2):119–140, 1982.

[80] D. L. Martell. Forest fire management. In A. Weintraub, C. Romero, T. Bjørndal, R. Epstein, and J. Miranda, editors, *Handbook of Operations Research in Natural Resources*, pages 489–509. Springer US, Boston, MA, 2007.

[81] D. L. Martell, L. S. Davis, A. Weintraub, et al. *Proceedings of a workshop on hierarchical approaches to forest management in public and private organizations*, volume 124. 1996.

[82] D. L. Martell, E. A. Gunn, and A. Weintraub. Forest management challenges for operational researchers. *European Journal of Operational Research*, 104(1):1–17, 1998.

[83] M. E. McDill and J. Braze. Comparing adjacency constraint formulations for randomly generated forest planning problems with four age-class distributions. *Forest Science*, 46(3):423–436, 2000.

[84] M. E. McDill, S. A. Rebain, and J. Braze. Harvest scheduling with area-based adjacency constraints. *Forest Science*, 48(4):631–642, 2002.

[85] T. McDonald, K. Haridass, and J. Valenzuela. Mileage savings from optimization of coordinated trucking. In *Mitchell, D.; Gallagher, T., eds. Proceedings of 2010 COFE: 33rd Annual Meeting of the Council on Forest Engineering*, 2010.

[86] T. P. McDonald, S. E. Taylor, R. B. Rummer, and J. Valenzuela. Information needs for increasing log transport efficiency. In *First International Precision Forestry Symposium*, 2001.

[87] S. Mitra. A parallel clustering technique for the vehicle routing problem with split deliveries and pickups. *Journal of the Operational Research Society*, 59(11):1532–1546, 2008.

[88] R. Montemanni, L. M. Gambardella, A. E. Rizzoli, and A. V. Donati. Ant colony system for a dynamic vehicle routing problem. *Journal of Combinatorial Optimization*, 10(4):327–343, 2005.

[89] A. T. Murray. Spatial restrictions in harvest scheduling. *Forest Science*, 45(1):45–52, 1999.

[90] A. T. Murray and R. L. Church. Analyzing cliques for imposing adjacency restrictions in forest models. *Forest Science*, 42(2):166–175, 1996.

[91] Natural Resources Canada | Ressources Naturelles Canada. *State of Canada's forest annual report*. 2014.

[92] D. I. Navon. Timber RAM... a long-range planning method for commercial timber lands under multiple-use management. Technical Report PSW-RP-070, Albany, CA: Pacific Southwest Research Station, Forest Service, U.S. Department of Agriculture, 1971.

[93] C. Novoa and R. Storer. An approximate dynamic programming approach for the vehicle routing problem with stochastic demands. *European Journal of Operational Research*, 196(2):509–515, 2009.

[94] M. Nowak, Ö. Ergun, and C. C. White III. Pickup and delivery with split loads. *Transportation Science*, 42(1):32–43, 2008.

[95] K. Öhman. Creating continuous areas of old forest in long-term forest planning. *Canadian Journal of Forest Research*, 30(11):1817–1823, 2000.

[96] L. Olsson. *Optimisation of forest road investments and the roundwood supply chain*. PhD thesis, Department of Forest Economics, Swedish University of Agricultural Sciences, 2004.

[97] M. Palmgren, M. Rönnqvist, and P. Värbrand. A column generation algorithm for the log truck scheduling problem. *Department of Science and Technology, LiTH-MAT-R, Linköping University, Norrköping Sweden*, pages 3–15, 2001.

[98] M. Palmgren, M. Rönnqvist, and P. Värbrand. A solution approach for log truck scheduling based on composite pricing and branch and bound. *International Transactions in Operational Research*, 10(5):433–447, 2003.

[99] M. Palmgren, M. Rönnqvist, and P. Värbrand. A near-exact method for solving the log-truck scheduling problem. *International Transactions in Operational Research*, 11(4):447–464, 2004.

[100] V. Pillac, M. Gendreau, C. Guéret, and A. L. Medaglia. A review of dynamic vehicle routing problems. *European Journal of Operational Research*, 225(1):1–11, 2013.

[101] H. N. Psaraftis. A dynamic programming solution to the single vehicle many-to-many immediate request dial-a-ride problem. *Transportation Science*, 14(2):130–154, 1980.

[102] A. Regan, H. Mahmassani, and P. Jaillet. Dynamic decision making for commercial fleet operations using real-time information. *Transportation Research Record: Journal of the Transportation Research Board*, 1537:91–97, 1996.

[103] A. C. Regan, H. S. Mahmassani, and P. Jaillet. Improving efficiency of commercial vehicle operations using real-time information: potential uses and assignment strategies. *Transportation Research Record*, 1493:188–198, 1995.

[104] P. A. Rey, J. A. Muñoz, and A. Weintraub. A column generation model for truck routing in the Chilean forest industry. *INFOR: Information Systems and Operational Research*, 47(3):215–221, 2009.

[105] G. Rix, L.-M. Rousseau, and G. Pesant. A column generation algorithm for tactical timber transportation planning. *Journal of the Operational Research Society*, 66(2):278–287, 2014.

[106] J. G. Rix. *Transportation Optimization in Tactical and Operational Wood Procurement Planning*. PhD thesis, Polytechnique Montréal, 2014.

[107] M. Rönnqvist. Optimization in forestry. *Mathematical Programming*, 97(1-2):267–284, 2003.

[108] M. Rönnqvist. O.R. challenges and experiences from solving industrial applications. *International Transactions in Operational Research*, 19(1-2):227–251, 2012.

[109] M. Rönnqvist, S. D'Amours, A. Weintraub, A. Jofre, E. Gunn, R. G. Haight, D. Martell, A. T. Murray, and C. Romero. Operations research challenges in forestry: 33 open problems. *Annals of Operations Research*, 232(1):11–40, 2015.

[110] M. Rönnqvist and D. Ryan. Solving truck despatch problems in real time. In *Proceedings of the 31st annual conference of the Operational Research Society of New Zealand*, volume 31, pages 165–172, 1995.

[111] N. Secomandi and F. Margot. Reoptimization approaches for the vehicle-routing problem with stochastic demands. *Operations Research*, 57(1):214–230, 2009.

[112] J. Sessions and J. Sessions. Tactical harvest planning using snap 2.03. In *Proceedings of the Conference of the International Union of Forestry Research*, 1991.

[113] S. A. Snyder, R. G. Haight, and C. S. ReVelle. A scenario optimization model for dynamic reserve site selection. *Environmental Modeling & Assessment*, 9(3):179–187, 2005.

[114] P. Toth and D. Vigo. *Vehicle routing: Problems, Methods, and Applications*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2014.

[115] R. H. Twito, S. E. Reutebuch, R. J. McGaughey, and C. N. Mann. Preliminary logging analysis system (plans): overview. *Tech. Rep. PNW-GTR-199. Portland, OR: U.S. Department of Agriculture, FForest Service, Pacific Northwest Research Station*, 1987.

[116] J. I. Van Hemert and J. A. La Poutré. Dynamic routing problems with fruitful regions: Models and evolutionary computation. In *PPSN VIII: 8th International Conference*, pages 692–701, Birmingham, UK, September 18-22 2004. Springer.

[117] F. B. Veliz, J.-P. Watson, A. Weintraub, R. J.-B. Wets, and D. L. Woodruff. Stochastic optimization models in forest planning: A progressive hedging solution approach. *Annals of Operations Research*, 232(1):259–274, 2015.

[118] A. Weintraub, F. Barahona, and R. Epstein. A column generation algorithm for solving general forest planning problems with adjacency constraints. *Forest science*, 40(1):142–161, 1994.

[119] A. Weintraub and A. Cholaky. A hierarchical approach to forest planning. *Forest Science*, 37(2):439–460, 1991.

[120] A. Weintraub, R. Epstein, R. Morales, J. Seron, and P. Traverso. A truck scheduling system improves efficiency in the forest industries. *Interfaces*, 26(4):1–12, 1996.

[121] A. Weintraub and C. Romero. Operations research models and the management of agricultural and forestry resources: a review and comparison. *Interfaces*, 36(5):446–457, 2006.

[122] A. Weintraub, C. Romero, T. Bjørndal, R. Epstein, and J. Miranda. *Handbook Of Operations Research In Natural Resources*. Springer Science & Business Media, Boston, MA, 2007.

[123] A. Weintraub and J. Vera. A cutting plane approach for chance constrained linear programs. *Operations Research*, 39(5):776–785, 1991.

[124] P. Wikström, L. Edenius, B. Elfving, L. O. Eriksson, T. Lämås, J. Sonesson, K. Öhman, J. Wallerman, C. Waller, and F. Klintebäck. The heureka forestry decision support system: an overview. *Mathematical and Computational Forestry & Natural Resource Sciences*, 3(2):87, 2011.

[125] J. Yang, P. Jaillet, and H. Mahmassani. Real-time multivehicle truckload pickup and delivery problems. *Transportation Science*, 38(2):135–148, 2004.

[126] R. Yousefpour and M. Hanewinkel. Modelling of forest conversion planning with an adaptive simulation-optimization approach and simultaneous consideration of the values of timber, carbon and biodiversity. *Ecological Economics*, 68(6):1711–1722, 2009.

[127] V. Zeimpekis, I. Minis, K. Mamassis, and G. Giaglis. Dynamic management of a delayed delivery vehicle in a city logistics environment. In V. Zeimpekis, C. D. Tarantilis, G. M. Giaglis, and I. Minis, editors, *Dynamic Fleet Management: Concepts, Systems, Algorithms & Case Studies*, pages 197–217. Springer US, Boston, MA, 2007.