Konrad-Zuse-Zentrum
für Informationstechnik Berlin

Takustraße 7
D-14195 Berlin-Dahlem
Germany

RALF BORNDÖRFER MARTIN GRÖTSCHEL MARC E. PFETSCH

Public Transport to the fORe!

Supported by DFG Research Center MATHEON "Mathematics for key technologies"

ZIB-Report 05-22 (March 2005)

Public Transport to the fORe![*]

Ralf Borndörfer[**] Martin Grötschel[**] Marc E. Pfetsch[**]

March 2005

Abstract

Can OR methods help the public transport industry to break even?

How would you build a public transport system? For example, have a look at Berlin. The BVG, Berlin's public transport company, maintains a network of 2,423 km, operates 197 lines with 3,286 stops, using 1,554 busses, 1,391 subway cars, and 599 trams from 12 depots, and has 13,409 employees [7]. The BVG currently transports about 800 million passengers per year and covers about 40% of the total non-pedestrian traffic volume of the city [18].

Does Berlin have a "reasonable" public transportation network? Does the BVG run a "good" transportation system? Is it "efficient"?

These are difficult questions. In fact, politicians, transportation managers, customers, taxpayers, etc. frequently employ judgments such as "good" and "efficient", but nobody can give a definition what this exactly means. Since almost every public transportation system in the world is in the red, the cheapest system is no public transportation at all. On the other hand, the most convenient system for the passenger – a stop in front of every house with direct connections to everywhere – is much too expensive. What is the right compromise? Operations Research has no good answer either – so far. But OR can improve aspects of public transportation significantly, as we want to demonstrate in the following.

Where Are We?

That OR can help to organize public transport is not a new idea. Many of the planning steps are instances or variants of classical combinatorial optimization problems, and transport optimization software based on mathematical

[*]Supported by DFG Research Center MATHEON "Mathematics for key technologies"

[**]Konrad-Zuse-Zentrum für Informationstechnik Berlin (ZIB), Takustr. 7, 14195 Berlin, Germany, `[surname]@zib.de`

programming has been developed for several decades. A good overview is provided by the proceedings series of the triannual international conferences on Computer-Aided Transit Scheduling (CASPT), of which we would like to cite the last five proceedings [10, 12, 8, 24, 23]. The 1988 volume contains a survey article [16] by Josef Hoffstadt, at that time CEO of the Hamburger Hochbahn, the public transport company of Hamburg, which lists already 33 industrial scheduling systems, of which four, namely, BUSMAN, HASTUS, HOT, and RUCUS, are singled out specifically for their mathematical excellence. "No mass transit company will be able to reject computer-aided planning in the medium or long term", predicted Hoffstadt.

While Hoffstadt's forecast has become true in large parts of the world in terms of data handling, the same can not be said about the OR part of computer-aided scheduling. Despite the progress in transport optimization, the vast majority of companies and authorities still do their planing essentially by hand and experience. As far as scheduling algorithms are concerned, the progress has been limited in many places to replacing pencil and paper by drag and drop. This does not mean that one cannot obtain good results in this way; in fact, the quality of traditional scheduling is often remarkable – provided that schedulers have experience with their scenarios and sufficient time. But that good tools can only improve scheduling quality is obvious.

Of course, Hoffstadt's prediction refers only to cities or regions where public transport represents an organizational challenge. It is such challenges that our article addresses. The concrete accentuation of the importance of planning problems certainly depends on the local cost structures, such as wages and fares, the availability of streets and tracks, etc. In this article we will give examples from our experience with projects in Germany. The results are probably valid for many other situations and countries.

Three Reasons Why OR is Entering Public Transport

The point of this article is to advocate that it is time for OR to make it into the public transport industry on a broad scale. The reason for our belief is that there is now a combination of favorable factors that have not existed to such an extent before.

- The first factor is the increase in computer storage capacity. Desktop machines with a large main memory are the standard today. They make it possible to keep and maintain all necessary data, to model large scenarios up to the last detail, and to provide real-time access to all necessary information. This eliminates, e.g., the argument that manual post-processing is indispensable.

- Factor two is the improvement in CPU speed combined with tremendous algorithmic progress. In a recent survey [3], Bob Bixby, founder of CPLEX Optimization, Inc., has estimated the performance improvements in computing machine speed from 1987 to 2002 to three orders of magnitude. This has to be multiplied with another three orders of magnitude of algorithmic improvements in linear programming, resulting in a speed-up factor of one million. "A model that might have taken a year to solve 10 years ago can now be solved in 30 seconds", concludes Bixby. These improvements carry over to the applications in public transport as well and make it possible, for the first time, to solve real scenarios with unprecedented precision and solution quality.

- The final factor, perhaps most important, is the increase of serious competition in the area of public transport. This brings cost management to the companies' focus of attention. It accelerates planning cycles and changes scenarios to such an extent that manual planning is just no longer feasible. In Europe, these developments are spearheaded by the deregulation policy of the European Union; here and elsewhere, the globalization leads to the formation of national and international enterprises that exert pressure on the existing local monopolistic transportation structures.

Nearly all public transport companies in the world operate at deficit. For instance, Dieter Ludwig, the chairman of the Union of German Public Transport Companies (VDV), in his opening speech for the 2003 VDV conference in Karlsruhe, specified the cost recovery ratio of all public transport companies in Germany by 70.5% [22]. This average is probably too optimistic.

Is it imaginable that OR can help this industry to break even? We want to give evidence in this article that there exist significant potentials in this direction, which can be harnessed by a combination of modern OR methods and local planning knowledge.

Our Experience

The transportation research group at the Zuse Institute Berlin (ZIB) has carried out investigations in the area of public transport optimization for about twelve years. Our work is on the mathematical foundations of public

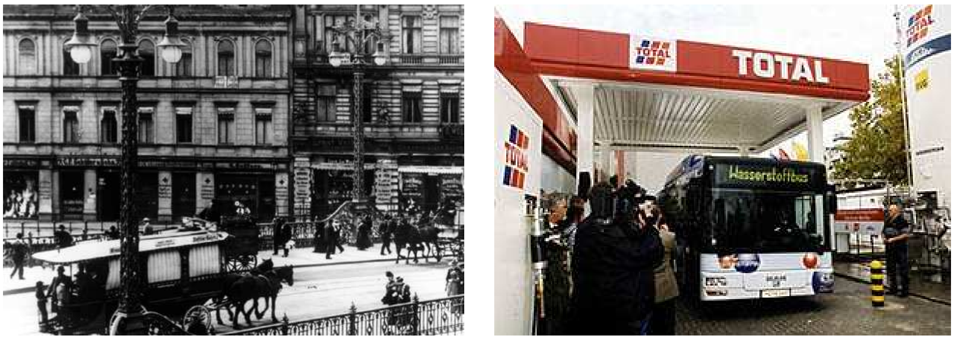

Figure 1: Progress in public transport: Horse bus 1897, H_2 bus 2002. What about OR?

transport optimization and on the transfer of these results into practice. We cooperate with software companies that produce public transport scheduling systems, such as IVU Traffic Technologies AG (IVU) and mentz Datenverarbeitung GmbH (`mdv`), as well as directly with public transport companies such as the BVG. Several of our optimization modules, maintained by the ZIB spin-off company LBW GbR, are currently in use at about 20 public transport companies throughout the world.

We observe that our optimization installations are gaining momentum. The same is reported by other companies that have kept advocating optimization, such as GIRO, Inc. Other software companies have also noticed this trend and are in the process of developing such modules for their systems.

OR Problems in Public Transport

The planning process in public transport is generally subdivided into the three major steps: *planning*, *scheduling*, and *dispatching*. Each of these consists of a series of individual tasks which are dealt with in sequence, see Table 1. *Planning*, sometimes also termed *strategic planning*, makes long and medium term decisions about the level of service that is offered to the public; *scheduling*, or *operational planning*, is concerned with the cost efficient organization of this service for the next timetable period; *dispatching* organizes the daily operation of the system.

OR has so far contributed most to the scheduling aspect. One reason is (probably) that there is only one objective function: cost. Money dominates everything else. The system design (planning) and operations control (dispatching) do not have such a clear-cut goal. Several, often conflicting criteria have to be taken into account.

Phase:	Planning	Scheduling	Dispatching
Horizon:	Long Term	Timetable Period	Day of Operation
Objective:	Service Level	Cost Reduction	Get it done
Steps:	Network Design Line Planning Timetabling Fare Planning	Vehicle Scheduling Duty Scheduling Duty Rostering	Crew Assignment Delay Management Failure Management Depot Management

Table 1: Planning problems in public transport.

Progress in OR Methods

Before we substantiate our claim on the OR perspective in public transport, we outline two important developments in this area.

- The first, and perhaps most important, development of the last decade is the *progress in large-scale methods*. Typically, public transportation problems are mathematically stated in terms of large graphs with tens of thousands of nodes and millions of arcs. In these graphs one searches for lines, vehicle rotations, or driver duties. The number of these structures is, although finite, astronomically large. Constructing them all is beyond question, even more so investigating all combinations in a line plan, timetable, vehicle or duty schedule. The trick in this situation is to consider only a small number of lines, rotations, or duties explicitly, generating additional specimen on the fly as needed and discarding the rest by implicit (i.e., without explicitly looking at them), but rigorous mathematical arguments. Linear programming (LP), whose enormous algorithmic advances were already mentioned above, is the key. LP can guide the generation of additional solution components in such a way that an overall improvement is guaranteed. We mention a few contributions in this general direction: *Lagrangean pricing* and *column generation* are two frameworks for the dynamic generation of parts of graphs and of large integer programs [2, 13, 20]. *Stabilization* is a speed-up technique for these methods [11, 14]. *Volume* and *bundle methods* can be used for the approximate solution of large LPs and integer programs [1, 17].

- This algorithmic progress has been employed for the second important development, namely, *integration*. There are two main approaches: *horizontal* and *vertical integration*.

A company typically subdivides its planning tasks regionally. Uniting such regions and combining the associated scheduling scenarios is called *horizontal integration*. Its ultimate goal is to treat each planning task in a single company-wide scenario. The level reached here is already high today: vehicle scheduling can be done for entire companies, duty scheduling for entire depots, etc.

Additional gains can be derived from the simultaneous solution of consecutive planning steps, such as timetabling and vehicle scheduling, vehicle and duty scheduling, or duty scheduling and rostering, etc.; we call this *vertical integration*. It is well known that vertical integration has considerable saving potentials. Various ideas have been tried to exploit this. *Sensitivity analysis* is such a technique. In the public transportation context, it denotes the slight shifting of timetabled trips in order to save vehicles. It has been applied at the borderline between timetabling and vehicle scheduling with remarkable success since the seventies [9]. Likewise, duties-first, vehicles-second approaches have been used to reduce negative effects of myopic vehicle scheduling decisions that turn out to be disadvantageous in the following duty scheduling step.

These and other approaches of this type address interdependencies between succeeding planning problems. They are, however, not integrated approaches in the sense that they allow the full use of all degrees of freedom that arise by combining two or more consecutive planning problems into one (most likely giant) task.

First steps of this kind have been made very recently; e.g., reports on the successful solution of integrated vehicle and duty scheduling problems have just appeared [6]; we discuss an example below. Further work on vertical integration is certainly one of the most promising future research directions.

Several groups all over the world in countries such as Canada, Germany, Great Britain, Italy, New Zealand, Portugal, Scandinavia, the U.S., and elsewhere are working on theoretical advancements as well as on the practical implementation of such techniques in the public transportation industry. We are now going to discuss three examples from our experience for which the companies have either published results themselves or entitled us to report them. These examples illustrate the developments in the area, the level and the limits that OR in public transit has reached today.

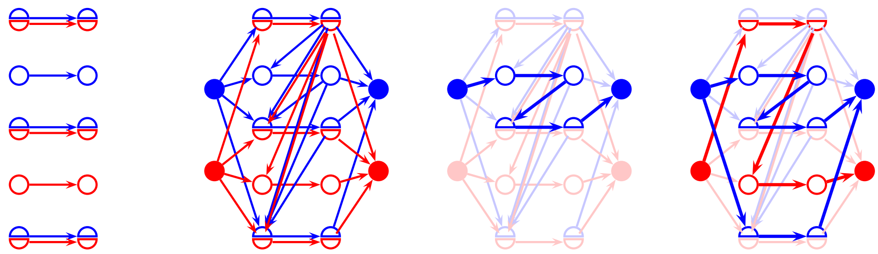

Figure 2: Vehicle scheduling in a small example. *From left to right:* trips; depots, trips and all legal connections; a rotation; a schedule.

Vehicle Scheduling at the Berliner Verkehrsbetriebe (BVG)

The vehicle scheduling problem at the bus department of the BVG is, as far as we know, still the largest problem of this type that has been solved to proven fleet optimality. The optimization was performed from September 1998 to June 1999 in a joint project of the BVG and our group at ZIB. At that time, the BVG executed about 28,000 timetabled bus trips per day. These trips were covered with about 1,500 busses of 14 different types, which operated from ten depots [21]. Not every bus can be assigned to every trip, since there are different types of busses, such as doubledeckers, low-floor busses, etc. Therefore, for every trip there is a set of feasible vehicle types. The problem was to cover the timetabled trips with vehicle rotations in such a way that each trip is assigned a vehicle of feasible type and such that the number of busses is minimized.

A vehicle rotation consists of a sequence of timetabled trips, linked by trip connections, see Figure 2. Such a trip connection can be a real vehicle movement without passengers, a so-called deadhead trip, a turnaround at the end of a line, or a layover at a depot. Clearly, all timetabled trips have to be covered. Savings can only be achieved by a smart choice of trip connections. In other words, the trip connections are the degrees of freedom. For the 28,000 timetabled trips in Berlin, there are about 400 million possible connections. Not all are admissible because, for instance, there may be no common vehicle type for two trips. In the case of the BVG, about 100 million connections had to be considered.

To deal with problems of this type and size, we developped a new solution technique, called *Lagrangean pricing* [20]. Its basic idea is as follows. Iteratively, mathematical simplifications of the problem, so-called relaxations, are considered, essentially by ignoring the vehicle types. This results in large, but simple min-cost flow problems. The trick is that these flow problems can be solved to optimality for all 100 million connections by special purpose methods in a few minutes. The outcome is used as a feedback to construct a new, better relaxation. This is repeated until some stopping criterion is satisfied. Of course, the final relaxation is not as accurate as the exact model, but it is good enough to identify those candidate connections that are important to construct a good vehicle schedule. As a simplification of the actual problem, the final relaxation also produces a lower bound for the cost of an optimal vehicle schedule, which in turn can be used to eliminate a large number of connections that can provably not be contained in an optimal schedule. Using this, one can construct a vehicle schedule on a smaller set. Comparing this solution with the lower bound gives a mathematically rigorous proof for the quality of the solution.

This Lagrangean pricing approach has been implemented and is available as the optimization module VS-OPT. In the case of the BVG problem, VS-OPT computed a fleet-minimal solution for the entire bus department, i.e., there was a mathematical proof that, with the given data, one could not save a single additional vehicle. The BVG has published that for the Spandau depot alone, 38 busses (about 20%) and 377 hours of operation could be saved [21].

Where did these savings come from? One reason is that scenarios of the BVG-class are hardly manageable by traditional scheduling methods, which are based on decomposing the problem into smaller pieces, thereby loosing "degrees of freedom". For instance, the BVG had developed a scheduling process that started from individual lines, coupled them, and finally distributed the resulting rotations to the depots. This process was fairly sophisticated, but it sacrificed optimization potential for overview. There was, in particular, one important "global" factor that was neglected by this bottom-up approach. Namely, the peak in the vehicle demand occured in the morning in the western part of the city, while East Berlin had an afternoon peak. A good scheduling idea was therefore to shift some busses from the eastern depots a little bit to the west in the morning and vice versa in the afternoon. Sounds easy, but it is not so easily done! A company-wide optimization, however, considers such balancing effects just as any other connections and has no special problems with it. In this sense, vehicle scheduling at the BVG is a good example for the potentials of horizontal integration.

Vehicle and Duty Scheduling at the Stadtwerke Bonn (SWB)

The SWB was the first public transport company in Germany that produced vehicle and duty schedules single handed with our optimization modules VS-OPT [20] and DS-OPT [4] in a concerted endeavour. In fact, a vehicle schedule is only as good as the duty schedule that is build on top of it, in particular, if one takes into account that vehicle and crew costs in Germany are roughly in relation 2:1 [19]. Duty scheduling is also combinatorially more difficult than vehicle scheduling, even though both tasks appear quite similar at first sight.

A duty is constructed from indivisible units of driving work, the so-called duty elements, which arise from cutting the vehicle rotations at the possible driver relief points. Analogous to the trips of vehicle scheduling, these duty elements are connected by links that represent transfers by foot and other means of transportation, breaks, or just the continuation of driving the same vehicle. The new aspects in comparison to vehicle scheduling are complicated rules for the legality and the cost of individual duties. The prototype is a break rule stating that a driver has to take a break after a certain maximum time of continuous driving. The point is that such a rule cannot be decided "locally" by looking at any two successive duty elements in a duty. Only the consideration of the entire time horizon gives information whether the rule is satisfied or not. Such and other rules of similar flavour are the reason for the construction of duties being much more difficult than the construction of vehicle rotations.

The SWB started their project in the beginning of 2000. At first their main effort was not in operating the optimization, but in setting up the necessary data. For vehicle scheduling these are: driving times, possible deadhead trips, vehicle types, and depot assignments. For duty scheduling one needs: trip-to-trip traveling times, duty scheduling rules and objectives. After they had collected and input their data, they started the first optimization runs at the end of 2000. Several months and iterations later, they had produced a vehicle and duty schedule that was acceptable for both the management and the workers' council. The result was put into production on June 30, 2001. At that point, the bus department had saved 5 out of 200 busses (2.5%), and 12 out of 280 duties (4.29%). The tram department couldn't save any vehicles, but they reduced the number of duties by 3 from 120 (2.5%).

The SWB, however, had more ambitious plans than that. From the beginning, they planned to use optimization as a decision support tool for their management. One project that they did in 2004 was the relocation of one of their depots. Using VS-OPT and DS-OPT, they could assess the precise cost impact of the changed pull-in and pull-out trips. This helped the management by greatly reducing the uncertainty about the operational impacts of the measure.

Another example was a cooperation with the neighboring Rhein-Sieg-Verkehrsgesellschaft (RSVG), which SWB started in 2003. During the negotiations, they optimized the two lines that at that time were already operated jointly by both companies, and saved one vehicle. Giving this saving to RSVG helped in the progress of the negotiations.

Integrated Vehicle and Crew Scheduling at the Regensburger Verkehrsverbund (RVV)

As large vehicle and duty scheduling optimization problems can nowadays be solved routinely by modern optimization codes, the next challenge is to integrate these two scheduling problems into one big optimization problem. In this way, one can avoid decisions for vehicle scheduling that are bad for the succeeding duty scheduling step. In fact, regional public transport companies often run into even bigger problems, because a traditional plain-vanilla sequential vehicles-first duties-second approach often results in long vehicle rotations without relief opportunities, which do not permit a feasible duty schedule at all. Regional companies have therefore always employed vehicle scheduling strategies that looked ahead to duty scheduling. They did, however, have virtually no tools available to help them. Research has taken up the topic only very recently; for instance, see the survey article [15]. Problems of practical significance have been considered for at best two or three years.

The RVV is a medium-sized public transport company that services a half-urban and half-regional network extending some 30 kilometers around the city of Regensburg, Germany. The split topography of Regensburg, which is dissected by the rivers Danube, Naab, and Regen, results in a star-shaped network with long lines that stretch out into the surrounding regions with very few (in fact, four) relief opportunities to exchange drivers. The RVV had always had difficulties with sequential planning approaches. In 2000, RVV decided to join a research project of ZIB and the German software companies IVU and `mdv` for the development of an integrated vehicle and duty scheduling optimizer IS-OPT [6].

The difficulty of solving integrated vehicle and duty scheduling is, of course, greater than that of the individual problems: The integrated problem contains all the degrees of freedom of vehicle scheduling and all the degrees

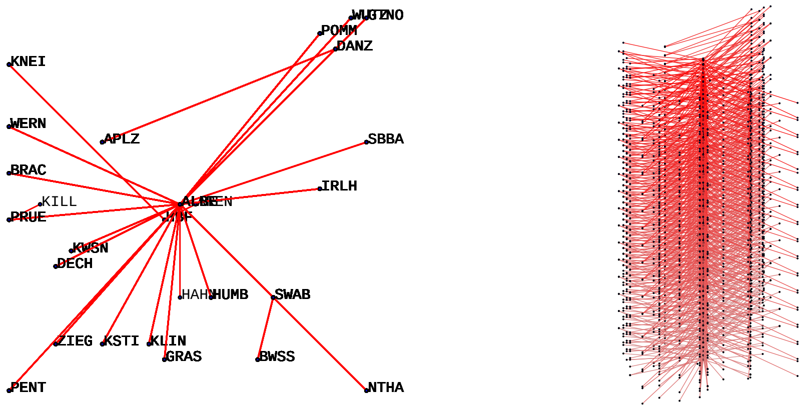

Figure 3: Timetabled trips (direct connections) of RVB Sunday scenario. *Left:* spatial plot; *right:* space-time plot

of freedom of duty scheduling. But the degrees of freedom, or variables, do not add, they multiply. How can one hope to tackle such a problem, when the individual problems are already hard? Perhaps not surprisingly, the trick is to take the problem apart again, but in a special way, which conveys information back and forth between the vehicle and the duty scheduling component of the model. An appropriate technique is the so-called *Lagrangean relaxation* of the coupling constraints between the models. It works as a cost manipulation mechanism as follows. Suppose you have constructed a vehicle and duty schedule independently. In these schedules some vehicle trips will not have a driver, and some drivers will operate phantom vehicles that are not in the vehicle schedule. One now adjusts the costs in such a way as to improve the synchronization; essentially, uncovered elements are made cheaper and overcovered elements more expensive. Mathematically these cost adjustments can be computed in such a way that a maximum of synchronization is attained; this maximum is, however, in general only some lower bound to the vehicle and crew costs. Nevertheless, using this information, one can construct actual schedules that are quite close to the bound.

Table 2 summarizes computational results for a Monday-to-Friday scenario (OPT) and for two variants of a Sunday scenario (OPT 1 and 2). The RVV's Monday-to-Friday scenario is, as far as we know, the largest integrated real-world scenario that has been attacked using this type of mathematical optimization technique. It marks about the limit of what can currently be handled in integrated vehicle and duty scheduling. The problem

IS-OPT	Mo-Fr	Opt	Su	Opt 1	Opt 2
Paid time	1103:48 h	1028:10 h	545:25 h	514:30 h	526:59 h
Breaks	211:53 h	127:06 h	112:36 h	55:11 h	75:06 h
Duties	140	136	82	78	65
Split duties	29	20	0	0	0
Pieces of Work	217	214	87	78	65
Vehicle rotations	91	80	36	33	33
Average	7:56 h	7:34 h	6:39 h	6:35 h	8:06 h
Computing time	—	50:34 h	—	09:30 h	8:35 h

Table 2: Integrated vehicle and crew scheduling at the Regensburger Verkehrsverbund (RVV).

has 1,414 timetabled trips, which have to be covered by about 80 vehicles of three different types, operating from a single depot. The RVV's primary optimization objective is paid time, but they also care for convenience and stability criteria, such as the number of split duties and the number of pieces of work. In roughly two days of computation time on a 3 Ghz PC with 2 GB of main memory, our code IS-OPT produced a solution with savings of 75 hours of paid time (6.83%), 4 duties (2.86%), and 11 vehicles (12.09%), see columns 2 and 3 of Table 2.

Perhaps even more interesting are the Sunday scenarios. They differ in the optimization objective and related scheduling rules as will become clear in a second. Opt 1 basically tried to produce a solution with similar characteristics as the one that the RVV already had: this means an emphasis on short duties such that RVV drivers can soon return to their homes on Sundays. It turned out that this was possible and that in addition 31 hours of paid time (5.68%), 4 duties (4.88%), and 3 vehicles (8.33%) could be saved. Opt 2 was an experiment that allowed for longer duties; it resulted in a dramatic reduction of 17 duties (20.73%), 18 hours of paid time (3.42%), and 3 vehicles (8.33%), at the cost of an increase in the average working time from 6:36 h to 8:06 h. The consequences of these results are clear: The drivers of the RVV now have the choice whether they want to continue to work short but many duties as they do today, or whether they want to work less often on Sundays, but then longer. Now that these schedules are known, management and workers council have entered a discussion to what extent and under which conditions some or all drivers are willing or may even demand to change to longer Sunday duties.

Conclusion

What is the message of these three examples? We have indicated that there are mathematical models that faithfully represent many of the important partial problems which arise in a public transport company. These models can be solved to optimality or near optimality in reasonable running time. Still, for some problems there is room for improvement and better algorithms need to be developed, e.g., for duty scheduling. The combination of the individual optimal solutions does not necessarily provide a system optimum. The mathematical challenge here is to design methods capable of handling global models that are formed by integrating the many subproblems of operational planning. Even if we could solve such integrated models to optimality, we have to keep in mind that the solutions heavily depend on the quality of the input data and the choice of the model parameters. By varying the data and parameters these models can provide excellent support in the decision process in public transport companies. Moreover, often big improvements can be obtained by changing old rules that have lost their meaning, either because nowadays more complex problems can be handled or the planning process has changed. In this way, mathematics can help public transport planners.

Can we now say what a "good" or "efficient" transport system is? Honestly, no! Although the present OR methods lead to significant progress in this respect, nobody really understands how to quantify "good" or "efficient" properly. This is in particular true if one wants to take into account the conflicting goals which customers, operators, and politicians may have and if, in addition, ecological and environmental aspects are also considered. The challenge here is not yet mathematical, it is to find adequate OR models.

A Guess at the Future

If we should make a guess at the future, we certainly see more progress in the solution of large-scale models and further success in the integration process. But we believe that some new aspects will come into play.

One is a greater impact of optimization in the area of (strategic) planning, i.e., the planning of the network, the lines, the timetable, and the fares. Today, this is done incrementally, in a highly political process, often using simulation tools. The political influence will, of course, remain, but we believe that quantitative optimization methods can make a contribution to arrive via a more transparent way at – hopefully – better decisions. We have, e.g., recently started a project on line optimization that aims at the construction of a line system that balances cost efficiency and a small total passenger traveling time [5].

Similarly, we see a great interest in online optimization methods to deal with failure events such as breakdowns, driver no-shows, lateness, etc. during the day of operation. Most public transport companies have now installed sophisticated operations control systems. These systems increase the awareness of failures and their costs, and they record megabytes of data that sit around waiting to be put to good use. In fact, what is the value of the best schedule if the actual operation is dramatically different and much more expensive? Intelligent online rescheduling methods as well as failure anticipating offline strategies will become more important in the future to address this basic problem.

Finally, we are surely missing something if our scheduling methodology continues to ignore the increasing importance of competition in the area of public transport. In fact, the entire planning process and all the models as established today have their roots in a time when monopolistic providers could organize the public transport in “their area” as they thought best. Basic tasks of vehicle and duty scheduling will, of course, always come up in a similar way, but it is almost sure that additional market elements will gain significant importance. Maybe it will not come so far as that we see competing public transport companies in the same service area everywhere, with auctions deciding on the kind, cost, and the particular provider of a service. However, something in this direction is in the air, and will lead to new game theoretic approaches.

Will public transportation break even? We think that the best use of the full power of OR techniques can bring us close. There is still a long way to go—let’s start!

References

[1] F. Barahona and R. Anbil, *The volume algorithm: producing primal solutions with a subgradient algorithm*, Math. Program., 87 (2000), pp. 385–399.

[2] C. Barnhart, E. L. Johnson, G. L. Nemhauser, M. W. P. Savelsbergh, and P. H. Vance, *Branch-and-price: Column generation for solving huge integer programs*, in Mathematical Programming: State of the Art, J. Birge and K. Murty, eds., University of Michigan, 1994, pp. 186–207.

[3] R. E. Bixby, *Solving real-world linear programs: A decade and more of progress*, Oper. Res., 50 (2002), pp. 3–15.

[4] R. Borndörfer, M. Grötschel, and A. Löbel, *Duty scheduling in public transit*, in MATHEMATICS – Key Technology for the Future, W. Jäger and H.-J. Krebs, eds., Springer Verlag, Berlin, 2003, pp. 653–674. ZIB Report 01-02.

[5] R. Borndörfer, M. Grötschel, and M. E. Pfetsch, *Models for line planning in public transport*, Tech. Rep. ZIB Report 04-10, Zuse-Institut Berlin, Takustr. 7, 14195 Berlin, 2004.

[6] R. Borndörfer, A. Löbel, and S. Weider, *A bundle method for multi-depot integrated vehicle and crew scheduling in public transport*, Tech. Rep. ZIB Report 04-14, Zuse-Institut Berlin, Takustr. 7, 14195 Berlin, 2004.

[7] `http://www.bvg.de/ueber/daten.html`, Dec. 2002.

[8] J. R. Daduna, I. Branco, and J. M. P. Paixão, eds., *Computer-Aided Transit Scheduling*, vol. 430 of Lecture Notes in Economics and Mathematical Systems, Springer-Verlag, Berlin Heidelberg, 1995.

[9] J. R. Daduna and M. Völker, *Fahrzeugumlaufbildung im ÖPNV mit unscharfen Abfahrtszeiten*, Der Nahverkehr, 11 (1997), pp. 39–43.

[10] J. R. Daduna and A. Wren, eds., *Computer-Aided Transit Scheduling*, vol. 308 of Lecture Notes in Economics and Mathematical Systems, Springer-Verlag, Berlin Heidelberg, 1988.

[11] G. Desaulniers, J. Desrosiers, and M. Solomon, *Accelerating strategies in column generation methods for vehicle routing and crew scheduling problems*, in Essays and Surveys in Metaheuristics, C. C. Ribeiro and P. Hansen, eds., Kluwer, Norwell, 2002, pp. 309–324.

[12] M. Desrochers and J.-M. Rousseau, eds., *Computer-Aided Transit Scheduling*, vol. 386 of Lecture Notes in Economics and Mathematical Systems, Springer-Verlag, Berlin Heidelberg, 1992.

[13] J. Desrosiers, Y. Dumas, M. M. Solomon, and F. Soumis, *Time constrained routing and scheduling*, in Network Routing, M. O. Ball, T. L. Magnanti, C. L. Monma, and G. L. Nemhauser, eds., vol. 8 of Handbooks in Operations Research and Management Science, Elsevier Science B.V., Amsterdam, 1995, ch. 2, pp. 35–139.

[14] O. du Merle, D. Villeneuve, J. Desrosiers, and P. Hansen, *Stabilized column generation*, Discrete Math., 194 (1999), pp. 229–237.

[15] R. Freling, A. P. Wagelmans, and J. M. P. Paixão, *An overview of models and techniques for integrating vehicle and crew scheduling*, in Wilson [24], pp. 441–460.

[16] J. Hoffstadt, *Computer-aided scheduling in urban mass transit companies: Past, present and future*, in Daduna and Wren [10], pp. 1–7.

[17] K. Kiwiel, *Proximal bundle methods*, Math. Program., 46 (1990), pp. 105–122.

[18] Land Berlin, Senatsverwaltung für Stadtentwicklung, Abt. VII Verkehr, *Nahverkehrsplan des Landes Berlin, Fortschreibung 2000/2001 und 2004*. http://www.stadtentwicklung.berlin.de, 2000.

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

[20] A. Löbel, *Vehicle scheduling in public transit and Lagrangean pricing*, Manage. Sci., 44 (1998), pp. 1637–1649. Available as ZIB Preprint SC 96-26 at http://www.zib.de/ZIBbib/Publications.

[21] Partner für Berlin, Gesellschaft für Hauptstadt-Marketing GmbH, *Berlinbrief 9/10 2003*. http://www.berlin-partner.de/index.php?page=99modaction=showItem&id=154%, 2003. ISSN 1611-3284.

[22] Verband Deutscher Verkehrsunternehmen, *Bericht des Vorsitzenden Dieter Ludwig*. Eröffnungsansprache zur Heureka 2003, Karlsruhe, 2003.

[23] S. Voss and J. R. Daduna, eds., *Computer-Aided Transit Scheduling of Public Transport*, vol. 505 of Lecture Notes in Economics and Mathematical Systems, Springer-Verlag, Berlin Heidelberg, 2001.

[24] N. H. M. Wilson, ed., *Computer-Aided Transit Scheduling*, vol. 471 of Lecture Notes in Economics and Mathematical Systems, Springer-Verlag, Berlin Heidelberg, 1999.