Transportation Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

TSL Dissertation Abstracts—2006 Transportation Science and Logistics Society Dissertation Prize Competition

Hickman Mark,

To cite this article:
Hickman Mark, (2007) TSL Dissertation Abstracts—2006 Transportation Science and Logistics Society Dissertation Prize Competition. Transportation Science 41(3):415-429. https://doi.org/10.1287/trsc.1070.0213

TRANSPORTATION SCIENCE
Vol. 41, No. 3, August 2007, pp. 415–429
ISSN 0041-1655 | EISSN 1526-5447 | 07 | 4103 | 0415

informs

DOI 10.1287/trsc.1070.0213

TSL Dissertation Abstracts

2006 Transportation Science and Logistics Society Dissertation Prize Competition

Mark Hickman
Chair, 2006 TSL Dissertation Prize Competition

The Transportation Science and Logistics (TSL) Society Dissertation Prize Competition is the oldest and most prestigious competition for doctoral dissertations in the field. The 2006 TSL dissertation prize committee consisted of Professor Michel Bierlaire (École Polytechnique Fédérale de Lausanne), Professor Ann Campbell (University of Iowa), Professor Alan Erera (Georgia Institute of Technology), Professor Donald Hearn (University of Florida), and Professor Mark Hickman (University of Arizona).

Eligibility was limited to doctoral dissertations completed between June 1, 2005, and May 31, 2006, in the general area of transportation science and logistics. To be considered, a dissertation must have been nominated by the author's thesis supervisor. This year we received 26 nominations from 24 universities and 10 countries, representing such diverse areas as vehicle routing, vehicle scheduling, traveler behavior, activity analysis, logistics, and supply chain management.

This year's pool of dissertations was exceptional in both its breadth and depth, showing remarkable theoretical and practical insights into many very interesting problems in transportation and logistics. With the large number of submissions and their exceptional quality, it was not easy for the committee to select a first place winner and an honorable mention. I would like to thank the authors, their advisors, and the prize committee for their part in this effort to recognize and honor the field's outstanding doctoral researchers. Awards were announced at the TSL meeting at the 2006 INFORMS meeting in Pittsburgh.

Perhaps without saying, the theses nominated for the TSL Dissertation Prize are distinguished by the fact that they were nominated for the award. Each dissertation makes a unique and valuable contribution to transportation science and logistics, and all entrants and their advisors deserve our sincere congratulations.

Abstracts for the 26 dissertations follow; awardees are listed first, and finalists' abstracts follow in alphabetical order by author.

First Place
Christian Liebchen, Technische Universität Berlin, Germany
Advisors: Rolf H. Möhring and Karl Nachtigall
"Periodic Timetable Optimization in Public Transport"

Honorable Mention
Anne Goodchild, University of California, Berkeley
Advisor: Carlos Daganzo
"Crane Double Cycling in Container Ports: Algorithms, Evaluation, and Planning"

Transportation Science 41(3), pp. 415–429, © 2007 INFORMS

Periodic Timetable Optimization in Public Transport

Christian Liebchen
Technische Universität Berlin, Germany
Rolf H. Möhring and Karl Nachtigall (Advisors)

"The timetable is the essence of the service offered by any provider of public transport" (Jonathan Tyler, "Philosophies of Timetabling, Definitions of Bottlenecks and the Usefulness of Spreadsheets: The experience of a practical strategic timetable planner, CASPT 2006, CD-ROM). Despite this observation, in the practice of planning public transportation, only a few months ago OR decision support has still been limited to operations planning (vehicle scheduling, duty scheduling, crew rostering). The present thesis describes the optimization techniques that were deployed in computing the very first optimized timetable that went into daily service: the 2005 timetable of the Berlin subway. Compared to the former timetable, with our timetable the passengers of Berlin subway are offered simultaneously improvements in three key criteria, which typically are conflicting: transfer waiting time, dwell time of trains, and number of trains that are required for the operation.

The basic graph model, the periodic event scheduling problem (PESP), is known for 15 years and it had attracted many research groups. Nevertheless, this thesis collects significant progress that has been made only recently on issues like solution strategies or modeling capabilities. The latter include additional timetabling modeling features (train coupling/train sharing, bundling of lines of similar speed in order to increase track capacity), theoretical evidence that the symmetry requirement for periodic timetables cannot be expressed by only using (integer-valued) PESP constraints, and even the integration of further planning tasks in public transport, such as line planning.

On the theory side, a more precise notion of the asymptotical complexity of the PESP is given by providing a MAXSNP-hardness proof as a kind of negative result. On the positive side, the design of more efficient algorithms gave rise to a much deeper understanding of cycle bases of graphs, another very hot topic in discrete mathematics during the last three years. In 2005, this culminated in the drawing of the complete map for the seven relevant classes of cycle bases, and the design of the fastest algorithms for the minimum directed cycle basis problem and the minimum 2-basis problem.

These theoretical considerations are complemented by an exhaustive computational study. For the first time, the performance of a cut-and-branch algorithm, constraint programming, and even a genetic algorithm were compared for periodic timetable optimization. the cut-and-branch algorithm and the genetic algorithm turned out to be the most convincing. On the one hand, the genetic algorithm has a slight advantage with respect to independence of parameter settings. On the other hand, the quality of the solutions obtained by the cut-and-branch algorithm was somewhat better.

Crane Double Cycling in Container Ports: Algorithms, Evaluation, and Planning

Anne Goodchild
University of California, Berkeley
Carlos Daganzo (Advisor)

Loading ships as they are unloaded (double cycling) can improve the efficiency of a quay crane and container port. This dissertation describes the double-cycling problem, and presents solution algorithms and simple formulae to estimate benefits. First, a lower bound to the number of cycles necessary to turn around a ship is developed. We then present a greedy algorithm that yields a tight upper bound. The problem is formulated as a scheduling problem, which can be solved optimally using Johnson's rule. Ships with deck hatches are also considered.

We then consider the longer term impact of double cycling on port operations. We use another double-cycling sequence that is operationally convenient, easy to model, and nearly optimum, and compare the performance of this sequence to those determined by a greedy algorithm, and Johnson's rule. A simple formula is developed to predict the longer term impact on turn around time. We also comment on strategies for reducing landside vehicle requirements, and altering port operations to support double cycling such as segmenting vessel storage and streamlining

traffic flows. We show that double cycling can reduce the amount of time required to complete vessel loading and unloading operations by 10%, improving vessel, crane, and berth productivity.

Finally we consider somewhat broader issues, including the interaction of double cycling and security regulations, as well as ship design and routing. We estimate the financial impact of the benefits, which total approximately $70.00 per container moved.

The research demonstrates that double cycling can create significant efficiency gains while simplifying some aspects of port operations. We offer an explanation as to why the method has not been more rapidly adopted. We demonstrate that complementary analytical methods can be used to gain broad insight, and that simple, general models, are often of more value than more complex, detailed models.

A Discrete Choice Modeling Framework for Pedestrian Walking Behavior with Application to Human Tracking in Video Sequences

Gianluca Antonini
École Polytechnique Fédérale de Lausanne, Switzerland
Jean Philippe Thiran and Michel Bierlaire (Advisors)

This work addresses the problem of pedestrian walking behavior modeling and its integration with computer vision and data analysis techniques for automatic tracking and counting of people in video sequences, providing contributions to the state of the art of each of those research fields. First, a new pedestrian walking behavior model based on discrete choice analysis has been proposed. The model captures different behaviors: individuals can walk under both a constrained and an unconstrained regime. In the first regime, interactions between people are captured by means of a leader follower and collision avoidance patterns. In the second, free flow patterns are considered. This approach makes the model adaptive, able to represent walking behavior at different levels of density. The model has been calibrated on real pedestrian data, collected manually tracking people in video sequences. We provide an interpretation of the model coefficient estimates that is coherent with the behavioral assumptions made at the beginning. Finally, in order to refine and better validate the model, a pedestrian simulator has been developed.

The second contribution consists in the integration of state-of-the-art computer vision techniques for object tracking with the proposed model, in order to perform automatic pedestrian tracking in video sequences. A simple correlation-based tracking method has been combined with the behavioral model under both deterministic and probabilistic frameworks. Our research shows that combining behavioral information with image-based methods represents an useful tool to analyze complex sequences.

A final contribution consists in the investigation of hierarchical clustering techniques in order to perform automatic counting of people in video sequences. Several data representations, metrics and grouping rules are tested in order to cluster the pedestrian trajectories tracked by the model-based system. The result is a better approximation of the real number of individuals present in the scene.

Solving the Multicommodity Flow Problem with the Analytic Center Cutting Plane Method

Frédéric Babonneau
Université de Genéve, Switzerland
Jean-Philippe Vial (Advisor)

The focus of this research is the implementation of efficient numerical methods to solve the multicommodity flow problems (MCF) that mainly arises in the areas of transportation and telecommunications. The main contributions of the thesis are the modification and the specialization of the analytic center cutting plane method (ACCPM) and the implementation of an active set strategy that permits to considerably reduce the dimension of the problem. It appears that the new proposed approach always improves the results published in the literature in the linear and nonlinear case and makes it possible to solve huge problems.

First, we point out that the Lagrangian dual objective function of MCF has two main components: a piece-wise linear one and smooth one that is the negative Fenchel conjugate of the congestion function. The latter is smooth and can often be computed in closed form. ACCPM has then been specialized

Transportation Science 41(3), pp. 415–429, © 2007 INFORMS

to handle explicitly the available smooth component. The new approach considerably improves the performance of the former implementation of ACCPM and makes it possible to solve all the instances.

The implementation of an active set strategy is also proposed to speed up the solution method. This is motivated in linear MCF by our observation that on practical problems, the number of congested arcs in an optimal solution is a small fraction of the total number of arcs in the graph. Consequently, the Lagrangian dual variables associated with these arcs must be null at the optimum. In practice, the set of saturated arcs at the optimum is not known. Thus we implemented an active set strategy that dynamically estimates this set during the solution method. Then we propose an approximation scheme that replaces the objective function near the origin by a linear one to extend the active set strategy to the nonlinear MCF.

Service Choice in Periodic Distribution Operations

Peter Francis
Northwestern University

Karen Smilowitz (Advisor)

This research investigates service-differentiated routing problems. It is motivated by the need to provide customers with appropriate levels of service in the face of operational constraints and rising competition among logistical service providers. We examine the impact of service differentiation on the operation of a distribution system, specifically in the case when service quality is defined by the frequency of service provided. We develop a new model called the period vehicle routing problem with service-choice (PVRP-SC) that fundamentally extends the basic period vehicle routing problem (PVRP) by allowing additional flexibility in choosing levels of service for customers. We show that such modeling results in better solutions than those from methods that allocate service levels to customers a priori, based on factors such as customer demand that may not provide sufficient operational flexibility.

We present an integer programming formulation for the PVRP-SC and show how modeling differentiated operations with service choice can result in more efficient routing solutions and higher levels of service benefit. The new model allocates service levels to customers in a manner that minimizes a weighted objective of operational costs net of service benefits.

We present an exact solution method for the PVRP-SC, which is the first known exact solution method for the PVRP. We show a heuristic variant of the exact solution method to find solutions for larger instances. We observe that efficiency gains from introducing service choice depend on the distribution of nodes in the problem instance.

The PVRP and PVRP-SC belong to the NP-hard class of problems. To solve instances of reasonable size using the exact solution method, certain assumptions and simplifications are required. Therefore, we develop a heuristic solution method to solve the PVRP-SC and investigate the effect of relaxing these assumptions. The Tabu search heuristic allows modelers to examine trade-offs between efficiency gains and the complexity of the solutions produced.

Responding to Casualties in a Disaster Relief Operation: Initial Ambulance Allocation and Reallocation, and Switching of Casualty Priorities

Qiang Gong
The State University of New York at Buffalo

Rajan Batta (Advisor)

This research is concerned with models for response to casualties in a disaster relief operation. Three problems are analyzed. The first is that of initial ambulance allocation to casualty clusters. The second is that of ambulance reallocation between casualty clusters. The third is that of switching casualty priorities. We briefly describe each contribution.

The first problem analyzes a deterministic ambulance allocation model for a post-disaster relief operation. Casualties in a disaster tend to be numerous and distributed in space, typically forming clusters. We formulate a deterministic model that depicts how a cluster grows after a disaster strikes. Based on the model, we develop methods to calculate critical time measures, e.g., completion time for each cluster. Then we present two iterative procedures to optimize the makespan and the weighted total flow time, respectively. Our methods are illustrated via a case study, which is based on an earthquake in Northridge, California. The main conclusion is that the optimal

ambulance allocation can be significantly dependent upon the desired performance measure.

The second problem analyzes the ambulance reallocation problem on the basis of a discrete time policy. The benefits of redistribution include providing service to new clusters and fully utilizing ambulances. We consider the objective of minimizing makespan. The complication is that the distance between clusters needs to be factored in when making an ambulance reallocation decision. Our model permits consideration of travel distance between clusters.

The third problem concerns servicing casualties with different priorities. We formulate a two-priority, preemptive, single-server queueing model with cutoff queueing discipline. Under this discipline the server responds only to high priority customers until the queue length of the other class exceeds a threshold L. After that the server switches to handle only the low priority queue. Steady-state balance equations are established for this system. Then we use two-dimensional generating functions to obtain the average number of customers for each priority class. We then focus on the preemptive resume case. We develop methodologies to obtain the optimal cutoffs for the situation when the weights of both queues are constant and the situation when the weights change linearly with the queue lengths.

Vehicle Routing Problems and the Issues of Integrating Production and Distribution

Che-Fu Hsueh
National Central University, Taiwan
Huey-Kuo Chen (Advisor)

The vehicle routing problem has been and remains a rich topic for researchers and practitioners. In the real world, however, there exist many uncertain or constantly varying factors, which affect the operation of vehicle routing plan and result in the higher cost or failure. These factors may include occurrence of customer, quantity of demand, travel time, deterioration of commodities, etc.

This thesis explores the extensional issues of the vehicle routing problem with time windows in three parts. In the first part, the demand of customers and travel times may vary unexpectedly and a real-time, online operational vehicle dispatching and guiding system is established, taking the latest information of demand and traffic condition into account. The predictable travel-time pattern is also considered as a time-dependent function in this model.

In the second part, the travel time is considered as a random variable and the optimal routes are found in the pretrip planning. The researches considering the stochasticity of travel time are relatively rare in the past. In this part a stochastic programming model is established, and the objective is proved to have a lower bound. An optimal solution can be found by raising the lower bound. In addition, the stochasticity of the arrival time at a customer will not be affected by the accumulative stochasticity of all previous link travel time in this model.

The third part consists of integration of production scheduling and distribution for perishable goods. The perishable goods deteriorate as soon as they were produced and keep decaying when being delivered, resulting in deduction of revenue. The vehicle route as well as its delivery time and the production time should be considered in the phase of production scheduling, to reduce the deterioration and increase the profit.

The Value of Time in Freight Transportation

Jérôme Massiani
Université Paris XII–Val de Marne, France
Rémy Prud'homme (Advisor)

The purpose of this thesis is to investigate the trade-offs made by firms regarding the use of time in their operations. We concentrate on two different types of firms: producers and hauliers. These two types of firms make use of different types of operations: procurement, transformation, and distribution for producers; travel operations and nontravel operations (stocking, cross-docking, border crossings) for hauliers. We distinguish two different types of goods: generic goods and specific goods, the latest ones being produced based on consumers specifications. These different types of goods correspond to very distinctive relationships with time. Generic goods can be produced in advance, giving rise to optimal stock policy.

Transportation Science 41(3), pp. 415–429, © 2007 INFORMS

Specific goods can only be produced once the consumer and producer reach an agreement. This latest situation will give rise to a trade-off in which the cost of the operations will depend on the duration dedicated to such an operation.

We investigate the features of the equilibrium that will result from the profit maximizing behaviour of shippers and hauliers. We also show how such equilibrium will have the minimum (constrained) travel duration as a parameter. Subsequently, we can show the consequences of a reduction in this minimum duration.

The formal development that is achieved gives a formal expression to the willingness to pay of the different agents (shippers and hauliers) for a time saving. The next step is to provide a numerical estimation of the willingness to pay and an estimate of the whole benefits that the economic system obtains from a reduction in travel time. We make estimates using two different data sets and two different techniques. First, we apply an hedonic regression to Revealed Preferences data collected in the Nord Pas de Calais Region. Second, we implement a mixed logit estimate to SP data collected in three Italian regions.

Combinatorial Optimization Treatment of the Unary NP-Complete Disassembly Line Balancing Problem

Seamus M. McGovern
Northeastern University

Surendra M. Gupta (Advisor)

Managing technology and its growth requires the management of components in the "reverse" supply chain to adequately address products—especially consumer electronics such as cellular telephones and personal computers—at the end of the product lifecycle. The growing amount of waste created by products reaching the end of their useful life poses challenges for the environment, governments, and manufacturers. Logistics alternatives include reuse, remanufacturing, recycling, storage, and disposal. With disposal the least desirable, the first process required by the remaining alternatives is disassembly. Just as a manufacturing assembly line is considered the most efficient technology to assemble a product, the disassembly line is the preferred method of de-manufacturing. Efficient techniques are required to minimize the number of workstations and sequence end-of-life products on the line. Challenges lie in the fact that this problem tends to have high calculation complexities due to rapid growth in computational runtime with incremental increases in the number of parts. The added complications of disassembly also typically foster multiple objectives.

The disassembly line balancing problem seeks a demanufacturing sequence that is feasible, minimizes the number of workstations, minimizes total idle time, ensures similar idle times at each workstation, as well as other, disassembly specific concerns. With exhaustive search quickly becoming prohibitively large, methodologies from operations research, specifically combinatorial optimization, hold promise. The disassembly line balancing problem is described then defined mathematically and proven to belong to the class unary NP-complete. Disassembly line versions of exhaustive search, genetic algorithm and ant colony optimization metaheuristics, a greedy algorithm, and greedy/hill-climbing and greedy/2-optimal hybrid heuristics are presented along with a new uninformed search heuristic. Four instances (two case study instances from consumer electronics (a personal computer and a cellular telephone), an imaginary product case study, and a known-optimal solution (a priori) data set) are developed and used in numerical studies to illustrate the implementations and allow for quantitative and qualitative analysis and comparison for reverse logistics management.

Designing Aviation Security Systems: Theory and Practice

Laura A. McLay
University of Illinois at Urbana-Champaign

Sheldon H. Jacobson (Advisor)

Designing effective aviation security systems has become a problem of national concern. Since September 11, 2001, numerous changes have been made to aviation security systems. Many of these changes have been politically driven, rather than driven by coordinated, systematic analysis and planning, and as a result, the analysis of how aviation security systems operate continues to lag well behind their actual implementation.

This dissertation applies operations research methodologies to aviation passenger screening problems, an important and highly visible aspect of aviation

security. Effective passenger screening systems optimally allocate and use security assets and technologies to prevent terrorist attacks. Designing effective passenger screening systems becomes more critical as new screening technologies are made available and as security intelligence improves. Two basic passenger screening models are identified, and algorithms for obtaining optimal and near-optimal solutions for these problems are reported. One of the passenger screening models is extended to handle dynamic passenger arrivals. Exact and approximation algorithms are developed for three knapsack problem variations whose structures are similar to those of the passenger screening models. Finally, a cost-benefit analysis is performed to quantify the trade-offs between having effective, expensive screening technologies and accurate security intelligence. The key implication obtained is that investing in new, expensive technologies for preventing successful attacks is warranted only if security intelligence is accurate.

Next-generation aviation security systems need not merely be makeshift political solutions for mending complex problems; they can be the result of modeling, analysis, and planning. This dissertation provides a systematic approach for designing and analyzing aviation security systems, which provides insight into the operation of passenger screening systems and guidance for the design of next-generation aviation security systems.

Dynamic Programming Applied to a New Formulation of the Stochastic Truckload Routing Problem

Virginia M. Miori
Drexel University

Hande Y. Benson (Advisor)

This thesis focuses on vehicle routing and scheduling, more specifically the truckload routing problem (TRP). We include the consideration of time windows (TRPTW). In the TRPTW, each vehicle carries a single load that fulfills either the cubic volume or weight capacity of the equipment. The demand for these loads is stochastic.

The continuous flow TRPTW looks beyond traditional out and back routing by building extended routes that combine origin/destination pairs. Rather than returning directly to its domicile, a vehicle will instead proceed on to pick-up and deliver a subsequent load, which may take it even farther from the domicile. Each vehicle stays out on the road for as long as allowable by the various scenarios presented in the Department of Transportation (DOT) hours of service (HOS) restrictions.

This research presents an approach to the formulation of the TRPTW which allows the use of optimization through penalty methods and dynamic programming. The traditional lane definition (origin/destination pair) has been extended: We create triplets composed of two lanes in succession which represent (1) a loaded movement and (2) the necessary follow up movement to position for the next load. The triplet provides a natural way of addressing the stochastic demand. The termination point of a triplet in a route always corresponds to the initial point of the subsequent triplet.

The value of this research exists in the presentation of a new model formulation of the TRP which integrates count data analysis into the approach while taking into account real-world constraints. The movement of freight in the TRP is modeled more intuitively using triplets, resulting in greater applicability. The value of this approach is also seen in extensions to other areas of supply chain management and network modeling.

Supply Chain Design on a Line with Lateral Transshipments

Hussein Naseraldin
Technion–Israel Institute of Technology, Israel

Yale T. Herer (Advisor)

We conduct research in the management approach. We investigate a system that is comprised of retail outlets and customers, both of which are located on a homogenous finite line segment. In each period, the customers on the whole line pose a stochastic demand that is normally distributed. We analyze this stochastic inventory system in an infinite-horizon setting and use a cost-based performance measure.

Managing such a system entails determining various decisions on different planning horizons. Decisions such as the number and location of retail outlets are part of the strategic design of a sup-

Transportation Science 41(3), pp. 415–429, © 2007 INFORMS

ply chain. Decisions such as the inventory replenishment levels and the lateral transshipment quantities are part of the operational planning of a supply chain. Though integrating decisions that affect different planning horizons, e.g., strategic and operational decisions, necessarily leads to a better solution than if they were determined separately, this integration is rare. This rareness is due to the difficulties in modeling the interrelated nature of these decisions.

We propose models that integrate strategic and operational decisions and accurately depict the interrelation between these decisions. We integrate the decisions of number and location of retail outlets with the decisions of replenishment levels and lateral transshipment quantities. We start with a system that does not practice lateral transshipments, where we examine the impact that a simultaneous approach in decision-making has on the solution and objective spaces. We quantify the advantage of our approach in terms of the problem and cost parameters and explore the robustness of the objective function.

Afterward, we incorporate lateral transshipments in the model. We integrate the decision of number and location of retail outlets, replenishment levels, and lateral transshipment quantities. We explore the asymptotic behavior of this system and investigate the advantages from practicing lateral transshipment. Several managerial insights are gained.

Static and Dynamic Approaches for Solving the Vehicle Routing Problem with Stochastic Demands

Clara M. Novoa
Lehigh University

Robert H. Storer (Advisor)

This research focuses on the vehicle routing problem with stochastic demands (VRPSD). The VRPSD occurs in the delivery of industrial gases and the restocking of retail stores. Because exact customer demands are uncertain until the vehicle arrives at a customer, a planned route may fail. If a route failure occurs, different recourses can be taken at extra cost. For example, the remaining customers can be skipped, re-scheduled in extra trips, or served by the same vehicle after it returns to the depot to refill.

We study the VRPSD from two perspectives. The first (Part I of this dissertation) solves the VRPSD with a two-stage approach and new recourse actions that include deliveries completed by other vehicles and the re-scheduling of extra trips. The model is a stochastic program with integer recourse. We demonstrate that our new recourse decreases routing cost about 5% over the traditional "go to the depot and resume" recourse when the number of unsatisfied customers per route is low.

The second perspective (Part II of this dissertation) constructs the route dynamically in multiple stages using the most current information. We model the problem as a Markov decision process and use a rollout algorithm (RA) to compute the solution. We assume the traditional recourse supplemented by proactive returns to the depot. We extend previous work on the RA by investigating the effect of different base tours, cost-to-go estimation methods, look-ahead policies, and pruning schemes. We find that cost-to-go estimation via simulation reduces the CPU time by about 60% when compared to exact cost-to-go computation. We compare our best rollout solution to the perfect information case being the confidence interval for the overall mean difference (4.565%, 5.27%). This work also assesses the advantages of these stochastic approaches over following deterministic solutions from VRP and the benefits of merging the studied perspectives.

System-Level Stability and Optimality of Decentralized Supply Chains

Yanfeng Ouyang
University of California, Berkeley

Carlos Daganzo (Advisor)

This thesis focuses on decentralized multi-echelon supply chains where a sequence of suppliers act independently (i.e., placing orders from their neighbors based on private inventory and orders information) to serve a final customer. Experience shows that in decentralized chains fluctuations in order sequence are usually greater upstream than downstream. This phenomenon, often called "the bullwhip effect," is both prevalent and very costly. It is a manifestation of system instability.

Unlike other supply chain studies, which focus on specific demand scenarios, we show in this thesis that instability is an inherent property of the system, e.g., of the ordering policies used by the suppliers, but is independent of customer demand. We demonstrate this fact, both for deterministic and stochastic chains.

Randomness is assumed to arise from uncertain customer demands, and also from unpredictably varying factors in the operating environment such as transportation and production lead times.

This thesis presents analytical stability conditions for deterministically operated chains driven by arbitrary (deterministic or random) customer demands. These chains are modeled in the frequency domain as single-input, single-output (SISO) systems, and their stability is evaluated with the H-infinity norm. When the customer demand process is known and ergodic, we characterize the stream of orders placed at any stage of the chain, and give an exact formula for the variance. We also develop stability conditions allowing for uncertainty in the operating environment (random lead time, transportation loss, etc.). These chains are modeled in the time domain as Markovian jump linear systems. Our new results reveal how key features of a system (such as "lead times" and ordering policies) affect stability.

The results have also been used to develop effective policies based on advance demand information. We design contracting schemes that mitigate the bullwhip effect while allowing every supplier to act in its own self-interest. Market-based negotiation and contracting schemes have also been developed.

The analysis framework presented in this dissertation is general. Thus, can be used to improve the performance of multi-commodity supplier networks, manage freeway flow, coordinate transit vehicles, regulate air traffic, and control other interconnected systems.

Optimization Models and Solution Methods for Intermodal Transportation

Michael Berliner Pedersen
Technical University of Denmark, Denmark

Oli B. G. Madsen and Otto A. Nielsen (Advisors)

This thesis is composed of three papers, each dealing with different aspects in optimization of intermodal transportation. In addition to providing a context for the three papers, the summary introduces congestion and environmental issues within the transportation business and why the European Union sees the reestablishment of the rail sector in an intermodal setting as the solution to these problems.

The first paper presents a mathematical programming model to determine optimal timetables for public transit systems with respect to passenger transfer waiting time. By adopting a value of time cost measure the model minimizes the total sum of passenger transfer waiting time cost. The model is solved using a Tabu search heuristic on a large scale network instance taken from the public transit system of the greater Copenhagen area. The results show that there is a potential benefit for passengers accounting for 4 million Euros measured in value of time.

The second paper presents an optimization model to determine intermodal train schedules in a European setting. Trains are assumed to run on a rail infrastructure divided into train paths. Furthermore, the model includes terminal operations on an aggregated level to capture the transfer operations and costs at terminals. The model also takes into account the value of time cost for freight to determine the trade-off between low operational cost and short freight transit time. The model is solved using Xpress-MP's mixed-integer programming (MIP) solver, which did not prove to be an efficient solution method.

The third paper presents a generalized version of the model from the second paper. The model is characterized as the fixed-charged capacitated multicommodity network design model (CMND) with additional design balance constraints denoted. These constraints require an equal number of open design arcs entering and leaving the network nodes and can be interpreted as vehicle balance constraints. The model is solved by a multi-phased Tabu search heuristic. The algorithm is tested on generated instances, compared with solutions obtained with the Xpress-MP MIP-solver, and shows convincing results for eventual applications.

Robust Airline Fleet Assignment

Venkata L. Pilla
University of Texas at Arlington

Victoria Chung-Ping Chen and Jay M. Rosenberger (Advisors)

The fleet assignment model allocates a fleet of aircraft to scheduled flight legs in an airline timetable, typically published 90 days prior to departure. The robust fleet assignment model addressed in this dissertation uses a two-stage stochastic programming framework along with the Boeing concept of demand driven dispatch to swap crew-compatible aircraft closer to departure, when most of the demand

Transportation Science 41(3), pp. 415–429, © 2007 INFORMS

has been realized. Crew-compatible aircraft have identical cockpits, allowing an airline to swap aircraft without swapping crews. The two-stage formulation assigns crew-compatible aircraft in the first stage, 90 days out, so as to enhance the demand capturing potential of swapping in the second stage, about two weeks out. The stochasticity of the demand is modeled by different demand scenarios in the second stage, and the average over the scenarios estimates the expected profit. The methodology is tested on a real airline network with 50 stations and 2,358 legs.

Traditional two-stage stochastic programming uses a Benders' approach; however, for large-scale problems, this can be slow to converge. This dissertation develops a two-phase design and analysis of computer experiments (DACE) approach to reduce the computation involved in conducting the optimization. The DACE phase uses first stage constraints in a multistep process to construct an experimental design within the feasible region, then builds a statistical model that approximates the expected profit function in the first stage of the stochastic program. The optimization phase solves the two-stage problem using the DACE expected profit approximation in the restricted master problem instead of solving many subproblems in every iteration. This greatly speeds up the optimization, compared to Benders', because the computation of the subproblems is shifted to the DACE phase. Thus, future work will focus on reducing the computation of the DACE phase using dimensionality reduction methods from statistical data mining.

A Hybrid Fixed and Flexible Transportation Service: Description, Viability, Formulation, Optimization and Heuristic

Luca Quadrifoglio
University of Southern California

Maged Dessouky (Advisor)

The mobility allowance shuttle transit (MAST) service is a new concept in transportation that merges the flexibility of demand-response transportation (DRT) systems with the low cost operability of fixed-route bus systems. In a MAST system vehicles follow a base fixed-route composed by a few mandatory checkpoints conveniently located at major connection points. Given an appropriate slack time, vehicles are allowed to deviate from the fixed path within a proper service area to pick up and drop off passengers at their desired locations.

The purpose of this research is to address the gap in the research community by studying this hybrid fixed and flexible type of service, providing insights on its challenges, and foreseeing its performance in large scale use as an alternative to conventional public services.

The system is defined as viable if the longitudinal velocity along the primary direction of the service is higher than a minimum threshold value to maintain the service attractive to customers. By using continuous approximations we develop a relationship between this velocity and the demand to assess the viability and aid in the setting of the main parameters of the service.

For the static operating scenario we provide a mathematical formulation of the MAST system as an integer linear program and we aim to find the optimal schedule. Because of the combinatorial nature of the problem we develop a set of valid inequalities to increase the lower optimality bound and efficiently speed up the search for the optimal solution.

For the dynamic operating scenario we develop a customized insertion scheduling algorithm, which includes control parameters to prevent the "wild" consumption of the slack time and significantly improve the performance of the algorithm. A comparison of MAST services versus conventional fixed-route systems show that are competitive with conventional ones and perform better under certain demand distributions.

Algorithms for Procurement in Transportation Networks

Samik Raychaudhuri
University of Wisconsin, Madison

Dharmaraj Veeramani (Advisor)

With the advent of the Internet, companies are getting interested in e-procurement to reduce cost and time of procurement operations. In this research, we address the procurement problem in the context of transportation networks. We consider two types of network procurement problems. For some industries, e.g., oil, electricity or natural gas supply chain, links of the networks (namely pipelines, trans-

mission lines) are the object of procurement. We call this scenario as edge procurement because edges are procured to facilitate formation of a network. For other types of businesses, a predefined underlying network is specified by the sources and destinations of transportation of goods or services. A route in such a network is a set of edges leading from the source node to the destination node. For these industries, the procurement process includes selection of transportation agents and routes. We refer to this as transportation service procurement.

In the context of edge procurement, our research considers a distribution network model where each edge is owned by a selfish owner and a customer wants to supply a product or a service to each of the nodes in the network using a set of edges that results in the lowest possible distribution cost, i.e., a minimum spanning supply tree. The owners incur costs when their edges are used by the customer, and this cost information is privately known only to respective edge owners. We characterize the equilibrium price space of such an economy and analyze a descending price auction mechanism that discovers one of the Walrasian equilibrium prices satisfying certain conditions. We then study the strategic behavior of the owners in such an auction and show that following a greedy strategy is a Nash equilibrium for the owners.

Transportation service procurement auctions are typically characterized by a limited number of rounds and low profit margin. In this context, our research addresses the bundle selection and bid determination problem on behalf of single-lane truckload carriers and combinatorial bids in the successive rounds of an auction, which maximizes the profit for the carriers. First, we develop a framework for calculating the cost of providing transportation service for a lane or a set of lanes. We also develop a methodology for calculating winning probabilities of a bundle of lanes, given the bid price and historical information about the same lane. Based on the cost framework and the methodology for calculating winning probability, we develop a model using constraint programming which generates bundles with a high probability of winning. We then use a mixed integer programming (MIP) formulation to select a set of bundles that maximizes the expected profit, given a specific set of profit multipliers. After we select the set of bundles to bid on at a particular round, we compute a profit factor for each of the bundles using the industry average price for a lane or lane bundle from the historical figures and the winning bids for the past round of an auction. We develop two methods to update the estimate of the present market rate of a lane or a lane bundle, and two methods to evaluate a profit factor, which will maximize the expected profit of the carrier. We perform numerical experiments to validate our results, and to determine which strategy combination works best for the carrier.

We also study an extension of the above process in which we relax the assumption of a carrier winning all the truckloads associated with the lane. A simple modification of the above models is used to analyze this case.

Information and Decentralization in Inventory, Supply Chain, and Transportation Systems

Guillaume Roels
Massachusetts Institute of Technology

Georgia Perakis (Advisor)

This thesis investigates the impact of lack of information and decentralization of decision-making on the performance of inventory, supply chain, and transportation systems. In the first part of the thesis, we study two extensions of a classic single-item, single-period inventory control problem: the "newsvendor problem." We first analyze the newsvendor problem when the demand distribution is only partially specified by some moments and shape parameters. We determine order quantities that are robust, in the sense that they minimize the newsvendor's maximum regret about not acting optimally, and we compute the maximum value of additional information. The minimax regret approach is scalable to solve large practical problems, such as those arising in network revenue management, since it combines an efficient solution procedure with very modest data requirements. We then analyze the newsvendor problem when the inventory decision-making is decentralized. In supply chains, inventory decisions often result from complex negotiations among supply partners and might therefore lead to a loss of efficiency (in terms of profit loss). We quantify the loss of efficiency of decentralized supply chains that use price-only contracts under the following configurations: series, assembly, competitive procurement, and competitive distribution.

In the second part of the thesis, we characterize the dynamic nature of traffic equilibria in a

Transportation Science 41(3), pp. 415–429, © 2007 INFORMS

transportation network. Using the theory of kinematic waves, we derive an analytical model for traffic delays capturing the first-order traffic dynamics and the impact of shock waves. We then incorporate the travel-time model within a dynamic user equilibrium setting and illustrate how the model applies to solving a large network assignment problem.

Activity-Based Modelling of Household Travel

Matthew J. Roorda
University of Toronto, Canada
Eric J. Miller (Advisor)

This thesis reports on three areas of innovative, practical research. First, a major in-depth longitudinal data collection effort has been undertaken, which has provided groundbreaking information on the travel/activity scheduling process. The three-wave panel survey involved a variety of methodological techniques aimed at capturing different components of the activity scheduling process. Analysis of these new data has significantly improved our understanding of behavioural processes surrounding household activity scheduling. Second, a new rule-based technique for simulating activity schedules has been developed. The model predicts participation in individual and joint activities within the household context and organizes those activities into consistent schedules. The model includes a method for resolving conflicts that arise when multiple activities are generated for the same point in time. Third, a tour-based mode choice model that incorporates crucial household interactions including vehicle allocation, ridesharing to joint activities, and dropoff/pickup scenarios has been estimated, and applied in a large-scale microsimulation framework. A random utility approach is adopted in this model to determine the choice among modal options. The model's complex nesting structure, however, requires that a simulation approach be used to estimate the likelihood associated with a parameter set. A genetic algorithm is used to determine the maximum likelihood parameter set.

The latter two pieces of work have been developed into a modelling software package entitled the Travel Activity Scheduler for Household Agents (TASHA). The TASHA modelling system is more behaviourally valid than current state-of-practice models, provides more precise outputs with no additional input requirements, and has strong potential for the analysis of contemporary demand-oriented policies. With some additional research, testing, and development, it is felt that this approach can evolve to become the "next generation" of travel demand modelling for major urban centres such as the greater Toronto area.

Heuristic and Exact Algorithms for Vehicle Routing Problems

Stefan Ropke
University of Copenhagen, Denmark
David Pisinger (Advisor)

This thesis focuses on the pickup and delivery problem with time windows (PDPTW). The first part of the thesis presents a very capable metaheuristic for solving the problem. We show how many classes of the vehicle routing problem (VRP) from the literature can be transformed to a PDPTW. Combining these transformations with the metaheuristic gives a very general heuristic that solves the most famous variants of the VRP and is competitive with recent, more specialized, heuristics for VRP variants. The generality and robustness of the heuristic makes it well suited for application to real life problems.

The heuristic is based on the large neighborhood search (LNS) framework. The thesis presents important improvements to the framework, which result in a more robust heuristic. These improvements lead to a new metaheuristic framework denoted adaptive large neighborhood search that can be applied to problem types outside the VRP domain.

The second part of the thesis studies exact methods for solving the PDPTW. Two methodologies are used: branch and cut, and branch and cut and price. Two new integer programming models are proposed to apply branch and cut to the problem. Compared to earlier models, these are more compact in terms of

the number of variables, but require an exponential number of constraints. Several new classes of valid inequalities are proposed. This leads to a branch-and-cut algorithm that is able to solve significantly larger instances compared to another recent branch-and-cut approach (96 requests versus 36 requests).

Four set-partitioning relaxations of the PDPTW are studied in the branch-and-cut-and-price approach. Some are new, and some have been studied in previous literature. This leads to a unique comparison of different relaxations for the PDPTW. Valid inequalities from the branch-and-cut approach are used to strengthen these relaxations. Overall the branch-and-cut-and-price approach is capable of solving more instances to optimality compared to the branch-and-cut approach.

Branch-and-Price Algorithms for Vehicle Routing Problems

Matteo Salani
Universitá degli Studi di Milano, Italy
Giovanni Righini (Advisor)

This thesis considers vehicle routing problems (VRPs) with additional constraints stemming from real applications and the design of branch-and-price algorithms to solve them to optimality. Because the computational effort required by branch-and-price algorithms is mainly affected by the solution to the pricing subproblem, the main part of the thesis is devoted to dynamic programming algorithms for pricing. In particular we present some new ideas to improve dynamic programming algorithms, namely bi-directional and bounded dynamic programming and decremental state space relaxation. Bi-directional and bounded dynamic programming allows us to exploit the symmetry of the pricing subproblem. This is a resource constrained shortest path problem: The route of the vehicle is built in two opposite directions, implicitly enumerating two smaller spaces instead of a larger one. We combine ideas from branch and bound and dynamic programming to make the resulting algorithm efficient. Decremental state space relaxation is a new technique that solves the pricing subproblem iteratively, tightening its initial relaxation only as needed to avoid cycles. Branching strategies and primal heuristics are also presented. Computational experiments are done for the capacitated vehicle routing problem (CVRP), the vehicle routing problem with distribution and collection (VRPDC), and the capacitated vehicle routing problem with time windows (CVRPTW) on a set of benchmark instances from the VRP literature. Some of these problems have been solved to proven optimality for the first time.

A Network Design Model with Lead Time and Safety Stock Considerations

Karthik Sourirajan
Purdue University
Leyla Ozsen and Reha Uzsoy (Advisors)

A company's success in getting business is significantly affected by the operational aspects of lead times and service levels faced by its customers. By including such aspects into the distribution network design models, we can potentially achieve a design that is more effective in capturing customer expectations than the traditional design models. The main contribution of this research is the explicit modeling of nonlinear relationships between the demand flows in the network, lead times, and safety stocks that allows us to capture the trade-off between safety stock risk-pooling benefits and lead time congestion costs. We consider a two-stage supply chain in which a production facility replenishes a single product at a given set of retailers. The objective is to locate distribution centers such that the sum of location and expected inventory (pipeline and safety stock) costs is minimized. We develop a Lagrangian heuristic to obtain near-optimal solutions in reasonable CPU times for large problems. We then develop genetic algorithms for the model and compare their performance to the Lagrangian heuristic. A novel chromosome representation that combines binary vectors with random keys is shown to provide near-optimal solutions. Finally, we develop a multi-product network design model to capture the effects of the demand mix across the various products on the costs. We develop a Lagrangian heuristic and a genetic algo-

Transportation Science 41(3), pp. 415–429, © 2007 INFORMS

rithm to approximately solve the model. We find that while congestion costs can be thought of as a global network property that depends on the total demand flows into a distribution center, the risk-pooling benefits can be thought of as a local network property that depends on the demand allocations of individual products. Thus, there is a trade-off between these global and local network properties that influences the distribution network design and is captured by our model.

Analytical Approaches to Railroad and Rail-Truck Intermodal Transportation of Hazardous Materials

Manish Verma
McGill University, Canada

Vedat Verter (Advisor)

Hazardous materials are potentially harmful to people and environment due to their toxic ingredients. Although a significant portion of dangerous goods transportation is via railroads, prevailing studies on dangerous goods transport focus on highway shipments. We present an analytical framework that incorporates the differentiating features of trains in the assessment of risk. Each railcar is a potential source of release, and hence risk assessment of trains requires representation of multiple release sources in the model. We report on the use of the proposed approach for the risk assessment of the Ultra-train that passes through the city of Montreal daily. The risk assessment methodology is then used to model the operations of freight trains in a network, wherein freight involves both hazardous and regular cargo. A bi-objective optimization model, distinct from conventional models, with an intermediate horizon is presented. A memetic algorithm based solution methodology was developed to solve the tactical planning problem. A realistic size case example based in Canada was used to gain numerical and managerial insights into the problem and the solution algorithm. In an effort to answer the question "rail or truck: which is safer?," and also to combine the economies of trains and efficiencies of trucks, we deal with rail-truck intermodalism for hazardous and nonhazardous cargo. A delivery-time based risk-cost analysis is presented for determining the best routes for these intermodal shipments. Two special cases and a general case of rail-truck intermodal transportation models, driven by the element of "time," are also presented.

Reliability of Railway Systems

Michiel Vromans
Erasmus University Rotterdam, The Netherlands

Jo van Nunen, Leo Kroon, and R. Dekker (Advisors)

The on-time running of trains is the most important goal of Netherlands Railways; it is an important characteristic to attract customers to railway transport. Moreover, the on-time running is also a central issue in contracts with the government and passenger organizations.

This leads to the main goal of this thesis: to develop rules and instruments for supporting the generation of more reliable timetables.

The first three chapters give an introductory description of the environment in which the research was conducted. In these chapters theory and practice, timetable construction and timetable evaluation, and deterministic planning problems and stochastically disturbed operations are discussed.

Running time supplements are included in public transport timetables to be able to absorb small delays. Chapter 4 describes this issue and concludes that the usual proportional allocation of the supplements in The Netherlands is not optimal. It is profitable to have less supplement on the first and last parts of a train line, and relatively more in the middle.

Small headways increase the probability of delay propagation. The headways can be increased by homogenizing the railway traffic, i.e. reducing the speed differences. Two new measures, sum of the shortest headway reciprocals and sum of arrival headway reciprocals, are introduced in Chapter 5. These measures not only quantify the heterogeneity of the traffic, but also provide an indication of delay propagation.

In Chapter 6 an innovative stochastic optimization model is presented. The model integrates a timetabling model with a timetable evaluation model. The symbiosis of these two parts leads to an optimal timetable, given the model environment. A case study shows that a delay reduction of up to 30% is possible within the model settings.

To the best of our knowledge, this is the first timetabling model that explicitly takes delay propagation into account, and directly optimizes the timetable without cumbersome trial-and-error procedures.

Finally, Chapter 7 summarizes the results of the thesis and raises some additional research questions.

Online Optimization in Routing and Scheduling

Michael R. Wagner
Massachusetts Institute of Technology

Patrick Jaillet (Advisor)

In this thesis we study online optimization problems in routing and scheduling. An online problem is one in which the problem instance is revealed incrementally. Decisions can (and sometimes must) be made before all information is available. We design and analyze (polynomial-time) online algorithms for a variety of problems. We use worst-case competitive ratio (and relaxations thereof), asymptotic, and Monte Carlo simulation analyses in our study of these algorithms.

The focus of this thesis is on online routing problems in arbitrary metric spaces. We begin with online versions of the traveling salesman problem (TSP) and the traveling repairman problem (TRP). We then generalize these basic problems to allow for precedence constraints, capacity constraints, and multiple vehicles. We give the first competitive ratio results for many new online routing problems. We then consider resource augmentation, where we give the online algorithm additional resources: faster servers, larger capacities, more servers, less restrictive constraints, and advanced information. We derive new worst-case bounds that are relaxations of the competitive ratio. We also study the (stochastic) asymptotic properties of these algorithms—introducing stochastic structure to the problem data, unknown and unused by the online algorithm. In a variety of situations we show that many online routing algorithms are (quickly) asymptotically optimal, almost surely, and we characterize the rates of convergence.

We also study classic machine sequencing problems in an online setting. Specifically, we look at deterministic and randomized algorithms for the problems of scheduling jobs with release dates on single and parallel machines, with and without preemption, to minimize the sum of weighted completion times. We derive improved competitive ratio bounds and show that many well-known machine scheduling algorithms are almost surely asymptotically optimal under general stochastic assumptions.

For both routing and sequencing problems, we complement these theoretical derivations with Monte Carlo simulation results.