Integrated timetabling and passenger satisfaction with disturbances

Naut Bulten [*]
Shadi Sharif Azadeh
Rommert Dekker

February 2019

[*]Bulten@ese.eur.nl

1   Introduction

The number of people making use of public transportation, and railway in particular, has seen a steady growth over the course of the past few decades. Population growth, environmental awareness, economic prosperity, and many other factors are contributing to the increase in demand for public transportation, which causes the need for robust planning methods for the Train Operator Company (TOC).

In our research we focus on a specific part of the planning process of a TOC, namely the Train Timetabling Problem (TTP). This stage of the planning process, which is generally executed at least once a year, comes down to determining the arrival and departure times of all vehicles at each of their stopping stations. The infrastructure and fleet size are exogenous to this problem, as well as the line planning. This means that it is given how many times each vehicle has to visit what station (in what order), and it has to be determined at what time they will do so.

The TTP can be formulated and solved in various ways, depending on which aspects the modeler wants to incorporate. In public transportation, punctuality has a prominent influence on the customers' satisfaction, so it is important that the timetable can be executed in the way that it was planned as much as possible. Besides reliability, other factors such as the total travel time and frequency of the departures are important for the utility of passengers. In this paper, the goal is to create a timetable which minimizes the *perceived* travel times of the passengers. Perceived travel time, otherwise known as generalized travel time, is a metric to measure the satisfaction of the passengers.

The framework of our research is to minimize the perceived travel times in the presence of *disturbances*. In the TTP the infrastructure is given and we know the minimum travel time between pairs of consecutive stations. We define a disturbance to be an unavoidable event causing a vehicle to require more than minimum travel time between a pair of stations. We assume the disturbances to be exogenous, and attempt to reduce the *delays* they cause using buffers. The difference between a disturbance and a delay is that the latter is measured as a deviation between the scheduled and realized departure times. An example to clarify: suppose the minimum travel time between stations A and B is 5 minutes. If at some moment a defect requires a given vehicle to take at least 8 minutes to travel from A to B, there is a disturbance of 3 minutes. However, despite the minimum travel time being 5 minutes, it might be that in the posted timetable there were 7 minutes scheduled for this vehicle to go from A to B. In this example there is only 1 minute of delay.

This example contains the crucial tool we will use to minimize perceived travel times: *buffers*. As we are deciding on the arrival and departure times, we have the option of scheduling more than minimum travel time between consecutive stations. In the given example, this means that two minutes of buffer were scheduled, which absorbed a large part of the disturbance.

The remainder of this paper is organized as follows: in Section 2 we discuss the research that has been done on related topics. After that, in Section 3, we give a more detailed explanation of the problem. Then, in Section 4, we explain our approach to solve this problem, and we conclude with the results in Section 5.

2 Related literature

The framework in the previously presented section is loosely based on Kroon et al. (2005). In this paper, the authors insert buffers (named supplements) in the timetable to reduce delays and solve the mathematical model to optimality using CPLEX. This is first done for a single line, single vehicle, visiting each station once, using Sample Average Approximation (SAA), Kleywegt et al. (2002). The model is then expanded to include multiple vehicles making use of the well known PESP formulation developed by Serafini and Ukovich (1989). The main differences with our project is that we are operating in non-cyclic framework, which allows for more freedom when distributing the buffers, as well as that we work with a given horizon which then implicitly gives an upper limit to the buffers that can be allocated whereas Kroon et al. (2005) work with a predetermined amount of buffer which then implicitly determines the horizon.

The integration of the train timetabling problem and passenger satisfaction is treated in Robenek et al. (2016). The authors create a mathematical model that incorporates passenger route choice as an endogenous factor to the timetabling problem. The objective is to maximize profits (ticket revenue minus driver and rolling stock costs), and a Pareto frontier with constraints on the minimum passenger satisfaction is created in a case study of the Swiss Federal Railways (SBB). A comparison between cyclic and acyclic timetables is made, but the robustness against disturbances is not considered.

An effort to minimize delays comes from Jovanović et al. (2017), who consider the slack in a cyclic timetable to be the difference between the period length and the minimum time required to execute all actions. In essence, they consider the travel times fixed and a timetable to be given, and shift the departure times by solving a knapsack formulation to create more than the minimum headway time as a buffer between consecutive trains, thus aiming to minimize delay propagation. Besides not requiring an existing cyclic timetable, our framework also differs by having the ability to reduce primary delays through the option of scheduling more than minimum travel time whereas Jovanović et al. (2017) only consider allowing more than minimum headway time between consecutive departures.

A similar idea is discussed in Cacchiani et al. (2012), where the timetable is also assumed to be given, but besides shifting departures of a vehicle it is also possible to 'stretch' the departures of a vehicle, meaning they insert both buffers between consecutive trains as well as insert time supplements. The formulation is based on Fischetti et al. (2009) and solved using a Lagrangian heuristic. Our formulation has the same tools at its disposal, but the main difference is that we do not work from a given timetable, whereas Cacchiani et al. (2012) do and are limited in how far they can deviate from it.

Lastly, Liebchen et al. (2010) have developed a method to integrate both the timetabling as well as the real-time delay management. Whereas most robust timetables do not incorporate any delay management, it is in fact a more realistic view to incorporate the possibility for the TOC to minimize delay propagation as there are researchers actively working on improving the best methods to do so.

Considering the elements in our research (delays, passenger utility, cyclic versus acyclic timetables), we can conclude that research has been done with respect to all of them, but not (to the best of our knowledge) with all of them integrated as

endogenous factors. This is precisely the gap we intend to fill by explicitly incorporating passenger satisfaction in the train timetabling problem with buffers for reducing both primary and secondary delays.

3 Problem description

The essence of our problem is to determine the departure and arrival times of all vehicles at each of their stopping stations. We incorporate the time spent at a station (dwell time) as required travel time between stations, and thus we only need to determine the departure times.

The input to our problem is a fleet of homogeneous vehicles, each of which is assigned to exactly one of the lines that are serviced, with a given number of round-trips to make in the horizon. For safety reasons, it is necessary that the vehicles do not drive too close to each other, known as the minimum headway constraint. Thus, we have to ensure that all departures are planned within the horizon, and that for any pair of consecutive departures at a station, they are at least a predefined amount of time apart (the minimum headway).

Furthermore, between each pair of consecutive stations, a minimum required travel time is given, and the timetable is constrained to allow for *at least* the minimum required travel time for every vehicle traversing between the stations. We allow ourselves to schedule more than the minimum required travel times to absorb potential disturbances.

Because of the disturbances, the vehicles will often require more than the minimum travel time. As the disturbances are stochastic, we work with sample average approximation to evaluate their effect on the timetable. We create one planned timetable and multiple realized timetables, where the latter are effected by means of realizations of the disturbances. The realized departure time can not be earlier than the planned departure time, and it must also hold that it is not earlier than the realized departure time at the previous station plus the minimum required travel time and the disturbance (if there is one).

Besides disturbances, it is also possible that the vehicles are slowed down due to secondary delays. Just like in the planned timetable, the vehicles have to adhere to the minimum headway constraint in the realized timetable as well, meaning that a vehicle which is behind on schedule can cause delays for subsequent vehicles as well. Here, the point of the buffers becomes clear: we can allow vehicles to have more than the minimum required travel time to traverse between the stations (to absorb disturbances), but we can also use the buffers to space out the departure times of consecutive vehicles (to reduce secondary delays). As the horizon and minimum travel times are given, we are limited in the amount of buffer we can use. The framework above specifies the requirements on the planned timetable and its evaluations, the realized timetables. Using these settings, we can create a schedule that minimizes delays, which we measure as the difference between planned and realized departure times. Preliminary results of this model are shown in the results section. This is an ongoing research and we are currently working on more complicated models as extensions.

The problem above describes a vehicle oriented model, where only the minimization of delay matters. Because we have not incorporated passenger flows yet, all delays have been penalized with equal weight. Furthermore, we use buffers to

schedule more than minimum required travel time at the cost of increasing the planned travel time, but we need to account for the negative effect of an increase in planned travel time on the utility of passengers. Finally, by taking the passenger satisfaction into account, we can incorporate the benefit of offering a high frequency during rush hour.

To integrate these elements, we include *passenger segments*. A passenger segment is a group of passengers sharing the same characteristics; origin, destination, and desired time of arrival. Rather than minimizing the delays, we focus on the *perceived travel times*, weighted by the size of the segments. The perceived travel times consist of the total travel time with the addition of linearly penalized delays and deviations from the desired as well as planned arrival time. The perceived travel times, however, depend on the chosen departure and the corresponding realized delays. To model this: we make the following assumptions:

(1) the passengers do not have knowledge of (or means to estimate) the delays,

(2) the passengers choose their departure based on the planned timetable (ergo assuming no delay will occur), and do so by minimizing their perceived travel time. In essence, we assume the passengers to choose the best planned departure, but the quality of the timetable is evaluated on the experience of the passengers with the realized timetables.

4   Methodology

The vehicle oriented framework has been formulated such that it is completely linear and all the variables are continuous. This means that it is efficiently solvable using existing linear programming methods. The usage of realizations of disturbances does come with the danger of the solution being made to fit specific historic values as well as the risk of using an unusual sample, which needs to be avoided by using a sufficient number of realizations. Convergence can be investigated through iteratively increasing the number of realizations and observing the obtained objective values becoming worse with smaller variances until the changes between consecutive iterations become negligible.

When we extend the framework to include the passengers, we need to model their choice of departure. The problem that arises, is that from a theoretical standpoint each passenger segment can choose any departure. We consider all departure choices to be feasible, and the optimal choice is not known beforehand as it is dependent on the timetable. This means that it is intractable to explicitly consider all possible departure choices for each segment as it would require too much runtime and memory to solve using branch and bound.

In order to deal with these issues, we are developing a column generation algorithm for large scale problems where we iteratively generate departure options for the passenger segments. In each iteration, we determine a planned timetable and corresponding realized timetables based on a subset of available choices for each segment (restricted master problem). Given this timetable we find, for each segment, the best departure that was not an option in the restricted master problem (pricing problem). We then add this departure to the subset of choices of the segment, and continue with the next iteration.

As the number of variables and constraints can be too big to solve linear relaxations of the master problem, we also need to develop a method for calculating

lower bounds. To do this, we split the master problem by the segments, allowing us to create a different timetable for different sets of passenger segments. The smaller such a set is, the easier the problem is to solve. Adding up the weighted perceived travel times of the sets of passenger segments gives a valid lower bound.

5 Results

We present the results of experiments using synthetic data of the vehicle oriented settings, work on the passenger oriented approach is underway. We are interested in the number of realizations required to get accurate results, as well as the effect of (not) imposing cyclicity. In our case, a cyclic timetable implies that there is a basic hour pattern. In other words, the schedule for any hour is the same as the schedule for the previous hour with all departures shifted by 60 minutes. In order to be able to compare acyclic and cyclic timetables, it is necessary that the number of vehicles multiplied by the number of trips per vehicle (= the total number of trips) is a multiple of the number of periods, otherwise only partially cyclic timetables can be constructed. Furthermore, as we are mainly concerned with rush hour planning, we need to have a horizon large enough to capture an entire peak. Therefore we used the following base case settings for our experiments:

Fleet size	Trips per vehicle	Stations	Horizon (min.)	Period (min.)
5	4	10	240	60

Figure 1 displays the change in mean delay (the objective) when the number of realizations is increased. Here we have used the uniform distribution with parameters 0 and 2 for all of the disturbances. The pattern shows two elements we want to avoid: using too few realizations leads to overfitting (low mean delay for low number of realizations), and using too few realizations is unreliable (high variance in mean delay for low number of realizations). The majority of both effects seem to disappear when using at least 40 realizations, but adding more realizations still improves stability at this point. Therefore, unless otherwise mentioned, we use 100 realizations in our experiments.
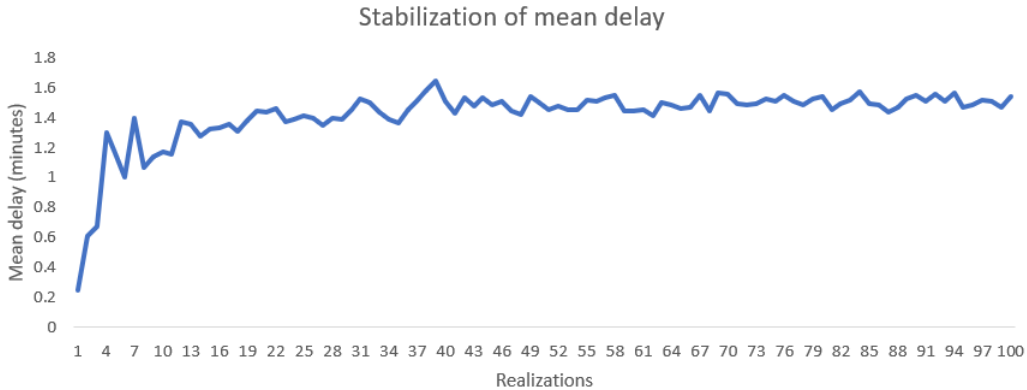

Figure 1: The effect on mean delay when the number of realizations is changed

To compare the objective values of the cyclic (C) and acyclic (A) timetables, we have solved the problem for both settings on four types of disturbance distributions

6

of which the results are displayed in Table 1. For each setting, we have created 5 instances, each having 100 realized sets of disturbances. In the base case we have generated all disturbances using the uniform distribution with parameters 0 and 2. Note that the reported standard deviation is of the delays within an instance.

Table 1: Results on vehicle oriented framework

	Mean Delay	St. dev. Delay	% Decr. Mean Delay
Base case C	3.6	2.9	-
Base case A	1.5	2.3	58.3
Exp(1) C	6.8	5.8	-
Exp(1) A	3.5	4.7	48.5
Not i.i.d. over stations C	5.2	4.3	-
Not i.i.d. over stations A	2.4	3.9	53.8
Not i.i.d. over trips C	13	12.2	-
Not i.i.d. over trips A	2.4	3.3	81.5

For the second case we used the exponential distribution with parameter 1 which gives the same expected disturbance, but a higher variance (1 versus $\frac{1}{3}$). This leads to an increase in mean delay for both the cyclic and acyclic timetable, but decreases the gap between them.

For the third experiment, we use different disturbance distributions for different parts of the network. For 15 out of 20 arcs we generate them uniformly with parameters 0 and 1, and for the remaining 5 arcs we use the uniform distribution with parameters 0 and 5. This implies that the expected disturbance over all the arcs is still 1. Both the acyclic and cyclic timetables are worse compared to the base case, but the gap between them (53.8) does not change much. The reason for this, is because the disturbance distribution is still i.i.d. over the trips, meaning both the cyclic and acyclic timetable can adjust to the change in distribution.

For the final experiment we generate the disturbances to be uniform with parameters 0 and 1 for all arcs of 3 out of 4 trips of each vehicle, and use parameters 0 and 5 for 1 out of 4 trips of each vehicle. As in the previous scenario the expected disturbance taken over all arcs is still 1, but now the disturbance distribution changes over time, not space. Here, the disadvantage of cyclic timetables becomes more clear. The mean delay is more than three times as big as in the base case, and even more than four times as big as the acyclic timetable for the same scenario.

Our current research is focused on the integration of passenger satisfaction. As mentioned before, this gives a large number of variables and constraints. We are currently implementing a column generation algorithm and improving the lower bounds obtained through splitting the master problem over the passenger segments. This method will be applied to the network of the RET, which operates the metro in Rotterdam, The Netherlands.

References

Cacchiani, V., Caprara, A., and Fischetti, M. (2012). A lagrangian heuristic for robustness, with an application to train timetabling. *Transportation Science*, 41(1):124–133.

Fischetti, M., Salvagnin, D., and Zanette, A. (2009). Fast approaches to improve the robustness of a railway timetable. transportation. *Transportation Science*, 43(3):321–335.

Jovanović, P., Kecman, P., Bojović, N., and Mandić, D. (2017). Optimal allocation of buffer times to increase train schedule robustness. *European Journal of Operational Research*, 256(1):44–54.

Kleywegt, A., Shapiro, A., and de Mello, T. H. (2002). The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization*, 12(2):479–502.

Kroon, L. G., Dekker, R., and Vromans, M. J. C. M. (2005). Cyclic railway timetabling: a stochastic optimization approach. *ERIM report series research in management*, ERS-2005-051-LIS.

Liebchen, C., Schachtebeck, M., Schöbel, A., Stiller, S., and Prigge, A. (2010). Computing delay resistant railway timetables. *Computers and Operations Research*, 37(5):857–868.

Robenek, T., Maknoon, Y., Azadeh, S. S., Chen, J., and Bierlaire, M. (2016). Passenger centric train timetabling problem. *Transportation Research part B*, 89:107–126.

Serafini, P. and Ukovich, W. (1989). A mathematical model for periodic event scheduling problems. *SIAM Journal on Discrete Mathematics*, 2(4):550–581.