Towards a Better Train Timetable for Denmark
Reducing Total Expected Passenger Time

Peter Sels · Katrine Meisch · Tove
Møller · Jens Parbo · Thijs Dewilde ·
Dirk Cattrysse · Pieter Vansteenwegen

Abstract With our Periodic Event Scheduling Problem (PESP) based time-tabling method we are able to produce a passenger robust timetable for all 88 hourly passenger trains running on tracks managed by the Danish Infrastructure Manager Banedanmark. The objective function of our model is the total expected passenger journey time in practice and is minimised. The result of this is that the produced timetable reduces the expected journey time of all corresponding train passengers together by 2.9% compared to the original timetable defined by Banedanmark. Our simulations show that the average probability of missing a transfer is also reduced from 11.34% to 2.45%. The computation of this timetable takes only 65 minutes. The major innovations of our approach are the addition of a complete objective function to the PESP model and the addition of a particular cycle constraint set that reduces computation times. In this paper, we demonstrate that these combined innovations result in a method that quickly generates cyclic timetables for a train network spanning an entire country and that these timetables also reduce the expected passenger travel time in practice.

P. Sels
Logically Yours BVBA
Plankenbergstraat 112 bus L7, 2100 Antwerp, Belgium
Tel.: +32-486-95 67 97, E-mail: sels.peter@gmail..com
Infrabel, Traffic Management & Services
Fonsnylaan 13, 1060 Brussels, Belgium

K. Meisch, T. Møller
Banedanmark, Amerika Plads 15, 2100 København, Denmark

J. Parbo
DTU Transport, Bygningstorvet 116B, 2800 Kgs. Lyngby, Denmark

P. Sels, T. Dewilde, D. Cattrysse, P. Vansteenwegen
KU Leuven, Leuven Mobility Research Centre, CIB
Celestijnenlaan 300, 3001 Leuven, Belgium

1 Introduction

This paper's topic is the automatic construction of a cyclic, macroscopic railway timetable. The word cyclic means that there is a timetable period, here 1 hour, by which every train repeats itself. The word macroscopic means that a standard value for the minimum headway times of 3 minutes is assumed and inside stations, the microscopic headway constraints that arise from the block sections staircase model are not enforced. We also assume that line planning is fixed including the halting pattern for each line. This means that for each train, for each station, only the arrival and departure time are to be determined. In other words, only ride and dwell supplements are to be chosen. Of course, many solutions exists, but these supplements have to be chosen so that the resulting timetable possesses some desirable proporties. We previously constructed a Periodic Event Scheduling Problem (PESP) based model which has as objective function: the total expected passenger journey time in practice over all passengers (Sels et al, 2015b). In Dewilde et al (2013), the authors conclude that, unlike to what is the case for some alternative definitions of robustness, this objective function is a *practical* method to obtain robustness and that the obtained robustness is ideal for passengers. Our objective function integrates and makes a trade-off between efficiency and robustness. It penalises supplements that are so big that they would lower efficiency too much but also penalise supplements that are so small that robustness would be compromised.

In Sels et al (2015b), this MILP model is generated for the set of all 196 hourly trains in Belgium. The main results were that a timetable, automatically generated in about 2 hours, saves about 3.8% of total expected passenger journey time. This timetable also significantly reduced the percentage of missed transfers from 13.9% to 2.6%. To study how generally applicable this model is to practice, we now also test it on the set of all 88 hourly trains using Banedanmark's infrastructure.

2 Timetabling Methodology and Assumptions

Our timetabling approach consists of the basic constraints of the popular *PESP* model (Serafini and Ukovich, 1989; Schrijver and Steenbeek, 1993; Nachtigall, 1996; Goverde, 1998a,b; Peeters, 2003; Kroon et al, 2007; Liebchen, 2007; Kroon et al, 2009; Caprara et al, 2011; Sparing et al, 2013) using a standard *event activity network*. We impose its classic *constraints* enforcing minimal ride times and minimal dwell times. As described in detail in Sels et al (2011), we automatically construct all *potential transfers*. By this, we mean that if two trains stop in the same station, a transfer edge will be added between the arrival time of the feeder train and the departure time of the target train. Currently, a minimum of 3 minutes is assumed for each transfer. Headway edges and the respective minimum headway time constraints are also automatically constructed between entry times of each pair of trains that enter the same

infrastructure resource and similarly also between all pairs of exit times. For single track sections, between each leaving and each entering train, a similar headway time constraint is imposed. The headway minimum time assumed on this macroscopic level is 3 minutes. This summarises all hard constraints in our model. For more details, we refer to Sels et al (2015b), where all these mandatory constraints and some supplementary ones that are merely intended to speed up computation are discussed.

We will also only give a qualitative description of our objective function here, as the main focus of this paper is the application of our timetabling model on the Danish train network. As derived formally in detail in Sels et al (2013b) and Sels et al (2013a), our *objective function* consists of the sum of the expected passenger time for each edge (action) in the event activity graph $G(V, E)$ that corresponds to a passenger activity. So, for each ride, dwell and transfer edge we model an expected passenger time. We express this expected passenger time of an edge as a function of its minimum time and its added *supplement time*. The *shape* of this function mainly depends on the expected primary delay distribution and consequently, so does the value of the supplement that should be ideally added. The *scale* of this function depends on the number of passengers involved. This indicates the relative importance of the expected passenger time of one edge compared to that of another and these are balanced by the objective function.

For the primary delays, as do Schwanhäußer (1974); Meng (1991); Ferreira and Higgins (1996); Goverde (1998a); Vansteenwegen and Van Oudheusden (2006); Kroon et al (2006) and Yuan (2006), we assume negative exponential distributions. These distributions have an average (=expected value) that can be set to a certain fixed percentage 'a' of the minimum time for that action. This average can in theory be determined by inspecting logs of trains as they are running in the current timetable. This has been described by Goverde and Hansen (2000) and Daamen et al (2009) for the Dutch and by Labermeier (2013) for the Swiss infrastructure. So, for example, if the minimum time of a ride action is 5 minutes from one stop to the next, if one sets 'a' to 10%, the average primary delay on that ride action is assumed to be 0.5 minutes. By this one parameter, the negative exponential distribution $p(d)$ of the primary delay d is unambiguously defined, as $p(d) = exp(-d/a)/a$. For now, we assume the same value of 'a' for all ride, dwell and transfer edges, for all trains and for all tracks. The value of 'a' is typically chosen in the range of 1% to 5% (Goverde, 1998a).

Depending on the action type that passengers participate in, the expected passenger time is another type of function of the supplements added to these actions. We now discuss these types of passengers and associate cost functions.

For *through* passengers, experiencing a ride and subsequent dwell action, the expected time, as a function of the added ride and dwell supplements s, as can be seen in the example in figure 1, is almost the function $f(s) = P \cdot s$, with P equal to the number of participating passengers. This is logical, since for whatever supplement is added to a ride or dwell action, the through passengers just have to sit it through. So high values of s are not beneficial to

Fig. 1 Through and arriving passenger expected time as a function of the chosen supplement. All time is given in in 6 second multiples.

these passengers. At low values of s, the slope of $f(s)$ is a little flatter because small delays occur more often than large delays and so, waiting for the end of s takes a smaller fraction of time s on average than for larger supplements. The larger the supplement, the smaller the fraction that common delay sizes form compared to it. So for larger supplements this secondary 'curving effect' diminishes. The situation is entirely similar for arriving passengers, experiencing a ride plus sink action, and so the cost function for arriving passengers is also similar to the one shown in figure 1. Note that the green vertical line shows that an 8 minute supplement was chosen by the solver. A supplement equal to 0 minutes would be locally optimal, but other hard constraints like headway constraints may forbid this here.

Note that all cost functions in figures 1, 2, 3 and 4 show an actual expected time cost function in green that is used in evaluation and a piecewise linear approximation of it in red which is used in linear optimisation. The green vertical line indicates an example of an actual chosen supplement. Its associated expected passenger time cost can then also easily be read from the graph. In each case, we see that the linearisation error is relatively small.

To *departing* passengers, experiencing a source plus ride action, it is beneficial when the train they get onto departs as scheduled. This is ensured by providing enough time buffers against primary delay on this train on the sections this train traverses *before* these departing passengers embark on it. The curve in figure 2 shows indeed that the selection of a larger buffer on the previous sections for this train statistically leads to lower expected delay for departing passengers than a lower buffer. However, it also demonstrates that a supplement larger than 10 minutes does not significantly increase the buffer-

Fig. 2 Departing passenger expected time as a function of the chosen supplement. All time is given in in 6 second multiples

ing effect compared to a 10 minute supplement. The green vertical line shows that the MIP solver decided to set the supplement to 8 minutes. This is not the local minimum, 60 minutes, but due to competition with other terms in the objective function this could be a reasonable choice. The value 8 minutes is the absis of the crossing of the two red segments which meet on the green curve so the linearisation error is 0 here.

For passengers who are *changing* between trains, experiencing a ride plus transfer action, we model an expected transfer time that depends on the chosen supplement for this transfer, on top of the minimum of 3 minutes. If the supplement is low, the probability that the transfer is missed is high. If the transfer is missed, we conservatively assume a penalty waiting time of the timetable period, here 1 hour. If the supplement is high, the probability of missing the transfer is low, but the transfer passenger will always have to wait until the supplement time has elapsed. The above means that the expected passenger time for a transfer is a U-shaped function of the supplement. An example of a transfer cost curve is given in figure 3. So there is a trade-off and a locally optimal value for the transfer supplement somewhere between 0 and 60 minutes. This supplement range is very broad and naturally very large supplements will rarely be added. Exceptionally, like when a transfer is only taken by very few people, and a small supplement on this transfer would mean a large supplement on an action with more people, a very large supplement on this less important transfer can occur though. The allowed range for supplements is defined as 0 to 60 minutes to avoid infeasibility problems. Note that a transferring passenger can be seen as the combination of both an arriving and a departing passenger and this is reflected in the cost function in figure 3 being

Fig. 3 Transfer passenger expected time as a function of the chosen supplement. All time is given in in 6 second multiples

the addition of the cost functions of figures 1 and 2. The vertical green line in figure 3 shows that the MIP solver was able to select a supplement equal to 4.5 minutes which minimises the local linearised expected transfer time. This also coincides with the minimum of the green curve.

As for secondary delays, or *knock-on delays*, our model already contains the graph edges associated to these. Indeed, they are the same edges as the headway edges, temporally separating pairs of trains that use the same infrastructure resource. So for each headway edge, we also add a term in the objective function that represents the knock-on time or secondary delay that passengers on the second train may experience in case the first train is delayed. In our model, as derived in Sels et al (2013a), this time depends on the delay distributions of both trains and on the number of passengers on the second train. Obviously, the total knock-on time is proportional to the number of passengers on the second train. Also, the expected knock-on passenger time forms a decreasing function of the train separating supplement $s_{i,j}$, since the higher the time separation between two trains i and j, the lower the expected knock-on delay. Figure 4 shows an example of a knock-on delay cost function. The horizontal axis shows the supplement between 0 and 60 minutes and on the vertical axis the expected knock-on time is given. Note that our MIP model optimises over all possible train orders. This means that when N trains use a common resource, for all train pairs, cyclically, $N(N-1)$ knock-on terms are added to the objective function. Knock-on costs are a major determinant for the optimal train orders, but major transfers will also play a role in this.

We could also consider the expected waiting time that passengers experience at their station of departure. This depends on the spreading between alternative trains in the timetable. In this paper we did not add these terms

Fig. 4 Shape of expected knock-on delay as a function of the chosen supplement. T is the timetable period, which is 60 minutes here. The vertical axis has no specific scale here.

to the objective function since our model developed to estimate this expected time does not scale well yet to networks with many trains (Sels et al, 2015a).

All types of objective function time terms described are seen as *objective time*. No subjective weights are added. This concludes our qualitative discussion of the objective function of our PESP MILP model representing the timetabling problem. In the next section, we apply our model to the train network of all passenger trains in Denmark and show the results.

3 Application to the Danish Railway System

Our complete method first constructs an event activity graph representing the train service network. Then, we route passengers over this graph to derive local passenger flow numbers for every ride, dwell and transfer action in this graph. We subsequently reschedule trains, deriving ideal arrival and departure times for all trains in all stations. We report results for each of these three phases.

3.1 Constructing the Event Activity Graph

For this project, Banedanmark started from the infrastructure they manage. This is 1956km or 79.5% of the the total of 2636km of railway track in Denmark. These tracks are visualised in figure 5. Subsequently, for an 'average' Wednesday in 2013, all trains running on this infrastructure were collected and slightly adapted, so that the timetable became exactly periodical with one hour. One representative hour for this network contains 84 passenger trains and 4 freight trains. Note that we do not schedule the suburban trains on the infrastructure of S-bane. The S-bane operates in the København area and is completely independent of the rest of the network, so it has no effect on our case. Some private operators run trains that briefly also use the Banedanmark infrastructure in just three places. These trains have not been modelled but are expected to have little influence on our main results. Freight trains were defined in the input only on sections where Banedanmark knows that there is a capacity bottleneck. For other sections, no freight trains were defined. It is assumed that they can be fitted between the scheduled passenger trains later.

We then generated the event activity network that corresponds to this service. This graph contains 88 trains, 264 stations, 3346 vertices and 9918 edges. The number of ride edges is 1541. Table 1 shows more problem instance statistics for this Danish event activity network.
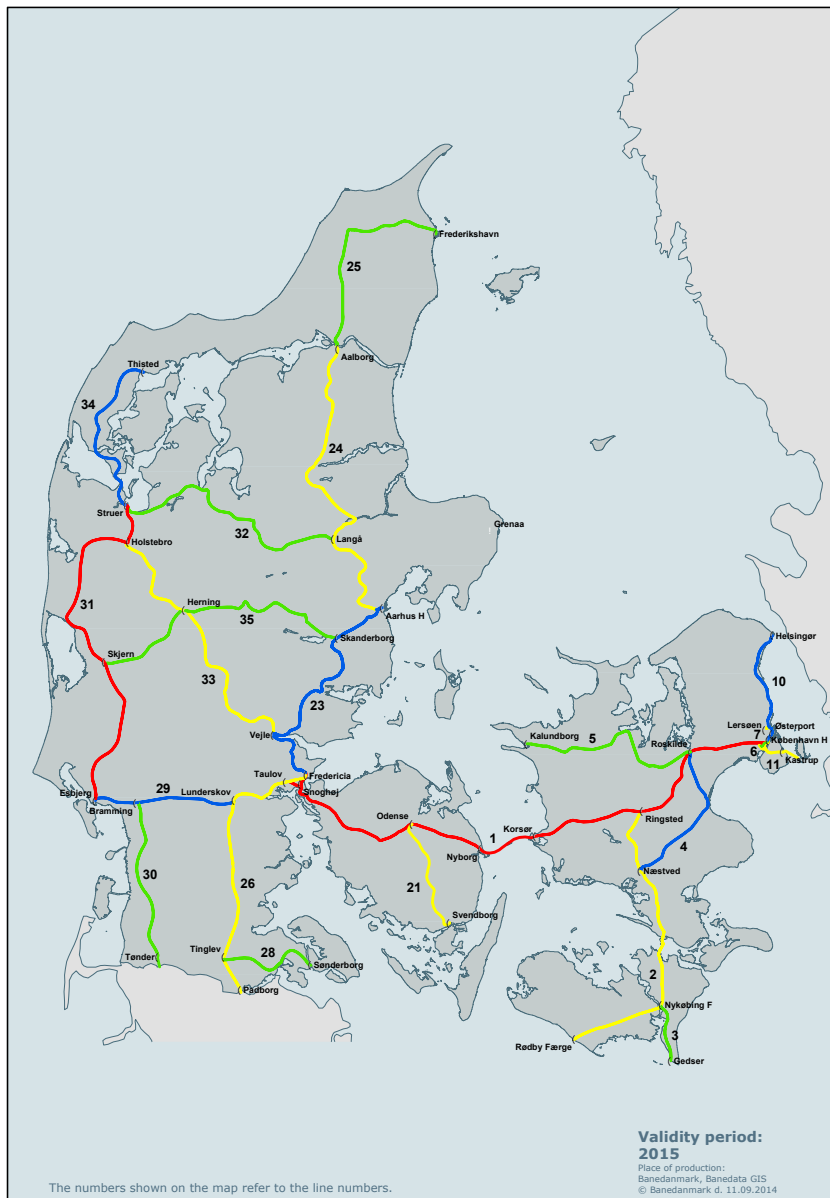
Fig. 5 Danish train infrastructure lines managed by Banedanmark

3.2 Routing: Reflowing

Now that the basic service graph is constructed, we mimmick the process were passengers decide what train to take if they go from an origin station (O) to a destination station (D). The number of commuters per day is 394377.

**Table 1** Graph and timetable MIP problem instance statistics

# ride edges =	1533
# dwell edges =	1445
# turn-around edges =	0
# knock-on(headway) edges =	13596
# major transfer edges =	4908
# model rows =	47335
# model columns =	32057
# model non-zero elements =	140516
# objective function terms for major flows =	16652
# objective function terms in post-optimisation evaluation =	21522

The morning peak OD matrix of these commuters is used to route passengers over this train service network, according to the routing algorithm described in Sels et al (2011). This is a modified Dijkstra algorithm implemented in C++. For efficiency, the modified Dijkstra algorithm was parallellised both on the core-level (using openMP, 2013) and the machine-level (using openMPI, 2014). For every OD-pair in the OD matrix, the best routings from O to D are calculated independently. First the modified Dijkstra algorithm is run to find the route with the lowest planned time, based only on the sum of minima for its ride and dwell actions. To avoid too many transfers in a route we penalise the choice of a transfer with 15 minutes. Note that the actual duration of a transfer is not known yet at this point. Next, all edges forming this route are eliminated from the graph and a new route search is performed. This route finding process is repeated until the new found route takes more than 20% more time than the first route found. At this point, it is assumed that no passengers will still opt for such a slower route. Passengers for a specific OD-pair are then distributed over the different OD-routes found, where more are assigned to the shorter routes than to the longer routes. Note that in our method, routing passengers comes before timetabling. This means that arrival and departure times are still unknown and so is their spreading out across one timetable hour. We simplify by assuming that these factors play no role in the passenger distribution over different routes for a given OD-pair (Jolliffe and Hutchingson, 1975). This assumption will be more realistic with good termporal spreading than with bad temporal spreading of alternative trains (Sels et al, 2015a). After the routing phase, which is parallellised for all OD-pairs, a non-parallellised merging phase, for each action (ride, dwell, transfer) on each link of the network is performed. Passenger numbers from the different OD-streams passing along an action are accumulated. We obtain the passenger number for every action (edge) in the event activity graph. Note that the freight trains in our system start in a technical station that passengers do not have access to. The freight trains also do not halt nor stop in passenger stations and so, in our routing algorithm, no passengers can get on or off these trains, as is the case in practice. This means that in our timetabling model, a freight train is treated like a passenger train with no passengers on, so it will

be of lower priority during scheduling. If one wants a higher importance, one could assign a virtual number of passengers to each freight train.

The results from the full passenger routing phase, accumulated per track section, are given graphically in figure 6. In this figure, the area of each circle
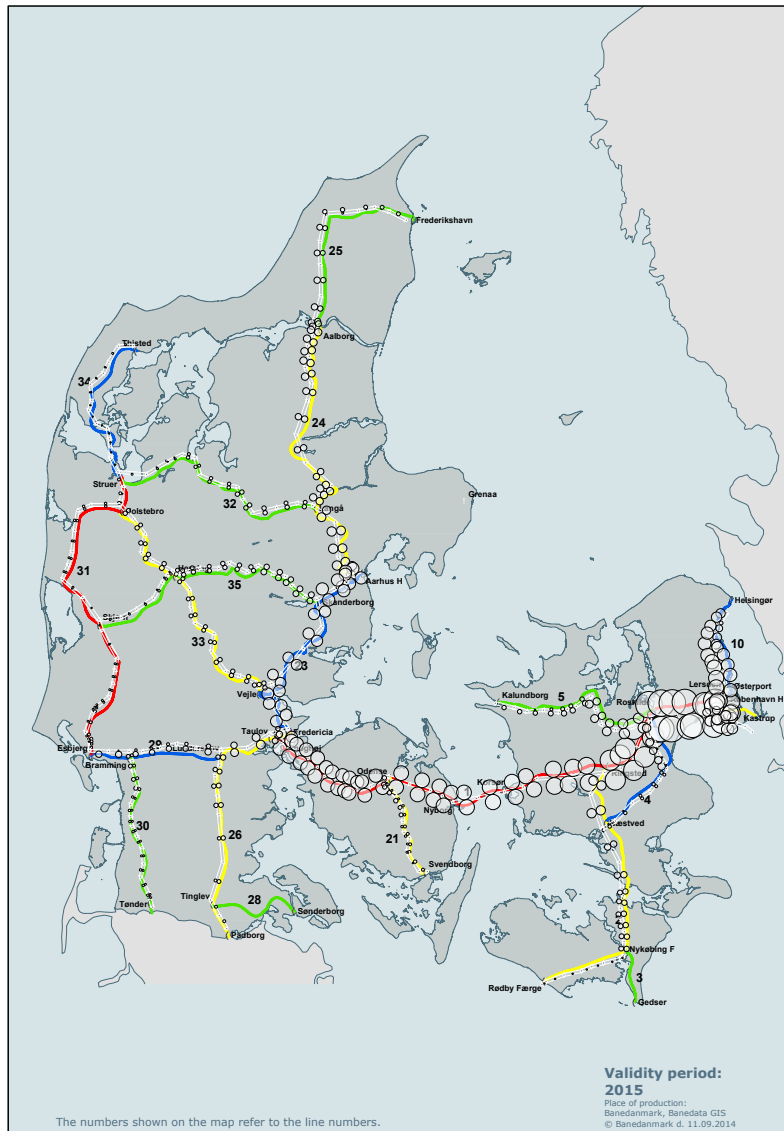
Fig. 6 Passenger flows in Denmark for a typical Wednesday morning peak

incident to a track section is proportional to the number of people traveling on trains that travel along that track section. It is clear that the set of trains in

the area around København transport the most passengers. All trains together going from Høje-Taastrup to Hedehusene, carry 29215 passengers in the morning peak. This is the maximum flow present in the graph. The second highest passenger flows occur on the tracks from København westwards to Odense and back and also from Fredericia North to Århus and back. It can be seen that the collected trains for other track sections in the rest of Denmark each transport a lot less passengers.

3.3 Scheduling: Retiming

Now that we know the number of passengers for each ride, dwell, transfer and knock-on action, we perform timetabling, according to the methodology described in section 2. We use the obtained local passenger numbers as fixed weights in the objective function.

4 Results

With different parameter settings, different MILP timetabling models were constructed. With each model, we construct a different timetable. Our software has a solver independent architecture, using the open source library milp-logic (Sels, 2012). This way, a simple solver setting and recompilation allows the software to call any solver supported by milp-logic. Currently, these are CPLEX, Gurobi, XPRESS. In this paper, we restrict ourselves to reporting of results with Gurobi. Each of our timetabling models was tackled by the MILP solver Gurobi version 6.0.0 on an Intel Xeon E31240 3.3GHz processor with 16GB of RAM. When constructing and optimising a MIP model, we noticed that computation times were sensitive to the amount of passenger flows we consider in the objective function. When all streams are considered, computation time becomes excessive so we defined a threshold of number of passengers. Streams with fewer passengers than this threshold are not considered in the objective function. The threshold of 210 passengers per morning peak gave manageable computation times. A further parameter is the required MIP gap. Setting this to 74% resulted in schedules with a lower total expected passenger time than the original schedule. Gap values lower than 74% result in better schedules but computation time also rises. For these parameter values 210 and 74% we get an optimised timetable. This is the timetable we report results for in sections 4.1 and 4.2. Section 4.1 describes that for this optimised timetable, there are no minimum headway time violations. Section 4.2 shows that large time supplements can be and are here assigned to train actions where no passengers are expected. Section 4.3 shows that for various parameter settings, the total passenger time in practice that is expected for the resulting optimised timetables, is always reduced compared to the current timetable.

4.1 No Collisions nor Headway Violations

The current and optimised timetables were verified by Banedanmark by visual inspection of space-time graphs per infrastructure line. Some examples of these graphs are given as figures 7, 8, 9 and 10.


Fig. 7 Space time graph for the original timetable for line 10


Fig. 8 Space time graph for the optimised timetable for line 10

Figure 7 shows the space-time diagram of trains running on the train infrastructure line 10 between København (KH) and Helsingør (HG) and back for the original timetable. Figure 8 shows the same trains but now for the optimised timetable. In figure 7 it can be seen that the original table gener-

ally leaves the required 3 or more minutes between each couple of subsequent trains except for 4 cases between København (KH) and Østerport (KK) and back as indicated by the red dashed circles C1 to C4. Indeed, in circles C1 and C4, train 828 (brown) and train 94423 (dark blue) only have a headway time of 2 instead of 3 minutes between them. The same happens between train 2514 (dark red) and train 74423 (semi-light blue) in circles C2 and C3. In the optimised timetable in figure 8, it can be seen that no such violations of the minimal headway time constraints of 3 minutes occur.

Fig. 9 Space time graph for original timetable for line 23

Fig. 10 Space time graph for the optimised timetable for line 23

Figures 9 and 10 show the space-time diagram of trains running on the train infrastructure line 23 between Fredericia (FA) and a Århus (AR) and

back, respectively for the original and the optimised timetable. One can verify that on this line, for both timetables, no single train collision nor violation of minimal headway time constraints occurs. For all other infrastructure lines, similar graphs were generated and verified as well and as such Banedanmark declared the optimised timetable as free of headway conflicts.

4.2 Large Dwell Times on Line 10 Explained

Figure 8 shows that, at the station Snekkersten (SQ), 4 trains heading for Helsingør (HG) are assigned large dwell times. These trains are (64421 (medium green), 72025 (light blue), 62029 (yellow-green) and 74423 (semi-light blue)) This is caused by the fact that our routing phase resulted in no passengers between Snekkersten (SQ) and Helsingør (HS). This can be seen in figure 6, where no white circle occurs between Snekkersten and Helsingør. This also means that these dwell times are not penalised in our objective function of our timetabling model. They can become arbitrarily large without having an effect on any passengers indeed.

Furthermore, it should be noted that our current timetable is only ideal for passengers traveling in the morning. Since one usually wants a timetable that is the same for morning and evening, one can express that by supplying an OD matrix that contains both morning and evening OD-pairs together. If then, all ride and dwell actions of all trains will have at least some passengers on them, in both directions, none of these dwell or ride times will stay unaccounted for in the objective function of our timetabling model. As such, all these actions will also have sensible supplements assigned to them.

Also note that in Snekkersten (SQ), in practice, there are not enough platform tracks in the station to allow simultaneous dwelling of 4 trains. Our timetabling model does indeed not take microscopic issues like this into account. Again, when some passengers would be assigned to these dwell actions, shorter dwell times will result and with that the number of simultaneously dwelling trains will most likely be significantly reduced.

4.3 Reduced Expected Passenger Time

By construction, our optimised timetables contain no single violation of hard (minimum run time, minimum dwell time, minimum headway time) constraints. For headway times this was illustrated in the previous sections graphically. In this section, we show that the optimised timetable also results in lower expected passenger time in practice than the original timetable. The relevant results are shown in table 2. Each of our timetabling models was tackled by the MILP solver Gurobi version 6.0.0 on an Intel Xeon E31240 3.3GHz processor with 16GB of RAM. Results for the different optimisations and their respective input parameter values are ordered from less to more demanding from top to bottom. By more demanding, we mean that either the required MILP gap (column 3) is lower or the number of transfers considered in the optimisation is higher or a combination of both. The transfer threshold (column 2) is the number of people that are required as minimum for a transfer to be considered

Table 2 Results for different timetable optimisations of all 88 hourly Danish trains. req. = required, obt.=obtained, exp. time = expected *passenger* time, red. = reduction, eval. = evaluation, orig.tt = original timetable, opt.tt = optimised timetable, rd. + dw. t = ride + dwell *train* time.

1 a (%)	2 transfer threshold	3 gap req. (%)	4 gap obt. (%)	5 solver time (s)	6 exp. time red.eval. (%)	7 missed transfers: orig.tt (%)	8 opt.tt (%)	9 planned rd.+dw. t red.(%)
2	420	75	74.92	19421	1.67	11.34	2.07	-4.88
2	420	74	73.63	62417	1.96	11.34	5.21	-3.95
2	210	79	78.07	1534	0.82	11.34	2.83	-7.77
2	210	77	76.62	2436	1.59	11.34	3.20	-4.89
2	210	75	74.96	2924	2.45	11.34	1.12	-3.08
2	210	74	73.83	**3922**	**2.90**	11.34	**2.45**	-2.53
2	210	73	72.96	20726	3.16	11.34	2.07	-2.05
2	195	76	≥76.8	≥101000				

in the optimisation. Column 6 shows the reduction in percent from original to optimised timetable of the expected time as evaluated over all streams, also the ones with fewer people than the threshold value. Column 7 shows the missed transfer probability in the original timetable as simulated over all streams and column 8 shows the same for the optimised timetable. Column 9 shows the reduction in percent of the planned ride and dwell supplements from the current to the optimised timetable.

We see that setting the transfer threshold to 420 makes that the solver spends a lot of time (19421 and 62417 seconds) before it finds a solution with an optimality gap below the required one. When the transfer threshold is lowered to 210 transfer passengers, resulting in more transfers considered in the optimisation, the model seems to become easier for Gurobi. When subsequently also lowering the required gap from 79% to 74% (column 3), timetable solutions are found within 1534 to 3922 seconds (column 5) and corresponding savings of total expected passenger time increase from 0.82% to 2.90% (column 6). Lowering the required gap further to 73% still improves the solution with a total reduction of expected passenger time of 3.16%, however, the computation time then increases significantly to 20726 seconds, being 5.76 hours. To test if lowering the transfer threshold further below 210 reduces computation time, we investigate whether a threshold of 195 combined with a not so demanding required gap of 76% gives us a good timetable quickly. The last line of table 2 shows that after 101000 seconds, no acceptable timetable solution was found yet, since the solver is still at a gap of 76.8%. So the value 210 as a transfer threshold somehow seems a good trade-off between giving Gurobi enough information about a good timetable and not too many terms in the objective function.

For the timetable that reduces the passenger time by 2.90% compared to the original one, we show the expected passenger time and its components graphically in figure 11. This figure stacks expected time components on top

Fig. 11 Reduction of expected passenger time of 2.90% compared to the original timetable.

of each other to reveal the total expected time for all passenger streams, large and small, for this optimised timetable. Expected time components can indeed be added together since all of them are expressed in the same units: (tenths of) passenger minutes. In figure 11, the left bar indicates the original timetable (orig) and the right bar indicates the optimised timetable (opt). The vertical dimension represents expected passenger time, also for its constituent components: ride (blue), dwell (yellow), transfer (orange), knock-on (purple). For dwell and transfer time, all ride time of the ride action preceding it, is convoluted with it, which is what the blue shading refers to. On the left of each bar, the percentages (orig.m and opt.m) indicate the ratio of the total *expected* passenger time part, that can be seen as the consequence of the planned minima (m), to its total bar height. Note that this part is equivalent to the *planned* passenger minimum time. On the right, the percentages (orig.s and opt.s) indicate the ratio of the total *expected* passenger time part, that can be seen as the consequence of the planned supplements (s), to its total bar height. This part is equivalent to the difference of the total *expected* time minus the total *planned* passenger minimum time. For each color, the minima are shown in a darker tone of the color and the supplements in a lighter tone of the same color. Figure 11 shows clearly that the obtained reduction of total expected time of 2.9% is caused by the net effect of three main changes. First, the amount of time spent in supplements on ride and dwell actions is significantly lowered from 7.51% to 4.57%. Second, the expected knock-on delay time is reduced from 3.14% to 2.59% of the total expected time. Third, the expected

transfer time is increased from 5.75% to 7.26% of the total expected time. In absolute terms, the transfer time increase is smaller than the sum of decreases in expected time spent in ride and dwell supplements and in knock-on events. This means the net result is a reduction in total expected passenger time.

We go back to table 2. For the best timetables found, its last column mentions that these possess between 3.08% and 2.05% more *train weighted planned* ride and dwell time than the original timetable. Even then, the total passenger time is reduced. This is possible due to a number of factors. Firstly, our method adds supplements to trains but weighs them by passengers. Secondly, supplements can cause extra robustness, so adding planned time can reduce experienced time in practice. Thirdly, classical manual timetabling uses rules of thumb like assigning a certain percentage of supplement *to each train*. To avoid knock-on delays, we expect these rules to perform worse than our rule of assigning supplements *between each couple of trains* sharing an infrastructure resource, even more so since we do this proportionally with the number of passengers on the second train and dependent on the expected delay distributions of both trains.

Table 2 also mentions that the expected missed transfer probability for all passenger streams together, both large and small, is 11.34% (column 7) for the original timetable while not more than 2.45% (1.12%, 2.45% and 2.07%, column 8) for our best three timetables. This is clearly a significant improvement that will be appreciated by the railway passengers. These results were obtained by a post optimisation calculation on the obtained timetables, for all passengers streams, small and large, where expected delays are accumulated and resulting in fractions of missed and non-missed transfers. For the original timetable the percentage is always the same, 11.34%, since the value of the transfer threshold plays no role is the missed transfer calculations. Indeed all passenger streams are considered here and not only streams with more than the number of passengers indicated by the transfer threshold.

4.4 Further Verification

Further verification of realistic parameter settings like the value of 'a' and the value of transfer minima is warranted for fair comparison with the current timetable. Also verification of other timetable quality criteria like the possible preference of some operators to avoid large inserted supplements, even for actions with very few passengers, is required and ongoing.

5 Conclusion

This paper demonstrates that our PESP based method with an objective function representing total expected passenger time in practice, improves the timetable for the whole train network of Banedanmark. Total passenger time in practice can be reduced by 2.9% and the average probability of missing a

transfer is reduced from 11.34% to 2.45%. The fact that, after our successful application to the Belgian train network, the application to a second country now delivers satisfying results as well, indicates that our approach is quite generally useful.

Thanks to the addition of a particular set of cycle constraints to the PESP model (Sels et al, 2015b), computation time stays limited to 65 minutes. This could lead to huge time savings in the current timetabling practice which, for the biggest part, is still carried out manually. Alternatively, the time spent on manual timetabling now, can instead be used to create more alternative line planning proposals which can be fed to our timetabling system. The line plan leading to the optimised timetable with the lowest total expected passenger time can then be selected. This would further improve passenger service.

6 Further Work

Even though the total expected passenger time of our optimised timetable is lower than the one for the original timetable, the total expected transfer time component of our optimised timetable increased. It would be interesting to see if our model could be adapted so that this expected transfer component is reduced while still also reducing the total expected passenger time.

Some degree of temporal spreading of alternative trains between origin and destination is beneficial to reduce the inter-departure waiting time for passenger travelling between these points. Also considering this inter-departure waiting time at the origin and inter-arrival-time at the destination would avoid potential bunching of trains and further generalise our method.

We now produce a timetable that respects headway time minima of 3 minutes everywhere in the network, which is the most common headway minimum value for macroscopic railway models. On a microscopic level, the actually needed headways can be derived from the blocking model (Hansen and Pachl, 2014) and depend on parameters like station infrastructure, train speed and train length. Per train pair, per station, the required minimum headway between these train pairs for that station can be calculated and these values can be substituted for the 3 minute macroscopic headway minima. When our method is used with these more accurate headway minimum values as input, a microscopically feasible timetable will result.

References

(2013) OpenMP Application Programmer Interface v4.0. URL http://www.openmp.org/mp-documents/OpenMP4.0.0.pdf

(2014) The OpenMPI API specification for MPI. URL http://www.open-mpi.org/doc/

Caprara A, Kroon L, Toth P (2011) Optimization Problems in Passenger Railway Systems. Wiley Encyclopedia of Operations Research and Management Science 6:3896–3905

Daamen W, Goverde R, Hansen I (2009) Non-Discriminatory Automatic Registration of Knock-On Train Delays. Networks and Spatial Economics 9(1):47–61

Dewilde T, Sels P, Cattrysse D, Vansteenwegen P (2013) Robust Railway Station Planning: An Interaction Between Routing, Timetabling and Platforming. Journal of Railway Transport Planning & Management (Young Railway Operations Research Award 2013 of IAROR) 3:68–77

Ferreira L, Higgins A (1996) Modeling Reliability of Train Arrival Times. Journal of Transportation Engineering 22(6):414–420

Goverde R (1998a) Optimal Scheduling of Connections in Railway Systems. Paper presented at the 8th WTCR, Antwerp, Belgium

Goverde R (1998b) Synchronization Control of Scheduled Train Services to minimize Passenger Waiting Times. Proceedings of 4th TRAIL Annual Congress, TRAIL Research School, Delft, The Netherlands

Goverde R, Hansen I (2000) TNV-Prepare: Analysis of Dutch Railway Operations Based on Train Detection Data. International conference on computers in railways VII:779–788

Hansen I, Pachl J (eds) (2014) Railway Timetabling and Operations, 2nd edn. Eurailpress

Jolliffe JK, Hutchingson TP (1975) A Behavioral Explanation of the Association between Bus and Passenger Arrivals at a Bus Stop. Transportation Science 9:248–282

Kroon L, Dekker R, Gábor M, Helmrich MR, Vromans M (2006) Stochastic Improvement of Cyclic Railway Timetables. Tech. rep., Erasmus Research Institute of Management (ERIM)

Kroon L, Dekker R, Vromans M (2007) Cyclic Railway Timetabling: A Stochastic Optimization Approach. Algorithmic Methods for Railway Optimization Lecture Notes in Computer Science pp 41–66

Kroon L, Huisman D, Abbink E, Fioole PJ, Fischetti M, Maróti G, Schrijver A, Ybema R (2009) The New Dutch Timetable: The OR Revolution. Interfaces 39:6–17

Labermeier H (2013) On the Dynamic of Primary and Secondary Delay. Proceedings of 5th International Seminar on Railway Operations Modelling and Analysis (IAROR): RailCopenhagen2013, May 13-15, Copenhagen, Denmark

Liebchen C (2007) Periodic Timetable Optimization in Public Transport. Operations Research Proceedings 2006:29–36

Meng Y (1991) Bemessung von Pufferzeiten in Anschlüssen von Reisezügen. Veröffentlichungen des Verkehrswissenschaftlichen Institutes der Rheinischen-Westfälischen Technischen Hochschule Aachen 46:29–36

Nachtigall K (1996) Periodic network optimization with different arc frequencies. Discrete Applied Mathematics 69:1–17

Peeters L (2003) Cyclic Railway Timetable Optimization. PhD thesis, Erasmus Research Institute of Management (ERIM), Rotterdam

Schrijver A, Steenbeek A (1993) Spoorwegdienstregelingontwikkeling (Timetable Construction). Technical Report, CWI Center for Mathematics and Computer Science, Amsterdam

Schwanhäußer W (1974) Die Bemessung von Pufferzeiten im Fahrplangefüge der Eisenbahn. Veröffentlichungen des Verkehrswissenschaftlichen Institutes der Rheinischen-Westfälischen Technischen Hochschule Aachen 20

Sels P (2012) milp-logic: a C++ MILP Solver Abstraction Layer with a C++ Boolean Modelling Layer on Top. URL http://github.com/PeterSels/milp-logic/

Sels P, Dewilde T, Cattrysse D, Vansteenwegen P (2011) Deriving all Passenger Flows in a Railway Network from Ticket Sales Data. Proceedings of 4th International Seminar on Railway Operations Modelling and Analysis (IAROR): RailRome2011, February 16-18, Rome, Italy

Sels P, Dewilde T, Cattrysse D, Vansteenwegen P (2013a) A Passenger Knock-On Delay Model for Timetable Optimisation. In: Proceedings of the 3rd International Conference on Models and Technologies for Intelligent Transport Systems (MT-ITS 2013), December 2-4, Dresden, Germany., pp 1–10

Sels P, Dewilde T, Cattrysse D, Vansteenwegen P (2013b) Expected Passenger Travel Time as Objective Function for Train Schedule Optimization. Proceedings of 5th International Seminar on Railway Operations Modelling and Analysis (IAROR): RailCopenhagen2013, May 13-15, Copenhagen, Denmark

Sels P, Dewilde T, Cattrysse D, Vansteenwegen P (2015a) Optimal Temporal Spreading of Alternative Trains in order to Minimise Passenger Travel Time in Practice. Proceedings of 6th International Seminar on Railway Operations Modelling and Analysis (IAROR): RailTokyo2015, March 23-26, Tokyo, Japan

Sels P, Dewilde T, Cattrysse D, Vansteenwegen P (2015b) Reducing the Passenger Travel Time in Practice by the Automated Construction of a Robust Railway Timetable. Submitted to Transportation Research Part B

Serafini P, Ukovich W (1989) A Mathematical Model for Periodic Scheduling Problems. SIAM Journal on Discrete Mathematics 2:550–581

Sparing D, Goverde R, Hansen I (2013) An Optimization Model for Simultaneous Periodic Timetable Generation and Stability Analysis. Proceedings of 5th International Seminar on Railway Operations Modelling and Analysis (IAROR): RailCopenhagen2013, May 13-15, Copenhagen, Denmark

Vansteenwegen P, Van Oudheusden D (2006) Developing Railway Timetables that Guarantee a Better Service. European Journal of Operational Research 173(1)(1):337–350

Yuan J (2006) Stochastic modelling of train delays and delay propagation in stations. PhD thesis, Technical University of Delft