Konrad-Zuse-Zentrum
für Informationstechnik Berlin

Takustraße 7
D-14195 Berlin-Dahlem
Germany

Ralf Borndörfer

Discrete Optimization in Public Transportation

Paper presented at the 1st Indo-US Symposium on "Advances in Mass Transit and Travel Behaviour Research (MTTBR-08)“, February 12-15, 2008, IIT Guwahati, Assam, India.

ZIB-Report 08-56 (February 2008)

Discrete Optimization in Public Transportation

Ralf Borndörfer
Zuse-Institute Berlin, Takustr. 7, 14195 Berlin, Germany, borndoerfer@zib.de

Abstract. The mathematical treatment of planning problems in public transit has made significant advances in the last decade. Among others, the classical problems of vehicle and crew scheduling can nowadays be solved on a routine basis using combinatorial optimization methods. This is not yet the case for problems that pertain to the design of public transit networks, and for the problems of operations control that address the implementation of a schedule in the presence of disturbances. The article gives a sketch of the state and important developments in these areas, and it addresses important challenges. The vision is that mathematical tools of computer aided scheduling (CAS) will soon play a similar role in the design and operation of public transport systems as CAD systems in manufacturing.

1 Introduction

Public transit is an area that is generally well suited to the application of mathematical optimization methods. In many municipalities, public transport is organized in a centralized, completely scheduled way, time and (at least in principle) data for planning is available, and, last not least, a number of important planning problems are close to classical optimization problems such as, e.g., multi-commodity flow problems. In fact, the influence goes also the other way round, i.e., scheduling problems in public transit have been among the driving forces behind the development of several areas of applied mathematics such as discrete and combinatorial optimization and integer programming for the last decades, see, e.g., Ball, Magnanti, Monma & Nemhauser (1995).

Service Design	Operational Planning	Operations Control
Network Design	Vehicle Scheduling	Vehicle Dispatching
Line Planning	Duty Scheduling	Crew Dispatching
Timetabling	Duty Rostering	Depot Management
Fare Planning		

Table 1: Planning problems in public transport.

There is a standard decomposition of the planning process in public transit into three major areas, which are in turn subdivided into a sequence of individual problems, see Table 1. On a long and medium time scale, *service design* deals with decisions about the level and type of service that is offered to the public in terms of the layout of the network, the line system, the timetable, and the fares; the aim here is to find a good compromise between quality and cost. *Operational planning* deals with the implementation of the resulting network at minimum cost with the available vehicles and crews; the time horizon is medium term. *Operations control* monitors the execution of the schedules in the presence of disruptions such as vehicle breakdowns, no-shows of crews, etc.; the time horizon is short term to online, the main objective is to get things done, but there is, of course, also an eye on costs.

Old Computer	New Computer	Speedup	Old Code	New Code	Speedup
Sun 3/50	Pentium 4, 1.7 GHz	800	XMP	Cplex 1.0	4.7
Sun 3/50	Compaq ES 40, 667 MHz	900	Cplex 1.0	Cplex 5.0	22.0
Intel 386, 25 MHz	Compaq ES 40, 667 MHz	400	Cplex 5.0	Cplex 7.1	3.7
IBM 3090/108S	Compaq ES 40, 667 MHz	45	XMP	Cplex 7.1	960.0

Table 2: Progress in linear programming 1987-2000 according to Bixby (2002).

Planning problems in public transit can be seen as *network optimization problems* that can be addressed with mathematical methods. At all times, researchers and practitioners have applied the methods available to solve as realistic and comprehensive models as possible. Today, progress in mathematical optimization and computing technology allows the solution of scheduling problems at unprecedented sizes and level of detail. A striking example to illustrate this progress is linear programming, the most important base technology in discrete optimization. Bixby (2002) has estimated a speedup factor of roughly one million in about a decade, of which a factor of 1.000 is due to better computers, and another factor of 1.000 to better algorithms, see Table 2. "A model that might have taken a year to solve ten years ago, can now solve in less than 10 seconds", he concludes.

These advances carry over to individual applications in public transit, in particular, in *operational planning*, such that we can produce today, for the first time, high quality solutions for a number of scheduling problems, most notably vehicle and crew scheduling problems; sometimes we can solve them even to proven optimality. Not surprisingly, these methods have made it into commercial scheduling systems and are used in many public transit companies to solve, on a routine basis, large-scale vehicle and crew scheduling problems, see, e.g. Daduna & Paixão (1995) for an overview. A natural next step is to try to overcome the artificial sequential decomposition of the planning process by *integrated planning*, i.e., by solving two or more optimization problems, that have hitherto been solved one after the other, together. This introduces additional degrees of freedom, but clearly also increases the computational complexity. In some cases, however, e.g., in regional scenarios, integrated scheduling is indispensable, because a sequential approach just doesn't produce feasible solutions. And in fact, the first successful steps into integrated scheduling have already been done, see Section 2.3 below.

In the areas of *service design* and *operations control*, mathematical optimization is more in a state of research. One reason is that the problems are complex and less clear defined. There are, e.g., usually competing objectives, such as quality and cost, there are behavioral and game theoretic aspects involved in the treatment of demand, etc. Therefore, a combination of models from different fields is needed, e.g., combinatorial methods to deal with the network, and behavioral models to predict the demand. Another reason is the lack of, or impossibility, or difficulty, or cost to obtain the data that one would like to use in such models, e.g., about travel choice. Similar arguments can be put forth for the area of operations control, which features additional constraints on the time that is available to make a decision.

However, considering the size of investments into public transit, their long term nature, their impact on the development of a city or region, and the complexity of the decisions, it seems that the application of advanced mathematical optimization models and methods for these problems is highly desirable. Discussing important developments in operational planning and service design, this paper makes a point that such methods can indeed make a contribution. Our vision is that mathematical methods of *computer aided scheduling* (CAS) will, in the future, play exactly the same role in the design and operation of public transportation systems as methods of computer aided manufacturing (CAD) in product design and assembly.

A broad overview on the development and the state of the art in mathematical optimization of public transit is given in the proceedings volumes of the triannual international conferences on computer-aided scheduling of public transport (CASPT), see Daduna & Wren (1988), Desrochers & Rousseau (1992), Daduna, Branco & Paixao (1995), Wilson (1999), Voss & Daduna (2001), and Hickman (2008).

2 Developments

We discuss in this section recent developments in operational planning and service design. Sketching five problems, namely, vehicle scheduling, duty scheduling, integrated vehicle and duty scheduling, line planning (sometimes also called transit network design), and fare planning, we attempt to illustrate the status quo and the perspectives of this field.

2.1 Vehicle Scheduling

This is perhaps the prime example for the successful application of mathematical optimization problems to scheduling problems in public transport. Given a number of vehicles of different types, a set of timetabled trips for a day of operation, and a set of possible deadhead trips that connect them, the problem is to construct a minimum cost set of vehicle rotations that cover all timetabled trips, such that each trip is covered by a vehicle of compatible type, see Figure 1 (left). This problem can directly be translated into a multi-commodity flow problem, one of the classical problems of combinatorial optimization, and formulated as an integer program, see, e.g., Desrosiers, Dumas, Solomon & Soumis (1995). The resulting models can be large. E.g., the vehicle scheduling problem that arises at Berliner Verkehrsbetriebe (BVG), the public transport company of Berlin, features 28.000 timetabled trips and almost 100 million possible deadhead trips, see Löbel (1997). Note, however, the important point that these scheduling graphs are acyclic, because vehicle trips proceed in time. Using methods that exploit the special network structure of the problem and deal with the underlying graphs and linear and integer programs dynamically by generating the relevant parts on the fly, such problems can today be solved on an ordinary desktop computer to proven fleet optimality in less than 30 minutes, see again Löbel (1997). One can, in this sense, therefore consider vehicle scheduling as a "solved problem". The community is now working on extended models that include multiple days of operation, maintenance constraints, or additional degrees of freedoms to slightly shift timetabled trips.

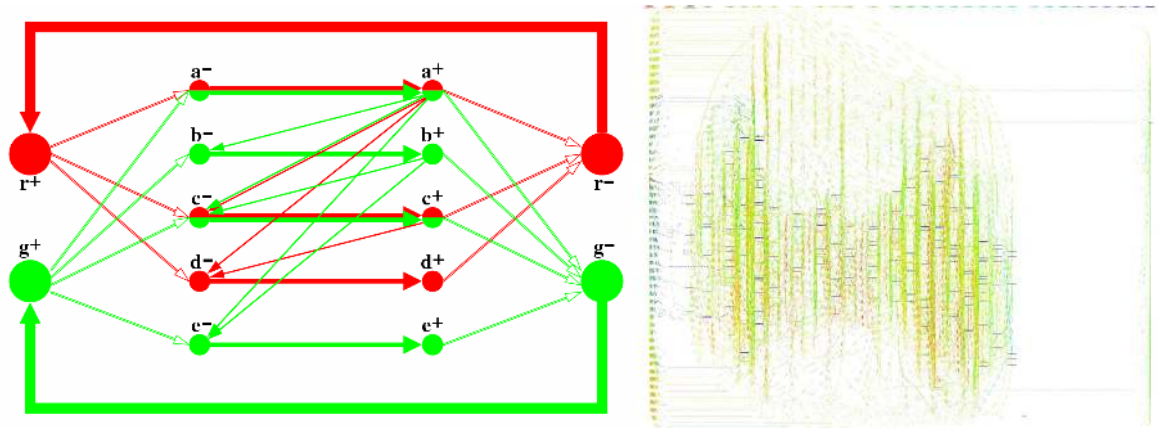

Figure 1: A vehicle scheduling graph for a 2-depot problem (left) and a duty scheduling graph (right); note the morning and the afternoon peaks on the right.

2.2 Duty Scheduling

A problem that looks very similar at first sight is the assignment of drivers to vehicle rotations. Cutting the rotations into elementary pieces, the so-called duty elements, one can again construct a scheduling graph that consists of all relevant units of work, connected by all possible transitions. The problem is then to cover the duty elements by paths that corresponds to legal drivers duties, in such a way as to minimize costs, see Figure 1 (right). The big difference to vehicle scheduling is the construction of the duty-paths, which must correspond to many legal restrictions, union agreements, and company regulations. In contrast to vehicle scheduling, where the legality of a rotation can be decided locally on the basis of feasible connections, compliance with, e.g., a break rule cannot be decided locally, but depends on the form of the entire duty. This difference has severe mathematical consequences, because the complexity of these rules make it essentially impossible to work with a duty scheduling model that is based on variables that correspond to arcs in the scheduling graph, and to express rules in terms of constraints. Instead, the standard approach to duty scheduling today is to work with so-called set *partitioning* or *set covering* models, that use a variable for each possible duty. As there is an astronomical number of these duties even for medium sized problems, such models can only be treated using solution techniques that generate the relevant duties dynamically. These are *column generation* methods, which also have also undergone a dramatic improvements in the last decades, see Desrosiers, Dumas, Solomon & Soumis (1995) and Barnhart, Johnson, Nemhauser, Savelsbergh & Vance (1994). Today, state-of-the-art systems can solve large scale problems involving an entire depot of a public transportation in a couple of hours, see, e.g., Borndörfer, Grötschel & Löbel (2003). Research in this area is concentrating on large-scale techniques such as stabilization of the dual variables, aggregation, special heuristics, etc. in order to solve bigger scenarios faster.

2.3 Integrated Vehicle and Duty Scheduling

In countries with high wages such as Germany, crew costs dominate vehicle costs by a ratio of 2:1, see Leuthardt (1998). It therefore seems not to be a good idea to schedule vehicles first, and crews second, because decisions that seem to be good in the first step can turn out to be bad later. But the problem is not only cost. For regional carriers, sequential vehicle and duty scheduling does in general not work, because very long vehicle rotations will leave the drivers out in the countryside at the point where a break has to be taken or the maximum shift length is reached. In such scenarios, vehicle rotations and driver duties must be constructed simultaneously.

From a modeling point of view, this is easy. One just takes a vehicle scheduling model and a duty scheduling model, and glues them together using coupling constraints on the deadhead trips, that force a driver to be used whenever a vehicle is used and vice versa. The problem is, of course, to solve such a model, because it is, in general, very large. In fact, there is one coupling constraint for every deadhead trip, such that even a medium sized regional scenario can easily feature half a million of coupling constraints. The method that suggests itself in this case is to use a Lagrangean relaxation of the coupling constraints, which decomposes the problem into a vehicle and a duty scheduling subproblem, which can be attacked with the standard methods. The main problem that remains is the multiplier update, which requires a method that converges fast, such not too many major iterations are necessary, and the construction of a feasible solution of good quality. A tool that has proved to work in this situation is to use a method from non-differentiable optimization, the bundle method, which converges much faster than standard subgradient methods, has desirable numerical features such as the possibility to reduce the problem dimension using duality, and an

automatic stabilization of the dual variables, and which fits well with specialized branching heuristics such as rapid branching. Using such methods, it is today possible to do an integrated vehicle and duty optimization of medium sized scenarios involving more than 3.500 timetabled trips, see Weider (2007). The sketched technique applies to coupled systems in general, such it is highly likely that we will see further successful applications of integrated scheduling in the future such as, e.g., integrated timetabling and vehicle scheduling, or integrated duty scheduling and rostering.

2.4 Line Planning (also called Transit Network Design)

Given the progress and success in mathematical vehicle and duty scheduling, it is a natural idea to try to extend these methods and apply them to the construction of a public transportation network itself. One possible problem is line planning, i.e., the construction of a set of lines, and corresponding frequencies, to be able to service some demand in a given time period. The construction of a line in an infrastructure network can again be reduced to a problem of constructing a path in a scheduling graph, similar to, say, duty scheduling, see Figure 2 (left), even if some technical difficulties come up, e.g., that the scheduling graph is no longer acyclic. There are however, some more severe issues. The first is that line planning has at least two objectives, namely, quality and cost, which are contradictory, i.e., line planning is a multiple criteria optimization problem. Moreover, service quality influences travel choices, because, e.g., passengers will not accept long or uncomfortable trips. This means that line planning does not only deal with the construction of lines, it should include a model of passenger behavior, which immediately makes the problem much more complicated. Today, integer programming approaches to line planning can deal with the combinatorial aspects of line construction and passenger routing including "fairness constraints" on the lengths of trips, see, e.g., Borndörfer, Grötschel & Pfetsch (2007). All these model assume, however, that the demand is fixed, i.e., passengers do not quit or enter public transit if the service level decreases or increases. Using such approaches, it is possible to construct optimized line systems for medium sized scenarios, see again Borndörfer, Grötschel & Pfetsch (2007). In this way, one can, e.g., investigate the trade-off between service quality and costs at the model's level of detail. This can already be a valuable contribution.

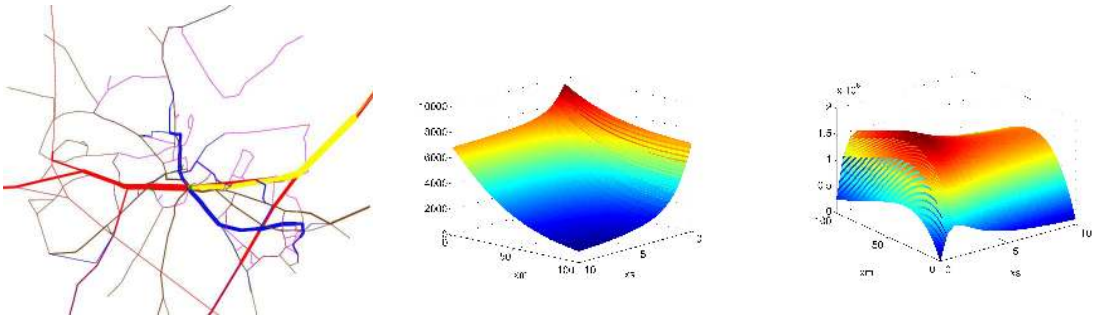

Figure 2: The line system of the city of Potsdam in Germany (left), a demand function depending on two fare variables (middle), and the corresponding revenue function (right).

2.5 Fare Planning

This is an example of an optimization problem that explicitly addresses forecasts on the behavior of users as an integral part of an optimization model. Fare planning is particularly suited for this type of investigation, because this problem involves very few variables, e.g., a price for a standard ticket and for a monthly ticket, see Figure 2 (middle and right), and, apart from user behavior, also few constraints, e.g., on maximum fares or subsidies. One well accepted way to forecast transportation demand are discrete choice logit models. They can forecast substitution effects. A technical pro is that the demand can be expressed in terms of closed formulas. Using these expressions, one can set up non-linear optimization models to maximize the revenue, profit, demand, or welfare, and it turns out that these models can be solved numerically using standard non-linear optimization solvers, see Borndörfer, Neumann & Pfetsch (2007). At present, the models use relatively simple estimates on, e.g., operation costs. But they give a very fine level of control over fares. The ultimate goal is, of course, to merge them with combinatorial models for line planning and network design.

3 Challenges

The discussion of the previous section has already indicated that *model integration* is one of the major challenges of public transport optimization. Starting with purely combinatorial models, e.g., in vehicle and duty scheduling, there is a need to integrate stochastic and combinatorial models to combine network design and demand forecast, and nonlinear model components to deal with inherently nonlinear quantities such as the revenue (revenue is price times demand (which depends on the price)). A very challenging topic would be

to investigate ways to use advanced behavioral models, such as activity based models, in an optimization model. Model integration is, in general, of course an extremely difficult task. For special problems in public transport, however, first successful steps have already been undertaken. Model integration is, in particular, a precondition for the successful treatment of problems in service design. We also mention that the data basis for service design must be improved.

Similar points can be made for the area of operations control. Dealing with schedule disruptions is clearly a topic that is most relevant to the customer. Ironically, the problem of disruptions is aggravated by the success in operational planning, which tends to eliminate buffers from the schedules, such that a plan, that looks good on paper, may turn out to be very expensive in practice. This is clearly not the intended effect, and there are currently a number of approaches to deal with these problems under investigation, namely, *online optimization*, *robust optimization*, and *stochastic optimization*. Online optimization works under a zero knowledge assumption; a main problem with this approach is that theoretical and empirical performance of online algorithms is often not related. On the contrary, stochastic optimization assumes complete knowledge on the likelihood of future events, which may be hard to obtain. The resulting models are theoretically satisfactory, but often extremely large scale and difficult to solve. Robust optimization attempts a compromise between these two extremes; there is, however, at present no general theory about what a robust analogue of a combinatorial problem is. In this area, substantial research is needed.

Finally, we are missing something if our models do not take *competition* into account. The European Union, e.g., is following a deregulation policy in the area of public transport, and there is already competition in parts of South America and Asia. Deregulation will change the organization of the planning process, and it will bring game theoretic aspects into play. We do not have a theory to find the right compromise between a monopolistic and a market organization in public transport. The research community can make a contribution to this discussion by proposing viable models and methods.

4 References

Ball M.O., Magnanti T.L., Monma C.L., Nemhauser G.L. editors (1995). "Network Models". Handbooks in Operations Research and Management Science, Vol. 7, Elsevier.

Barnhart C., Johnson E.L., Nemhauser G.W., Savelsbergh M.W.P., Vance P. (1994). "Branch-and-Price: Column Generation for Solving Huge Integer Programs“. In Birge J.R., Murty K.G. (Eds.) *Mathematical Programming: State of the Art*. Univ. of Michigan, pp. 186-207.

Bixby R.E. (2002). *Solving Real-World Linear Programs: A Decade and More of Progress*. Oper. Res. 50(1), pp. 3-15.

Borndörfer R., Grötschel M., Löbel A. (2003). "Duty Scheduling in Public Transit". In Jäger W., Krebs H.-J. (Eds.) Mathematics – Key Technology for the Future. Springer, Berlin, pp. 653-674.

Borndörfer R., Grötschel M., Pfetsch M.E. (2007). *A Column-Generation Approach to Line Planning in Public Transport*. Transp. Sci. 41(1), pp. 123-132.

Borndörfer R, Neumann M., Pfetsch M.E. (2007). "Optimal Fares for Public Transport". In Waldmann K.-H., Stocker U.-M. (Eds.) *Proc. Operations Research 2005*. Springer, Berlin, pp. 29-36.

Daduna J.R., Branco I., Paixão J.M.P. editors (1995). "Computer-Aided Transit Scheduling. Proceedings of the 6th Int. Conf. On Computer-Aided Scheduling of Public Transport“. LNEMS, Springer, Berlin.

Daduna J.R., Paixão J.M. (1995). "Vehicle Scheduling for Public Mass Transit – An Overview". In Daduna, Branco & Paixão (1995).

Daduna J.R., Wren A editors (1988). "Computer-Aided Transit Scheduling. Proceedings of the 4th Int. Conf. On Computer-Aided Scheduling of Public Transport“. LNEMS, Springer, Berlin.

Desrochers M., Rousseau J.-M. editors (1992). "Computer-Aided Transit Scheduling. Proceedings of the 5th Int. Conf. On Computer-Aided Scheduling of Public Transport“. LNEMS, Springer, Berlin.

Desrosiers J., Dumas Y., Solomon M.M., Soumis F. (1995). "Time Constrained Routing and Scheduling“. In Ball, Magnanti, Monma & Nemhauser (1995), Chap. 2, pp. 35-139.

Hickman M. editor (2008). "Computer-Aided Scheduling of Public Transport Proceedings of the 9th Int. Conf. On Computer-Aided Scheduling of Public Transport“. LNEMS, Springer, Berlin.

Leuthardt H. (1998). *Kostenstrukturen von Stadt-, Überland- und Reisebussen*. Der Nahverkehr 6, pp. 19-23.

Löbel A. (1997). "Optimal Vehicle Scheduling in Public Transit". Pd.D. thesis, Technical University of Berlin, Germany.

Weider S. (2007). "Integration of Vehicle and Duty Scheduling in Public Transport". Pd.D. thesis, Technical University of Berlin, Germany.

Wilson N. editor (1998). "Computer-Aided Transit Scheduling. Proceedings of the 7th Int. Conf. On Computer-Aided Scheduling of Public Transport“. LNEMS, Springer, Berlin.

Voss S., Daduna J.R. editor (2001). "Computer-Aided Scheduling of Public Transport. Proceedings of the 8th Int. Conf. On Computer-Aided Scheduling of Public Transport“. LNEMS, Springer, Berlin.