View metadata, citation and similar papers at core.ac.uk

brought to you by CORE

provided by Online Research Database In Technology

Technical University of Denmark

An Integrated Rolling Stock Planning Model for the Copenhagen Suburban Passenger Railway

Technical Report

Thorlacius, Per; Larsen, Jesper; Laumanns, Marco

Publication date:
2015

Document Version
Peer reviewed version

Link back to DTU Orbit

Citation (APA):
Thorlacius, P., Larsen, J., & Laumanns, M. (2015). An Integrated Rolling Stock Planning Model for the Copenhagen Suburban Passenger Railway: Technical Report. DTU Management Engineering.

DTU Library
Technical Information Center of Denmark

General rights
Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

An Integrated Rolling Stock Planning Model for the Copenhagen Suburban Passenger Railway

Technical Report

Per Thorlacius[*], Jesper Larsen[†], Marco Laumanns[‡]

November 13, 2015

Abstract

A central issue for operators of passenger railways is providing sufficient number of seats for passengers while at the same time minimising operating costs. This is the task of rolling stock planning. Due to the large number of practical, railway specific requirements that a rolling stock plan has to take into account, rolling stock plans are often constructed in a step-by-step manner, taking some requirements into consideration in each step. This may make it difficult in the final step to produce a plan that is feasible with regard to all of the requirements and at the same time economically attractive.

This paper proposes an integrated rolling stock planning model that simultaneously takes into account all practical requirements for rolling stock planning at DSB S-tog, the suburban passenger train operator of the City of Copenhagen. The model is then used to improve existing rolling stock plans using a hill climbing heuristic.

Experiments show that the heuristic used in the integrated rolling stock planning model is able to produce feasible solutions within minutes of computation time starting from infeasible rolling stock plans. Furthermore, the heuristic is able to improve the economic attractiveness of typical rolling stock plans with an average of 2%.

Keywords: Integrated rolling stock planning, data modelling, heuristics, resource constrained shortest paths, passenger railway.

[*]Corresponding Author, Danish State Railways (DSB), Dept. of Planning, Telegade 2, DK-2630 Tåstrup, Denmark

[†]Technical University of Denmark (DTU), Dept. of Management Engineering, DK-2800 Lyngby, Denmark

[‡]IBM Research, Säumerstrasse 4, CH-8803 Rüschlikon, Switzerland

1 Introduction

1.1 Background and Terminology

Rolling stock planning is the process a passenger railway operator performs in order to plan how to use the rolling stock for the conveyance of passengers. The goal of the rolling stock planning process is to provide sufficient seats for passengers while at the same time keeping operating costs as low as possible. This goal is of course a highly important matter for operators of passenger railways since it is the core question of their very existence: *Can the passenger railway convey its passengers at an acceptable price?*

A passenger railway operates a *timetable* of *train services* for the conveyance of passengers for *revenue*. Rolling stock planning is performed by assigning individual *train units* to the train services from the timetable.

When producing rolling stock plans for a passenger railway, a large number of practical, railway specific requirements need to be taken into account. These requirements relate to the railway infrastructure, the timetable, the rolling stock itself, the passenger demand, maintenance scheduling and a large number of other aspects of the railway operation.

Due to the large number of practical, railway oriented requirements and their complexity, rolling stock planning is often performed in a step-by-step manner, taking only some of the many requirements into consideration in each step. This is also the case in the rolling stock planning system currently used at DSB S-tog, the suburban passenger train operator of the City of Copenhagen. DSB S-tog is considered as case study for this paper.

In the rolling stock planning system of DSB S-tog, as it is typical for the industry, the first step is to decide how much seating capacity should be allocated to each train service. This step is called *composition planning*. Based on this, in the next step, individual train units are assigned to train services in a process called *rotation planning*. Finally, in the last step it is decided where the train units are to be parked in the *depots* when not in use. This step is called *depot planning*. Needless to say, the step wise approach may produce plans that are neither optimal nor feasible.

For DSB S-tog this is especially the case due to the very limited space in the depots where train units are parked when not in use. For this reason, the most constraining requirement for the rolling stock planning at DSB S-tog is that of being able to move the train units in and out of the depots. Planning this as the **last** step may prove highly problematic, since decisions taken in the earlier steps may limit the degrees of freedom for the depot planning steps to an extent that no feasible solution can be found. Such infeasible plans will have to be corrected manually, most often incurring extra cost.

Other suburban passenger train operators may have similar, challenging conditions that make sequential planning equally problematic. For this reason, an integration of **all** the different rolling stock planning processes is essential if an automated model is to produce plans that are usable in practice. This is achieved by the integrated rolling stock planning model proposed in this paper.

The combined process of composition planning, rotation planning and depot planning is called *circulation planning.* The circulation planning phase of rolling stock planning has a tactical scope and is conducted months before the plan is set into motion. The process of setting a circulation plan into motion is called *train unit dispatching.* This is the operational, short-term or real-time phase of rolling stock planning where last minute changes are made based on which physical train units are available, whether delays or disruptions have occurred, etc.

1.2 Literature Review

Until recently, operations research (OR) techniques have been applied to a wide range of specific problems in the railway industry, which are summarised in various surveys [16, 3, 25, 13, 28]. The challenge for the adaptation of OR techniques to the railway industry now no longer seems to lie in finding solutions to each specific problem, but much more in integrating the individual solutions to the (often highly interconnected) specific problems into holistic, integrated models. By integrating the specific models with each other, sub-optimal solutions can be avoided. The tendencies for the integration of models are currently also seen in the airline industry [38].

Table 1 shows an overview of characteristics of selected and reviewed, recent literature for rolling stock planning. The characteristics are grouped as follows: The overall topic of article; The railway planning processes it addresses; The type of the model proposed; The properties of the model graph (all reviewed models feature a graph); The railway specific requirements the model integrates; The objective of the model; And finally, the solution method applied.

As may be seen from Table 1, a large portion of the reviewed methods use an *arc based multi-commodity flow* or similar modelling scheme. In such a scheme the flow of train units or locomotives is modelled in a flow graph, with flow conservation constraints on each vertex of the graph making sure the flow into the vertex equals the flow out of it. Arc based flow models are typically relatively low in complexity, a presumed reason for their widespread use. In arc based flow models, however, it adds to complexity to model sub-path constraints such as recurring maintenance at regular distance intervals.

In *path based multi-commodity flow models* on the other hand, each potential sequence of movements of the individual train unit or locomotive is modelled (e. g. by enumeration), making it easier to also take recurring distance related constraints such as maintenance into account.

In Table 1, the literature reviewed is also categorised according to the properties of the graphs involved in the models. As may be seen, most models use a *space-time graph* type (also called time-expanded graph type), where each vertex in the graph is an event in space and time, e. g. the arrival of a train at a given time at a given station (in space). Correspondingly, arcs in a space-time graph may e. g. represent train services. Such a graph is also an *event-activity graph*, referring to the vertices as events and the arcs as activities. This type of graph is also well known in the airline industry [37, 6, 35, 38].

Some authors use the *edge-to-vertex dual* or *line graph* [24] graph type, conjugated from the space-time graph mentioned before. In the line graph type, vertices represent train services whereas arcs may e. g. represent the possibilities of a train unit to perform train services in sequence.

Similar "line graph" style graph types are used to model train composition changes at stations, the approach direction in depot planning and maintenance constraints. In two papers a hypergraph is used [8, 9].

As may be seen from Table 1, the models in the reviewed literature integrate a different number of railway-specific requirements, with a slight tendency that recent models integrate more requirements than earlier ones.

In most of the reviewed models the objective is to minimise operational costs and/or penalties. Some models are used to minimise seat shortages or the number of train units needed etc. Since depot planning is mainly a matter of fixed costs, some depot planning models have only feasibility as their objective.

All solution methods applied in the reviewed literature involve commercial solvers and/or heuristics. In addition, decomposition techniques such as column generation, branch and price and Benders decomposition are used in some cases.

The following is a brief overview of the size of the data instances used in the experiments

Table 1: Overview of characteristics in selected and reviewed, recent literature specific to rolling stock planning. Characteristics of the integrated rolling stock planning model proposed in this paper are listed at the bottom for comparison. *) Common requirements include: Timetable, overall infrastructure, rolling stock, passenger or freight demand requirements. Note that subsidiary requirements like cyclicality, robustness and disruption recovery with minimal changes have been omitted.

		Topic		Process				Model Type						Graph Properties					Requirements Integration							Objective							Solution Method				
Authors	Year	Passenger railway	Freight railway	Composition planning	Rotation planning	Depot planning	Train unit dispatching	Assignment etc.	Set partitioning	Arc multi-commodity flow etc.	Hyper arc multi-commodity flow	Hamiltonian cycle	Path multi-commodity flow	Vertex is space-time event	Vertex is train service	Vertex is composition	Vertex is approach type	Vertex is maintenance	Common requirements *)	Train composition order	Maintenance etc. (by time)	Maintenance etc. (by distance)	Personnel on duty	Depot capacity	Depot topology	Minimise seat shortage	Minimise cost	Minimise penalties	Minimise # of train units etc.	Maximise benefits	Minimise # of depot movements	Feasibility	Commercial solver etc.	Heuristics etc.	Column generation etc.	Branch and price etc.	Benders decomposition
Cordeau [14]	2000	●		●	●					●				●					●								●						●				●
Cordeau et al. [15]	2001	●		●	●					●				●				●	●		●	●					●						●				●
Brucker et al. [10]	2003	●		●	●					●				●					●					●			●	●						●			
Abbink et al. [2]	2004	●		●				●											●							●							●				
Razen [36]	2004		●	●	●					●				●					●										●				●	●			
Ahuja et al. [4]	2005		●	●	●					●				●					●								●	●					●	●			
Freling et al. [21]	2005	●				●			●							●	●			●				●	●			●					●		●		
Alfieri et al. [5]	2006	●		●	●					●				●	●	●			●	●							●						●	●			
Fioole et al. [19]	2006	●		●	●					●				●	●	●	●		●							●			●		●		●	●			
Føns [20]	2006	●				●		●								●	●			●				●	●							●	●			●	
Maróti and Kroon [32]	2007	●					●			●				●					●		●	●						●					●	●			
Jha et al. [27]	2008		●	●	●								●	●					●								●	●					●	●			
Kroon et al. [29]	2008	●				●		●								●	●			●				●	●			●					●				
Peeters and Kroon [34]	2008	●		●	●					●					●	●			●	●							●						●			●	
Vaidyanathan et al. [40]	2008		●	●	●					●				●					●		●						●	●					●				
Dirksen [18]	2010	●				●		●												●		●		●	●							●	●				
Beygo [7]	2011	●				●			●											●				●	●							●	●				
Borndörfer et al. [8]	2011	●		●	●						●			●	●	●		●	●			●					●	●					●	●			
Cadarso and Marín [12]	2011	●		●	●					●				●					●					●			●	●					●				
Borndörfer et al. [9]	2012	●		●	●						●			●	●	●		●	●			●		●			●	●					●	●			
Cacchiani et al. [11]	2012	●		●	●			●							●	●			●										●					●			
Giacco et al. [22]	2014	●		●	●							●			●			●	●		●	●							●				●				
Haahr et al. [23]	2014	●					●						●	●					●					●			●	●					●		●		
(Model proposed in this paper)	2015	●		●	●	●	●						●	●					●	●	●	●	●	●	●		●	●		●				●			

reported in the reviewed literature. [14, 15] assign train units to 300 train services for the period of one week. [10] assign 6 different types of train units to 200 train services. [2] circulate train units for 188 passenger train services, while [36] circulates train units for 86 freight train services. In [4] 3,324 train services have up to 1,600 train units assigned to them by 5 types. [21] distribute 600 train compositions with 1,100 train units on 19 depot tracks. In [5] 12 train compositions have train units assigned to them on an intercity line with 30 min. frequency. [19] assign 85 train units in 2 categories to 67 train services. [32] perform maintenance planning of 47 train units serving 800 train services per day for a period of up to 5 consecutive days. [27] assigns 1,200 collections of train units to 350 freight train services. [28] assigns up to 600 train services with 1,100 train units to 19 depot tracks. In [40] 388 freight train services have 6 train unit types in 8 possible compositions assigned to them. [8, 9] circulate the entire German ICE high speed fleet for a week, resulting in a model graph with more than 60 million hyperarcs in it. [12] assign rolling stock to up to 400 train services, and [11] up to 76 train units in up to 10 categories to up to 600 train services. [22] plans maintenance for a timetable with up to 104 train services.

In comparison, DSB S-tog has approximately 1,350 train services per day with 122 train units in operation, routed to 53 depot tracks. In the integrated rolling stock planning model proposed here, this yields a model graph with approximately 28,000 arcs. The following references use DSB S-tog data: [20, 18] perform depot planning by each depot, [7] circulates rolling stock and [23] recover from a disrupted rolling stock plan (but disregard the individual depot tracks).

To the best of our knowledge, no integrated planning model for the requirements we are considering has been worked out or used, neither in the literature nor in industry practise.

On a general note, the terminology used in rolling stock planning literature shows great variation. In some articles a terminology is used that has weak connotations to the actual railway operation. Needless to say, this neither facilitates the understanding nor the comparison of methods across literature.

1.3 Scientific Contribution

The scientific contribution of the proposed integrated rolling stock planning model is that of integrating into one process, processes which are normally solved separately. In particular, the integration of train unit to train service assignment, maintenance planning (by distance) and depot planning is not known from literature. The processes are integrated using a heuristic framework.

A further scientific contribution is the development of special side constraints to a (resource constrained) shortest path algorithm that can handle the individual order of train units in train compositions so that no train unit will obstruct the movement of another. We call this new concept *unit order flow conservation.*

In addition to this, experiments are conducted using real data instances with all the peculiarities of actual production data in order to prove the scientific viability of the model in realistic conditions.

1.4 How This Paper Is Structured

Section 2 formulates the problem to be solved and presents an overview of the solution concepts. Moreover, an overview of the proposed model is presented, along with a mathematical formulation. Next, the different parts of the solution approach are presented in detail: Section 3 describes the data model for the timetable and the infrastructure (the space-time graph), Section 4 presents the data model for train units, Section 5 the path finding algorithm and Section 6

the surrounding heuristic framework. Section 7 describes the real-world data instances used in the evaluation of the integrated model and the results obtained. Lastly, Section 8 discusses the implications of the proposed methods and outlines further research.

2 The Integrated Rolling Stock Planning Problem

2.1 Problem Formulation

Seen from an overall business perspective, the goal of the rolling stock planning problem is to provide seats for passengers while at the same time keeping operational costs to a minimum. Seen from a more detailed operational perspective, rolling stock planning is about deciding which individual train unit should be assigned to which train service. By doing so, one has implicitly assigned seating capacity to the train services. At the same time, it must also be decided when and where train units should be parked at the depots when not in use. All these decisions must be taken in such a way that operational costs are minimised and all of the practical, railway specific requirements are adhered to.

2.2 Solution Concepts

The underlying solution idea to the problem presented here is to look at the above mentioned assignment of train units to train services in an aggregated way: A rolling stock plan may be entirely described by the movement of its train units in space and time. The movement of a particular train unit in space and time for a particular period of time is called a *train unit trajectory*. A typical train unit trajectory starts off with the train unit being parked at a depot track before being *shunted* to the platform. From the platform, the train unit may then be assigned to a revenue train service starting at this *origin station*. At the *terminal station* of that particular train service, the train unit may *turn around* to be assigned to another train service in the opposite direction. This train service may be a *non-revenue train service* running without passengers with the purpose of positioning the train unit for later use. Typically, a train unit trajectory ends by having the train unit being shunted back into a depot track for parking for the remaining time of the given period.

As such, a train unit trajectory describes which train services the train unit in question is assigned to for the given period, including information of used turnaround times between train services, and at which depot tracks the individual train unit is parked when not in use. Formulated in the context of this solution idea, one can say that the rolling stock planning problem is to decide the individual train unit trajectories of the train units, thereby offering enough seating capacity for the passengers and at the same time keeping cost at a minimum and adhering to all railway specific requirements.

To be able to find new candidate train unit trajectories that are attractive (that is, new ways that the individual train units should move in space and time), we need a measure of the attractiveness of each train unit trajectory. This measure is called the *additional net value*, defined as the additional benefit that may be achieved by assigning a train unit of a given type to perform the operations represented by the train unit trajectory in question, minus the incurred penalties and factual costs for doing so. Penalties are awarded for undesirable aspects of the rolling stock plan.

A positive additional net value for a given train unit trajectory indicates that there is good "value for money" in letting the train unit perform the given train unit trajectory, since the

benefits of doing so outweigh the costs and penalties. A negative value would indicate that the costs and penalties outweigh the benefits, in most cases an unattractive option.

2.3 Requirements Overview

The following is a brief overview of the practical, railway specific requirements for rolling stock planning at DSB S-tog. For a list of all requirements, see the left part of Table 4 on page 11. For a full description of all the requirements, see [39].

In the long term circulation planning phase of rolling stock planning, the following requirements must be taken into account: The physical railway **infrastructure** must be adhered to, e. g. depot track capacities, the rules of the train control system and the order in which train units may be parked so as not to obstruct each other's movements; All trains of the **timetable** must have a least one train unit assigned; Only the available **rolling stock** can be used in the plan; The plan should provide seating capacity according to the **passenger demand** and provide an even distribution of flexible space for bicycles etc.; Planned shunting operations in the depot should have sufficient **personnel on duty**; Train units must undergo **interior** and **exterior cleaning**, **surface foil application** and **winter preparedness** treatment at regular time intervals; At regular service distance intervals train units must undergo **scheduled maintenance** etc., and consumables such as **friction sand** must be refilled; Certain train services must have train units with additional **train control system equipment** installed, special **passenger counting equipment** installed and/or perform predefined **exposure of commercials**.

In the short-term or real-time train unit dispatching phase of rolling stock planning, additional requirements include: **Exterior graffiti removal** and **unscheduled maintenance** on demand and sometimes within a given time frame; Make available train units to meet **surveillance video requests from the police** within a given time frame.

2.4 Model Overview

The integrated rolling stock planning model proposed here integrates all the mentioned requirements using four main components. The first two components constitute a data model for the rolling stock plan. The last two components are algorithms applied to modify a given rolling stock plan in order to improve it. The four components are:

1. **The combined timetable and infrastructure data model:** A space-time graph with extended arc and vertex attributes, describing the timetable, the infrastructure, passenger demand, personnel on duty, which train service has which train unit assigned to it and in which individual, relative order, etc. This component is described in Section 3;
2. **The data model for train units**, interconnected with the space-time graph, describing the activities of the train units, e. g. which train unit is assigned to which train service. This component is described in Section 4;
3. **A special-purpose resource constrained shortest path algorithm with side constraints** operating on the space-time graph. This algorithm is used to find new candidate train unit trajectories taking into account the maximum service distance a train unit may perform as a resource constraint. As a side constraint, the individual, relative position of the train unit in relation to the other train units in the space-time graph is handled, determining which movements the train unit can perform based on its relative position. Also the flexible space distribution is handled as a side constraint. This whole component is described in Section 5;

4. **A heuristic framework** to accept or reject the candidate train unit trajectories found using the previously described components. The overall concept of the heuristic framework is to remove a number of train unit trajectories from an existing rolling stock plan and then, one by one, to create a new trajectory and insert it into the plan. The newly inserted trajectories are accepted if they produce an increase in the objective function value; if not, they are rejected, and the previous ones are re-inserted. The heuristic component is described in Section 6.

For an overview of the first three components, their aspects and which requirements they implement, see Table 4 on page 11.

2.5 Mathematical Formulation

The mathematical formulation of the proposed integrated rolling stock planning model presented here is based on sets (with corresponding indices) and parameters as defined in Tables 2 and 3.

The heuristic component of the integrated rolling stock planning model (Component 4 in the overview in Section 2.4) is governing the program flow. It works by (in each iteration) selecting k number of train units U^* and removing each of their original train unit trajectories $j_u^- \in J$ (Component 2) from the graph G (Component 1). Next, for each train unit $u \in U^*$ a new train unit trajectory $j_u^+ \in J$ (Component 2) is then found using the shortest path algorithm (Component 3). This train unit trajectory is then inserted into the graph. (As will be seen in Section 5, the path finding algorithm takes into account all practical, railway oriented requirements, including train unit order, so no additional feasibility check is needed prior to the successive insertion of each newly found train unit trajectory.) This process is repeated for all train units in U^*.

Next, the heuristic component checks the change in objective value for the iteration in question, the *iteration net value increase* z_Δ. If z_Δ is positive, each newly created train unit trajectory j_u^+ for the selected train units U^* is kept in the graph G. If z_Δ is not positive, for each train unit $u \in U^*$, the newly found train unit trajectory j_u^+ is removed from the graph, and the original j_u^- reinserted, equivalent to the changes in the current iteration being rolled back, i. e. the train units U^* having their original train unit trajectories reinstated.

3 Timetable and Infrastructure Data Model

The first component of the integrated rolling stock planning model is the combined data model for the timetable and the railway infrastructure in the form of a directed, acyclic, space-time graph G. Space-time graphs are well known in the railway industry, the first use for timetabling is attributed to French engineer Ibry prior to 1885 [31]. Space-time graphs are also known from other industries, for an example of a recent application in the airline industry, see [6].

In the proposed space-time graph, the arcs A represent the possibility of a train unit to move in space and time or in time only. The vertices V represent space and time events, that is, points in space and time where a train unit may perform different movements later on. For instance, a train unit arriving as a revenue train service to one of the platforms of a station is an event, after which the train unit may either stay at the platform and turn around to the next departure of a train service (one outgoing arc from that vertex) or be shunted to a depot track (another outgoing arc from the same vertex).

A schematic illustration of the principles in the space-time graph is shown in Figure 1.

Table 2: Sets and their corresponding indices in the mathematical formulation of the integrated rolling stock planning model, their domains and definitions, ordered alphabetically by symbol. Sets have symbols in upper case, and their respective indices have the same symbol in lower case without subscripts or superscripts.

Symbol	Description	Index, domain, definition
A	Arcs in the space-time graph G, each arc going from one vertex to another. As such each arc also has a corresponding time interval $p \in P$	$a \in A$; $a = (v_1, v_2)$ $v_1, v_2 \in V$
A_j	The arcs of train unit trajectory j	$a \in A_j \subset A$
G	The directed and acyclic space-time graph with vertices V and arcs A. The graph has extended attributes as described in Section 3 on page 8	$G = (V, A)$
I	Train unit types	$i \in I = \{\frac{1}{2}, 1\}$
J	All possible train unit trajectories. A train unit trajectory is a path through the space-time graph G representing the movement in space and time of a train unit	$j \in J$ $j = a_1, a_2, ..., a_{\|j\|}$ $a \in A$
P	All possible time intervals, a time interval being a sorted 2-tuple of point in time. Each arc $a \in A$ represents a time interval	$p \in P$; $p = (t_1, t_2)$ $t_1, t_2 \in T$ $t_1 < t_2$
Q	Points in space, being the union of each depot track at every station, each side track at every station, all platform tracks [as a whole] at every station	$q \in Q$
T	Points in time	$t \in T$
U	Individual train units currently available	$u \in U$
U^*	The set of train units selected for train unit trajectory subtraction, creation and addition in the heuristic	$u \in U^* \subset U$ $\|U^*\| = k$
V	Vertices in the space-time graph G, each vertex being a *point in space, point in time* tuple	$v \in V$; $v = (q, t)$ $q \in Q$; $t \in T$

Table 3: Parameters in the mathematical formulation of the integrated rolling stock planning model, their domains and definitions, ordered alphabetically by symbol. All parameters have symbols in lower case.

Symbol	Description	Domain, definition
$b^-(a,u)$	Benefits lost by subtracting train unit u from the arc a	$\in \mathbb{R}_0^+$
$b^+(a,u)$	Benefits gained by adding train unit u to the arc a	$\in \mathbb{R}_0^+$
$c^-(a,u)$	Costs saved by subtracting train unit u from the arc a	$\in \mathbb{R}_0^+$
$c^+(a,u)$	Costs added by adding train unit u from the arc a	$\in \mathbb{R}_0^+$
j_u^-	The original train unit trajectory belonging to train unit $u \in U^*$ scheduled for subtraction from the graph G	$j_u^- \in J$
j_u^+	The new train unit trajectory belonging to train unit $u \in U^*$ scheduled for addition into the space-time graph G	$j_u^+ \in J$
k	The number of train units to select in order to subtract their original train unit trajectories and add their newly found train unit trajectories	$k \in \mathbb{N}_1$ $k = \lvert U^* \rvert$
n	Number of test runs performed on each data instance for algorithm performance testing, see Table 5 on page 27	$n \in \mathbb{N}_1$
$p^-(a,u)$	Penalties saved by adding **or** subtracting train unit u to or from the arc a	$p^-(a,u) \in \mathbb{R}_0^+$
$p^+(a,u)$	Penalties awarded by adding **or** subtracting train unit u to or from the arc a	$p^+(a,u) \in \mathbb{R}_0^+$
z_Δ	The iteration net value increase, calculated according to Equation (1) on page 22 for the subset of trajectories that have been changed in the iteration	$z_\Delta \in \mathbb{R}$
$z^-(j_u^-)$	The subtractional net value, i. e. the net value of subtracting train unit trajectory j_u^- from the space-time graph G, calculated according to Equation (2) on page 22 before the train unit trajectory is subtracted	$z^-(j_u^-) \in \mathbb{R}$
$z^+(j_u^+)$	The additional net value, i. e. the net value of adding train unit trajectory j_u^+ into the space-time graph G, calculated according to Equation (3) on page 22 before the train unit trajectory is added	$z^+(j_u^+) \in \mathbb{R}$

Table 4: Overview of the requirements for rolling stock planning at DSB S-tog and how they are implemented in the integrated model.

		Timetable and infrastructure data model (Section 3)															Train unit data model (4)			Path finding algorithm (sect. 5)	
		Arc attributes										Vertex attributes									
Requirement Category	Requirement Detail	Length of track repr. by arc	Service Distance	Assigned train units limit	Assigned train composition	Transition positions	Depot drivers required	Train Line	Additional/subtract. benefit	Additional/subtract. cost	Additional/subtract. penalty	Limit on # of shuntings	Assigned # of shuntings	Date and time	Depot drivers on duty (aggr.)	Graph Topology	Allowed on lines	Start/finish time and space	Service distance limit	Resource constraints	Side constraints
Infrastructure	Adhere to track length capacities	●			●																
	Handle order of train units in train composition				●	●															●
	Use platform tracks for temporary parking										●					●					
	Use side tracks for temporary parking															●					
	Adhere to train control system rules				●	●															
	Adhere to coupling and decoupling rules				●	●															●
	Keep train unit balance in depot over time																	●			
	Only one shunting per arrival/departure											●	●								
	Handle split depots and track usage rules										●					●					
Timetable	Assign train units to all revenue train services				●						●										
	Enable non-revenue services for positioning															●					
	Adhere to braiding and train service seq. rules															●					
Rolling Stock	Adhere to platform lengths by train line	●						●													
	Adhere to rules on # of train units per train			●	●																
	Handle train composition flexible space distribution				●	●															●
Passenger Demand	Provide seats according to demand								●												
Personnel on Duty	Perform shuntings only when personnel avail.						●						●	●	●						
Scheduled Maintenance	Get train unit to workshop within distance limit		●																●	●	
	Even out the flow of train units to workshop		●															●	●	●	
Unscheduled Maintenance	Get train unit to workshop within time limit																	●			
Friction Sand	Get train unit to facility within distance limit		●																●	●	
Exterior Cleaning	Get train unit to workshop within time limit																	●			
Graffiti Removal	Get train unit to workshop within time limit																	●			
Interior Cleaning	Allow time to clean train units																	●			
	Put newly cleaned train units into service																	●			
Winter Preparedness	Get train unit to facility within time limit																	●			
Exposure of Commercials	Expose commercials in certain regions							●									●				
Surveillance Video Requests	Get train unit to workshop within time limit																	●			
Surface Foil Application	Get train unit to facility within time limit																	●			
Passenger Counting Equip.	Assign specific train unit to specific train lines							●									●				
Train Control System Equip.	Assign specific train unit to specific train lines							●									●				
Operating Costs	Minimise energy costs		●							●											
	Minimise maintenance costs		●							●											
	Minimise infrastructure usage costs		●							●											
	Minimise train driver costs									●											
	Minimise depot driver costs									●											

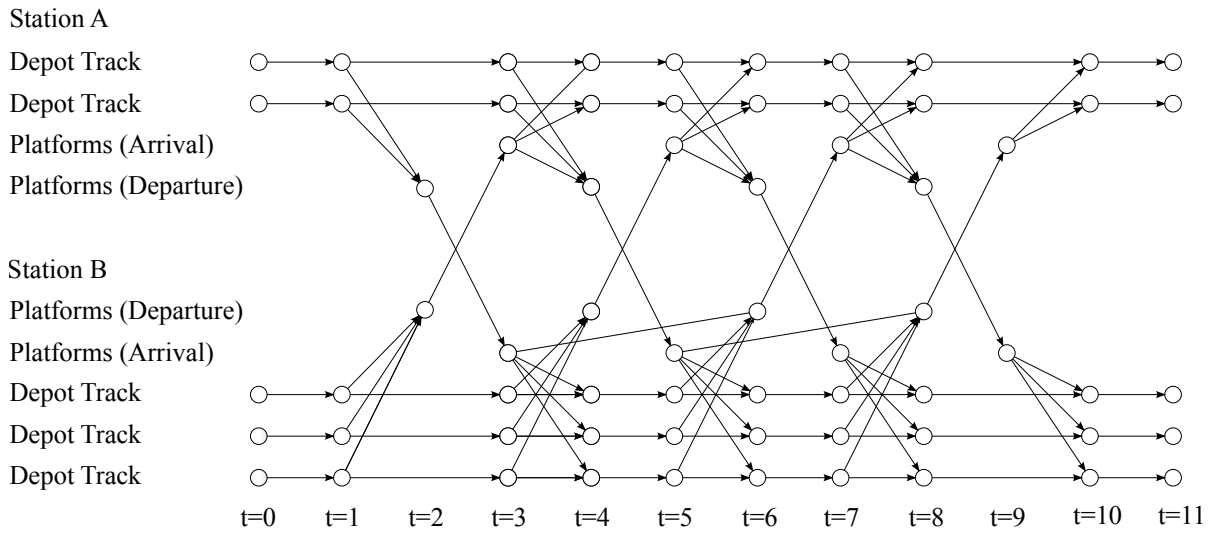

Figure 1: A schematic illustration of the space-time graph G. The time increases from left to right. Each vertex $v \in V$ represents an event, i. e. a space-time tuple $v = (q, t)$. Each arc $a \in A$ in the graph represents the possibility for a train unit to move in time and space (as a train service or train shunting) or in time only (as being parked at a depot track, or in the process of being turned around at a platform track). For example, the arc departing from station A at $t = 2$ arriving at station B at $t = 3$ is a train service. This arc has two corresponding vertices, the first one being the departure from the platforms at station A at $t = 2$, the second one the arrival at station B at $t = 3$. Note that the flow arcs going from the station source and to the station sink are omitted in this drawing. Also omitted are arcs representing overnight parking at the platform tracks.

In the graph, all platform tracks of a given station are treated as a whole. However, to be able to model how one train unit may turn around from one train service to the next at the platform tracks, vertices representing arrivals to the platform tracks are separate from those representing departures, even though their point in time may be the same.

In the model, non-revenue train services for the positioning of train units are static in the sense that they are not decided when solving the model, but read from the timetable data as given input.

The space-time graph has three different aspects which are used in the integrated rolling stock planning model. These aspects are:

1. **The topology of the graph**, that is, which arcs are connected to which other arcs by their common vertices. This is described in Section 3.1.
2. **The arc attributes**, describing features related to train services, train shuntings, depot tracks, etc. This is described in Section 3.2;
3. **The vertex attributes**, describing features related to depot drivers on duty, number of shuntings assigned to each arrival and departure, etc. This is described in Section 3.3;

As mentioned earlier, the space-time graph is used as the data model for a resource constrained shortest path algorithm with side constraints. In order for this to work, the space-time graph exhibits two features to the path finding algorithm for each arc. These features are:

1. **Arc feasibility:** *Is it feasible for a given train unit to traverse the arc?* This is determined by factors like how many train units already traverse the arc compared to the capacity of the arc, etc.;

2. **Resources consumption:** *What is the resources consumption of traversing the arc with a given train unit?* In the model, two resources are defined: The first one is the net value of traversing the arc, measured as the benefit of doing so minus the cost and penalties incurred. The second resource is the service distance travelled, this resource is constrained since a train unit may e. g. only travel a certain distance after which it has to undergo maintenance.

The path finding algorithm (see Section 5) is relaying feasibility and resource consumption requests to the space-time graph through these two features. The features are exhibited based on the values of the arc and vertex attributes and the graph topology as described in the next sections. The attributes are described in the same order as they appear in Table 4.

3.1 Topology

The topology of the space-time graph reflects the allowed movements of train units as stated by the current business rules at DSB S-tog. For example, there are no arcs connecting a depot track with another depot track, since a current business rule prohibits depot internal shunting. The graph also incorporates turnaround times between train services.

3.2 Arc Attributes

This section gives an overview of the different arc attributes and how they are utilised to implement the different requirements.

The maximum length that train units assigned to an arc may utilise is governed by the *track length* attribute, set according to the minimum platform length on the stations visited by train service, or the actual track length for platform, side and depot tracks.

The *service distance* attribute represents the real-world service distance travelled in physical space for that arc. This attribute is used to calculate the distance related cost for energy, maintenance and infrastructure usage, etc.

The *assigned train units limit* attribute is set to two for arcs representing train services, since at DSB S-tog, no more than two train units may be coupled together when running on the main line tracks. For other arc types the train unit capacity is infinite.

The main arc attribute in the space-time graph is the *assigned train composition* attribute, indicating which train units are assigned to the activity represented by the arc, and in what individual order. This attribute is used to check for train unit order feasibility (see Section 4.2) and overall resources consumption in comparison with other arc attributes.

The attribute for *transition positions* keeps track of at which positions couplings and decouplings may take place for the train composition assigned to the arc in question (see Section 4.2).

The attribute for *depot drivers required* keeps track of whether a depot driver is required for the activity represented by the arc. Only train shunting arcs to and from depot tracks feature this attribute, shunting to and from side tracks is performed by the train drivers. Also on the first and last train services in the weekend, the train drivers themselves perform the actual shunting to and from the depot.

The *train line* attribute keeps track of which train line (as defined in the timetable) a given train service belongs to in order to make sure that platform lengths are adhered to, commercials are exposed and that the technical equipment of the train is in accordance with the requirements for that train line.

The *additional benefit* attribute quantifies the benefit of assigning (adding) another train unit of a given type to the arc in the space-time graph. The *subtractional benefit* is equivalently the

benefit of removing (subtracting) a train unit from the assigned train composition of that arc. As mentioned in Section 2.2, the additional net value is calculated as the additional benefit minus the additional penalties awarded and additional costs incurred. As such, the additional benefit attribute is the only positive driver of the integrated rolling stock planning model. The attribute represents the economic value of providing seats to passengers that demand them. The value of the additional benefit is calculated according to stated preference time penalties for having no seat, the specific value of time for commuters and the duration of an average travel [33]. For the case of DSB S-tog, this yields a value of 0.44 DKK/min. This is roughly equivalent to the actual specific revenue gained by DSB S-tog by means of ticket sales for conveying an average passenger. Note that only arcs representing revenue train services can provide a benefit since they are the only type of arc that represent the conveyance of passengers.

The *additional cost* attribute is a resources consumption attribute quantifying the factual total cost for the activity represented by the arc, as incurred by the assigning (adding) another train unit of a given type, calculated as the sum of the cost for train or depot drivers, the cost for technical maintenance, the energy cost and the cost for using the railway infrastructure. The *subtractional cost* attribute is analogously quantifying the cost saved when removing (subtracting) a given train unit from the assigned train composition of that arc.

The *additional penalties* attribute quantifies the penalties awarded and saved by assigning (adding) a given train unit to the arc. The *subtractional penalties* attribute quantifies the penalties awarded and saved by removing (subtracting) a given train unit from the train composition assigned to the arc. Penalties are awarded for aspects of the plan that are not desired, e. g. uncovered train services or train shuntings from depot tracks to platform tracks where the train unit in question has to pass the main line tracks under way.

3.3 Vertex Attributes

All of the vertex attributes relate to the process of train shuntings to and from depot and side tracks. The limit on number of train shuntings quantifies the DSB S-tog business rule of an upper limit of one train shunting per train service departure or arrival. Vertices not being part of a train service have no limit on the number of train shuntings being performed.

The *assigned number of train shuntings* attribute keeps track of how many train shuntings are actually performed by assigned train units passing through the vertex.

The *date and time* attribute is used to map the time interval $p \in P$ of the arc being used for train shunting to and from depot tracks with the time intervals of the depot drivers on duty. If a train shunting arc is used, the date and time attributes on its vertices are matched with the supply in time of depot drivers for that particular depot station. If the demand from the current arc and the demand from other train shunting arcs being used at intersecting times may be supplied by the depot drivers on duty, the arc is feasible. If not, the arc is infeasible.

The *depot drivers on duty* attribute keeps track of the duties of the depot drivers by time and station. This is an aggregated attribute in the sense that the depot drivers on duty are shared between all the vertices on a given depot station for the time frame of that duty.

4 Train Unit Data Model

The second component of the proposed integrated rolling stock planning model is the data model for the train units. Its different attributes are described in the following.

The *allowed on lines* attribute keeps track of which lines (as defined in the timetable) the train unit in question may be be assigned to. This is in order to adhere to restrictions on installed

equipment and the exposure of commercials.

The *start/finish time-and-space* attribute is used to set the start and finish points in space and time for each train unit $u \in U$. This way the model can handle where train units are at the beginning and end of the plan period in order to e. g. send train units to maintenance at a given time (and by way of the service distance attribute also within the service distance limit). Furthermore, the balance of train units at each depot can be controlled using this attribute.

The *service distance limit* attribute keeps track of the service distance that the train unit in question can perform before it has to go into maintenance or have consumables refilled.

Two important features of the train unit data model are the train unit trajectories and the train unit order. These features are described in Sections 4.1 and 4.2.

4.1 Train Unit Trajectories

A train unit trajectory $j \in J$ keeps track of which activities the train unit in question is assigned to. As such, train unit trajectories are paths (i. e. ordered collections of arcs, $j = a_1, a_2, ..., a_{|j|}$) from the space-time graph representing the movement of their respective train units in space and time or time only. See Figure 2 on the next page for an illustration of train unit trajectories in the space-time graph. Figure 2 shows an example of how four train units move through time and space along their train unit trajectories. If two train unit trajectories use the same arc, this means that the two train units are coupled together as a *train composition*. Depending on the type of arc, they may either be coupled together and running as a train service, or be coupled and parked on at depot track. The order by which the train units are coupled is maintained by the space-time graph. Note that in the example the red and blue train units exchange places in the execution of the plan, however the balance of train units on the individual depot tracks remains constant.

The space-time graph is used to find new candidate trajectories for train units to improve the plan. Since the graph already contains information on the existing trajectories assigned, only candidate trajectories feasible in conjunction with the existing ones can be found.

New candidate trajectories are found using a specially constructed resource constrained shortest path algorithm with side constraints. This algorithm is described in Section 5.

4.2 Train Unit Order

The train unit data model also keeps track of the order of the individual train units relative to each other. In the following, the logic for coupling and decoupling train compositions in relation to the order of the train units will be explained. This logic is specific to the business and train control system rules currently in effect at DSB S-tog.

The simple explanation is this: *At the platform, train units are being coupled and decoupled in the direction that is facing the depot. At the depot, train units are being coupled and decoupled in the direction facing the platform.* This is illustrated in Figure 3.

A *train composition* is the ordered sequence of one or more train units coupled together. At DSB S-tog, train compositions consisting of one or two train units may be assigned to train services. Train compositions of more than two train units may be formed when parking train units at a depot (and thereby coupling them).

Current business and train system control rules at DSB S-tog state that when a decoupling takes place at the platform, the train composition moving away must move to the depot. It may not be assigned to a train service. Furthermore, when a coupling is to take place at a platform, the train composition moving in to couple must come from the depot. It may not come from the main line, i. e. from a train service.

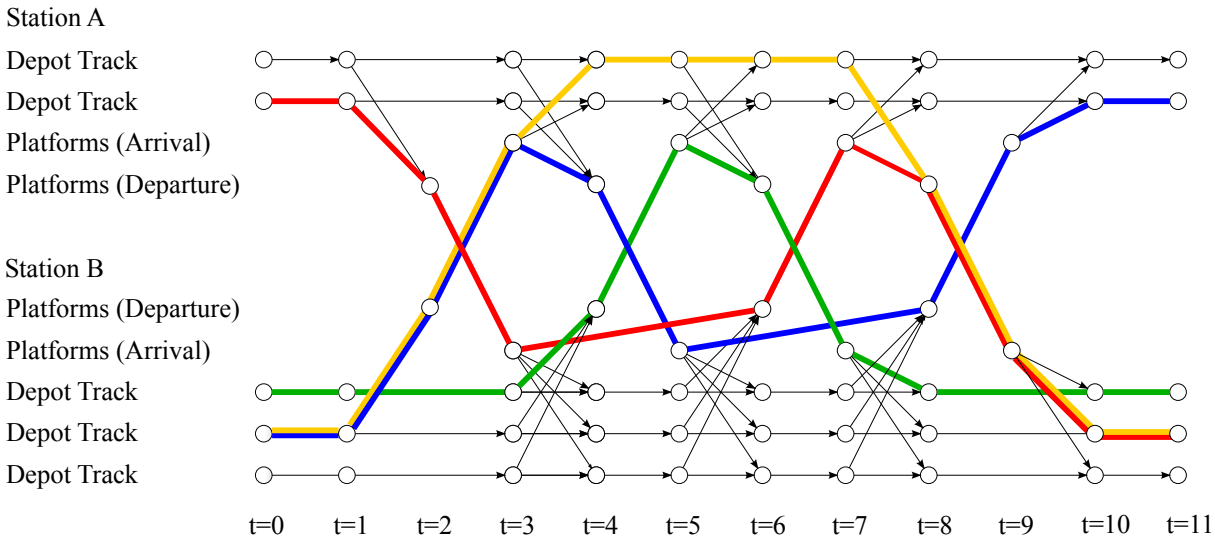

Figure 2: An example of train unit trajectories in the space-time graph, represented in colours red, green, blue and yellow. In this example, the red train unit starts at station A on depot track A2 at time $t = 0$. At $t = 1$ the red train unit is shunted to the platform, from which it departs as a train service at $t = 2$, arriving at station B at $t = 3$. At station B no arc exists to connect the arrival with the departure at $t = 4$. This is because the time difference between arrival and subsequent departure is less than the minimum turnaround time. For this reason, the red train unit waits at the platform from which it departs at $t = 6$, arriving at station A at $t = 7$. Station A has a shorter turnaround time and the red train unit may thus depart again at $t = 8$. Prior to departure it is coupled with the yellow train unit that is being shunted in from depot track A1 at $t = 7$. See Figure 3 for examples of how train units are parked at different times.

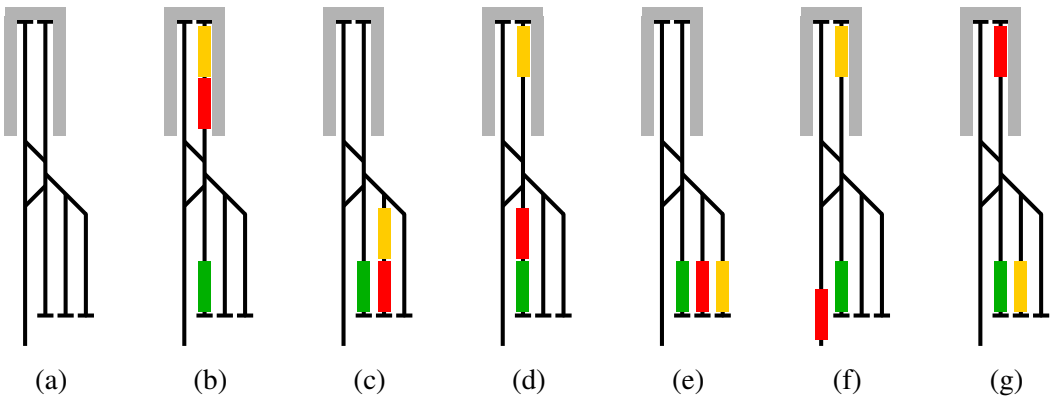

Figure 3: Examples of situations at Farum Station, equivalent to Station B in Figure 2. (a) shows the station layout: To the North there are two platform tracks, to the South there is the main line track (left) and three depot tracks (right). (a) also shows the situation at $t = 5$ for Station B in Figure 2. (b) shows the situation at $t = 9$ with two train units parked at the platform, the yellow one to the North of the red one. (c) shows the situation at $t = 10$, in which both train units have been shunted to the middle depot track. Business and train control system rules state that feasible transitions are (b) to (c), (b) to (d) and vice versa. The following transitions (and vice versa) are infeasible: (b) to (e) would require two shunting movements; (b) to (f) since the moving red train unit is on the main line track; (b) to (g) is infeasible because the red train unit obstructs the movement of the yellow one.

The term *platform train composition* is used to denote the train composition in the operation that is facing the platform. Similarly, the *depot train composition* is the train composition that is facing the depot.

The term *relative position* denotes how an object (platform track, depot track, train composition, train unit) is oriented relative to another. For DSB S-tog, the relative position can be either *North* or *South*. For example, the relative position of a depot track to a platform track at its corresponding station may be South, meaning that to reach the depot track from the platform track, train units must move towards the South. Equivalently, a train unit may have the relative position South of another train unit. At the same time, this also means that the other train unit has the opposite relative position, i. e. North, of the first one.

The individual train units in a train composition have a relative position to each other. With the proposed definition, the relative position of train units in compositions is conserved in all feasible coupling and decoupling operations.

Picture the **coupling** of two train compositions at a **platform track**, one being a platform train composition, the other being a depot train composition. The situation before the depot train composition is shunted in from the depot is depicted in Figure 3d, the situation after coupling in Figure 3b. After coupling, the original depot train composition will have the relative position to the original platform train composition (in the new train composition) equal to the relative position of the depot to the platform (in other words, equal to the relative position of where the train shunting movement **started**). If, like in the example in Figure 3, the relative position of a depot track is to the South of a platform track, then the original depot train composition being shunted in to the platform on this station will have the relative position South to the original platform train composition in the new train composition. This is the result of the transition between the situations from Figure 3d to Figure 3b.

In the case of a **coupling** taking place at a **depot track**, the relative position of the platform train composition (which is the one undergoing movement in the operation) will also be the same as the relative position of the place from which the shunting movement **started**, in this case the platform track. If, like in the example above, the relative position of a depot track is to the South of a platform track, then the relative position of the platform track is to the North of the depot track. The original platform train composition after coupling will then have the same relative position to the original depot train composition as the platform track has to the depot track, which in this case is North. This is the result of the transition between the situations from Figure 3b to Figure 3d.

In the case of a **decoupling** taking place at a **platform track**, only the train units at the same relative position to the others in the original train composition as the relative position of the depot track to the platform track may be decoupled to form the new depot train composition to be shunted into the depot. This is equivalent to the relative position of where the shunting movement **ended**.

The ordering of train units in a train composition can be found by sorting the individual train units according to their relative positions.

5 Path Finding Algorithm

The third component of the integrated rolling stock planning model is the path finding algorithm, used to find new candidate train unit trajectories. The path finding algorithm is operating on the space-time graph described in Section 3 to find a path between a start vertex and a finish vertex. In the context of the integrated rolling stock planning model, the goal of the algorithm is to find the path through the space-time graph between these two vertices, having the largest

additional net value (i. e. for which it is most advantageous to add a new train unit trajectory).

For convenience, the weights on the arcs in the space-time graph are set as the **negated** additional net value. This makes the algorithm work as a shortest path algorithm. The path finding algorithm concept is thus that of a *single-source shortest path for a directed, acyclic graph* [17], with resource and side constraints. It is implemented as a label setting algorithm and is using dominance to keep the set of potential paths to a minimum [26]. The algorithm traverses the space-time graph and creates potential paths through the graph by setting and processing labels on the vertices it traverses. Upon termination, the shortest path may be found by backtracking the processed labels.

The path finding algorithm consist of two parts: The outer part of the algorithm which is traversing the space-time graph (described in Section 5.1, shown as Algorithm 1) and the inner part which is processing the labels (described in Section 5.2 shown as Algorithm 2). The inner part is called in each iteration of the outer part.

When finding a new path, that is, a new train unit trajectory, the algorithm must ensure that the found train unit trajectory is feasible. This is ensured in three ways:

1. **Resource constraints:** The total resource consumption of potential paths is checked in each iteration so that no resources are exhausted (meaning that the service distance limit of the individual train unit is never exceeded). This check is handled in the inner, label processing part, Algorithm 2;
2. **Side constraints:** The check for train unit order feasibility in decoupling operations is built into the path finding algorithm itself. So is the check for flexible space distribution for bicycles etc. These two checks constitute the side constraints in the algorithm. Both these checks are handled in the inner, label processing part, Algorithm 2;
3. **Space-time graph constraints:** The handling of all other practical, railway specific requirements is relayed to the space-time graph as previously described in Section 3. The relaying is performed when traversing the graph in the outer part, Algorithm 1 (see Line 3).

In the path finding algorithm, each vertex from the space-time graph has associated to it a number of labels. The labels are used to mark potential paths through the space-time graph as the algorithm progresses. Each label refers to one arc on the potential path. Furthermore, each label belongs to a vertex (the to-vertex on the arc to which the label refers), but there may be many labels to the same vertex, since different, potential paths may pass through the vertex and since more arcs may be connected to it. Each label carries with it the following information:

1. **The arc in the space-time graph** from which a part of the potential path represented by the label passes through. This information is used to put together the shortest path when the finish vertex has been reached;
2. **The previous label on the potential path** (the predecessor). This information is used to backtrack the labels to find the shortest path once the finish vertex has been reached;
3. **The resources consumed** in order to reach the vertex of the label, starting at the start vertex. This information is used to ensure no resources are exhausted. It is also used for dominance, i. e. to keep the number of potential paths small;
4. **The ordered train composition** of the arc, i. e. which train units are assigned to the arc and in what individual order. This information is used in the side constraints of the algorithm to check for train unit order feasibility and flexible space distribution for bicycles etc.

The inner workings of the two parts of the path finding algorithm are described in the following Sections 5.1 and 5.2.

5.1 Space-Time Graph Traversing

The first, outer part of the path finding algorithm is the space-time graph traversing part, shown as Algorithm 1.

Algorithm 1: Resource constrained shortest path label setting algorithm with side constraints for a space-time graph with arc resource consumption data and vertices sorted in topological order. The resource and side constraints are checked in Algorithm 2.

```
Input: The from vertex and the to vertex in the graph.
Output: The resource constrained shortest path between the given vertices, feasible with
        regard to all practical, railway oriented requirements.
1  foreach (toVertex in vertices) do
2      foreach (thisArc in toVertex.getIncomingArcs()) do
3          if (thisArc.isFeasible()) then
4              foreach (fromVertexLabel in thisArc.getFromVertex().getLabels()) do
5                  toVertex.add(getLabel(thisArc, fromVertexLabel));        /* See alg. 2 */
6              end
7          end
8      end
9      toVertex.getLabels().removeDominated() ;
10     if (toVertex.isFinishVertex() && toVertex.isReached()) then
11         processShortestPath(toVertex);                /* Backtrack to set shortest path */
12         return shortestPathArcs;
13     end
14 end
15 throw NoPathFoundException;
```

Algorithm 1 starts its main for loop with the next vertex in topological order (Line 1). For each of the incoming arcs to this vertex (Line 2), a check is performed to see if the incoming arc is feasible (Line 3). This check is relayed to the space-time graph via the arc. If the arc it is not feasible, no processing of labels is occurring, and the algorithm is not proceeding further along that potential path. If the space-time graph responds that the arc is feasible, a loop over each of the labels of the from vertex of the incoming arc is started (Line 4).

Inside this loop, the inner, label processing part, is called, see Algorithm 2. Labels are only created if the side constraints are not violated. If not violated, the new label is added to the current vertex, otherwise no label is added.

Algorithm 1 then continues with the next incoming arc of the current to vertex. When all incoming arcs have been processed, the labels of the to vertex that are dominated are removed (Line 9).

If the current vertex is equal to the finish vertex, and this vertex has been reached, the shortest path is found by backtracking the labels starting with the label at the finish vertex having the least resource consumption for the path finding resource (in this case the negated additional net value).

5.2 Label Processing

The second, inner part of the path finding algorithm is the label processing part, shown as Algorithm 2.

Algorithm 2: The label processing part of the path finding algorithm. This part of the algorithm makes sure that no resources are exhausted and that the train unit order and flexible space distribution for bicycles etc. is feasible.

Input: The current arc; The current label of the from vertex of the current arc; The candidate train composition (i. e. information on which train unit(s) the algorithm is currently finding a new trajectory for); The resources available for the path.

Output: The algorithm creates a new label if constraints are not violated, no label is created if they are.

```
thisFromVertexTrainComposition ← getTrainCompositionAt(thisFromVertex,
fromVertexLabel);
thisArcOldTrainComposition ← thisArc.getTrainComposition();
thisArcNewTrainComposition ←
thisFromEventTrainComposition.getPreserveOrderSumOf( candidateTrainComposition,
thisArcOldTrainComposition);
if (thisArc instanceof TransitionPositioner) then
    transitionPositioner ← (TransitionPositioner) thisArc;
    if (!thisFromEventTrainComposition.canDecouple(thisArcNewTrainComposition,
    transitionPositioner.getDecouplingPosition()) ||
    !thisArcNewTrainComposition.hasLegalFlexibleSpaceDistribution()) then
        return null;
    end
end
thisConsumedResources ← resourcesPool.get();
thisConsumedResources.setSum(thisArc.getConsumption(),
fromVertexLabel.getConsumedResources());
if (thisConsumedResources.exhaust(availableResources)) then
    return null;
end
thisLabel ← labelPool.get();
thisLabel.set(thisArc, fromVertexLabel, thisConsumedResources,
thisArcNewTrainComposition);
return thisLabel;
```

The first part of Algorithm 2 relates to finding the ordered train composition of the extension of the current potential path. This part follows the hereby proposed principle of *(train) unit order flow conservation*, in which not only the inflow of train units to a vertex is conserved in the outflow (like in arc based multi-commodity flow models, as mentioned in Section 1.2), the train unit order is also conserved.

As such, line 1 calculates the train composition at the from vertex by coupling the train composition already assigned to the incoming arcs in the graph with the train composition of the predecessor label. This coupling is conducted using the topological information, the transition positions, as described in Section 3.2. As such the inflow order of train units to the from vertex is calculated for this particular potential path.

Next, line 2 sets a variable for the further reference to the train composition on the current arc consisting of train units already assigned to the arc in the space-time graph.

Line 3 then uses the order found in the inflow to the vertex in question to add the candidate train composition to the train composition already assigned to the current arc in the space-time graph. As such, the outflow from the vertex in question is determined with the correct order on the current arc.

The algorithm then proceeds to reject cases where the inflow order and the outflow order are not compatible: Line 4 tests if the current arc is of a type where transitions occur. If so the current arc is type cast (line 5) to be able to query it for the decoupling transition position (in line 6). If decoupling cannot take place while preserving the inflow order, or, if the flexible space distribution of the ensuing train composition is not feasible, the algorithm terminates by not returning a new label (line 7).

Next (in line 10) a new resources object is retrieved from the pool [41]. This resources object is set to the consumed resources being the sum of the consumption of the resources on the incoming arc and the previously consumed resources of the current from label.

If the consumed resources exhaust the available ones, the algorithm also terminates without returning a label. If not, a label is retrieved from the label pool, set with relevant data and returned.

Note that the logic to determine feasibility of the order of the train units in the composition applies to decoupling only. This is because it is only in the process of decoupling that a train unit may obstruct the movements of another. A coupling process will always conserve the relative position of the individual train units. A feasible decoupling will also conserve the relative position of the train units involved, however, an infeasible coupling, if it could occur, would not.

6 Heuristic Framework

The fourth and last component of the integrated rolling stock planning model is the heuristic framework used to accept or reject candidate trajectories found with the path finding algorithm described in Section 5. The heuristic framework used is that of hill climbing [30].

The overall concept of the heuristic framework in the integrated rolling stock planning model is to remove a number of train unit trajectories from an existing rolling stock plan and then, one by one, to create a new trajectory and insert it into the plan.

By generating and inserting new train unit trajectories into the plan one at a time, it is assured that each new train unit trajectory is feasible in conjunction with the existing ones in the plan.

The objective function of the heuristic is described in Section 6.1, the inner workings of the hill climbing heuristic itself is described in Section 6.2, and Section 6.3 describes the flow of changes to the objective value in one iteration of the heuristic.

6.1 Objective Function

The objective function of the proposed heuristic is the *net value* of a rolling stock plan. The net value is defined as the benefit a rolling stock plan provides minus the costs for providing it and the penalties awarded for undesirable features. As such, the benefit and the costs plus penalties are competing terms in the objective function and a rolling stock plan may be improved by maximising the benefits and/or minimising the costs plus penalties.

The objective value may be calculated for the entire rolling stock plan, however, for performance reasons the net value of removing individual trajectories (the subtractional net value) and adding new ones (the additional net value) is used, since this involves fewer calculations.

The *iteration net value increase* z_Δ is thus calculated as the sum over all selected train units U^* of the subtractional net value $z^-(j_u^-)$ of each train unit trajectory j_u^- removed plus the sum over all selected train units U^* of the additional net value $z^+(j_u^+)$ of each train unit trajectory j_u^+ inserted (1). Note that the subtractional net value is calculated before removing each original train unit trajectory and the additional net value is calculated before inserting each new train unit trajectory.

$$z_\Delta = \sum_{u \in U^*} z^-(j_u^-) + \sum_{u \in U^*} z^+(j_u^+) \tag{1}$$

The calculation of z_Δ described above yields the same result as the difference in total net value for the entire rolling stock plan before and after an iteration, but involves fewer calculations.

The subtractional net value is calculated in (2) as the costs saved $c^-(a, u)$ minus the benefits lost $b^-(a, u)$ minus the penalties awarded $p^+(a, u)$ plus the penalties saved $p^-(a, u)$ as a result of subtracting train unit u from arc a .

$$z^-(j_u^-) = \sum_{a \in A_j} c^-(a, u) - b^-(a, u) - p^+(a, u) + p^-(a, u) \tag{2}$$

The additional net value is calculated in (3) as the benefits gained $b^+(a, u)$ minus the costs added $c^+(a, u)$ minus the penalties awarded $p^+(a, u)$ plus the penalties saved $p^-(a, u)$ as a result of adding train unit u to arc a .

$$z^+(j_u^+) = \sum_{a \in A_j} b^+(a, u) - c^+(a, u) - p^+(a, u) + p^-(a, u) \tag{3}$$

Benefits represent the fulfilment of an unfulfilled seat demand for each individual arc and depend upon seat demand by time interval $p \in P$ and train unit type $i \in I$. Benefits are gained if train units are added to arcs that represent revenue train services demanding additional seats. Benefits are lost if train units are subtracted from the same arcs, if this results in more unfulfilled seat demand.

Costs are incurred for energy, maintenance, infrastructure usage and personnel (train drivers and depot drivers) and depend upon service distance, arc type (revenue train services, train shunting operations etc.), time interval $p \in P$ and train unit type $i \in I$. Costs for energy and maintenance are added if train units are added to arcs that represent train movements. Energy and maintenance costs are saved if train units are subtracted from the same arcs. For the case of infrastructure use and personnel, costs are only added for the first train unit added to the train movement arc in question. These costs are also only saved when the train unit being subtracted is the last train unit assigned to the train movement arc in question.

Penalties are awarded for unwanted features of the plan and depend upon time interval $p \in P$ and occurrence. A penalty for not having a revenue train service covered is awarded when the

last train unit assigned to the arc in question is subtracted. The penalty is saved when the first train unit is added to the same arc again. A penalty is also awarded for unwanted shunting operations. Here the concept is the opposite: This penalty is awarded when the first train unit is added to the train shunting operation arc, and saved when the last one is subtracted.

Note that penalties can be both awarded and saved by both the subtractional as well as the additional operation, this is the reason for both $p^+(a,u)$ and $p^-(a,u)$ being present in both (2) and (3).

6.2 Hill Climbing Heuristic

The concept of removing one or more train unit trajectories from the plan and reinserting them one at a time, operates within a hill climbing heuristic framework in which all of the inserted new trajectories are either accepted (and kept in the plan) or rejected (and removed from the plan and old trajectories re-inserted, yielding a plan identical to before any modifications were performed).

Normally in heuristics, changes are evaluated **before** being inserted into the solution. However the modification scheme described here is necessary to ensure feasibility: One can not generate a second candidate trajectory without inserting the first one into the plan, because otherwise the plan cannot determine if the second candidate trajectory conflicts with the first one.

The selection of train unit trajectories to remove from the plan is conducted at random.

The heuristic component continues until a given stopping criterion is met. For the calculations given in Table 5 on page 27 this was given as a negative value of z_Δ for the past 5 minutes, equivalent to the convergence curve (see Figure 6 on page 29) flattening out.

6.3 Objective Value Flow

Figure 4 shows a Sankey diagram of the flows for benefits and costs (including penalties) in the proposed heuristic. The figure is read from left to right showing the flow of costs (bottom) and benefits (top) through one iteration of the heuristic. The term *actual benefit* is used to describe the benefit that is in the current plan. The term *latent benefit* is used to denote a benefit that is not in the current plan, but one that could be in the next iteration through the modifications of the heuristic. The *total benefit potential* is thus the sum of these two types of benefit. On the cost side the same terms apply. The general idea of the heuristic is to maximise the benefits by turning latent benefits into actual ones and to minimise the costs (including penalties) by turning actual costs into latent ones, thus maximising the total net value of the plan.

7 Experimental Results

The integrated rolling stock planning model proposed in this paper has been tested on a number of data instances. The purpose of the experiments has been to benchmark the performance of the heuristic with plans produced manually. The conditions were chosen so as to make the benchmarking on as equal conditions as possible.

The integrated rolling stock planning model presented here has been implemented in the programming language Java 1.8 with approx. 15,000 lines of code. Apart from the library Joda-Time 2.8.2, only standard libraries have been used.

The tests were performed on a Dell PowerEdge T610 equipped with 16 Intel Xeon E5620 CPUs at 2.40 GHz and 16 GB RAM running Ubuntu Linux 14.04 LTS. Parallel processing was

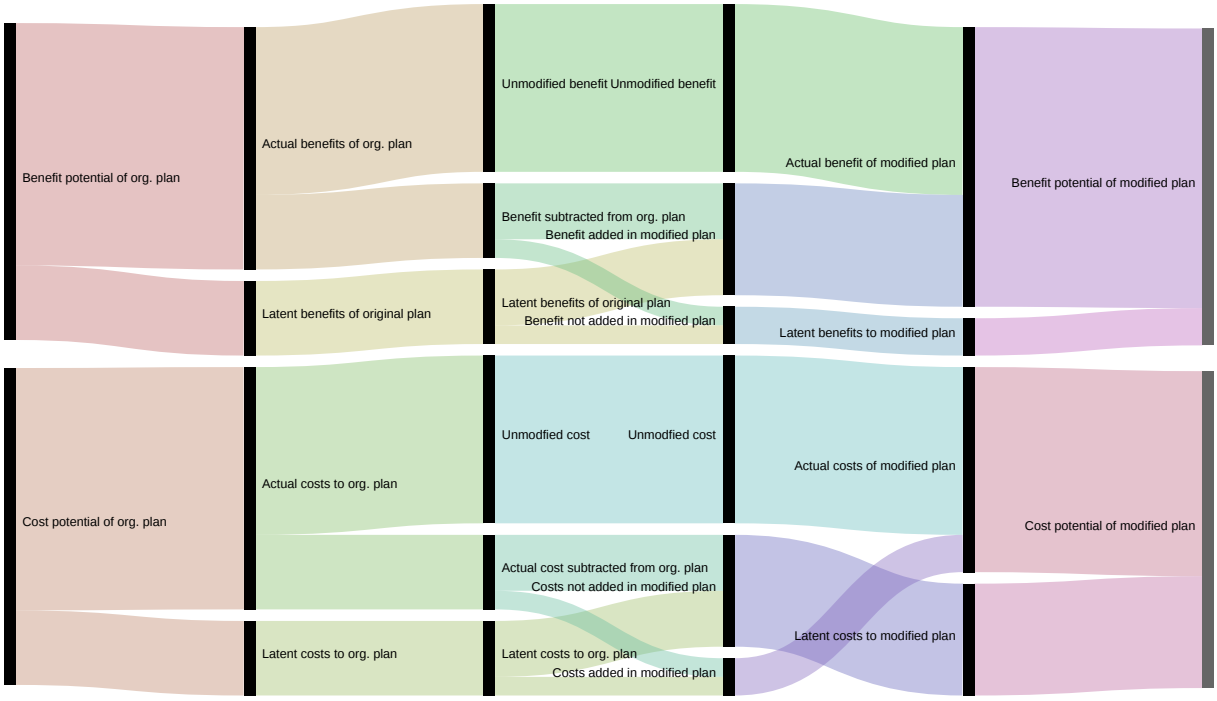

Figure 4: Sankey diagram of the flows for benefits and costs (including penalties) in the proposed heuristic.

used in which each individual test was run in its own thread on one CPU, parallel to other tests.

7.1 Data Instances

The proposed heuristic has been tested on 15 different rolling stock plan data instances as shown in Table 5. All the data instances represent long-term circulation plans (as opposed to short-term train unit dispatching plans). Each individual data instance represents a particular date, e. g. 2012-10-19 and a particular weekday, e. g. Friday.

Most of the data in the instances are actual, real-world data. This includes infrastructure data, timetable data, passenger demand data and data on personnel on duty. How the individual parts of the real world data vary between instances is described in the following.

The same timetable is in effect from Monday to Friday, but a different one is used on Saturdays, and again a different one on Sundays. Night train services operate on mornings after Fridays and Saturdays. The timetable also differs between the years 2012 and 2014.

The depot driver duties differ by weekday. The reason for this is that the start up procedures on Monday morning are different from the ones on Tuesday, since there is a change of timetable between Sunday and Monday, but not between Monday and Tuesday.

In the data instances, the passenger demand is represented by running the DSB S-tog passenger prognosis model with the actual, measured passenger data for those days. This is possible because the data instances are in the past. In a real planning situation prognosis passenger demand would be used.

The instance of 2012-10-19 is special since it occurs during the autumn holiday and also represents an extraordinary plan with infrastructure maintenance works on a parallel, long distance line. For this reason this plan provides extra capacity on the one of the lines. (It has turned out that, in hindsight, this plan provides far too much seating capacity, which can be seen by how much the heuristic may increase its net value by removing this excess capacity again.) The other dates represent normal plans with no extraordinary features.

All of the data instances are complete rolling stock plans. In the experiments, the infeasible train unit trajectories (if any) are removed prior to running the heuristic. The start and finish points in space and time of the original train unit trajectories are kept. As such, new trajectories constructed using the heuristic have the same origin and destination stations. This preserves the depot balance between the original and rolling stock plan when using the heuristic.

Since the current rolling stock planning procedures at DSB S-tog still involve some degree of manual work, data is not available for all aspects of manually produced rolling stock plans. Parts of the data instances not currently available include data on the depot movements of the train units and the grouping of anonymous train units so as to determine the number of train units in the plan.

Firstly, the only information currently available on the movement of train units is the assignment of train units to train services. As such, the provided train unit trajectories have gaps in them. In order for the integrated rolling stock planning model to work, this information has to be generated artificially by *retrofitting* each manually produced plan to the data model described in Sections 3 and 4.

In some cases, all train unit trajectories can be retrofitted. However, as shown in Table 5 on page 27 in the column *# of infeasible trajectories*, in most cases, the retrofitting process is not able to retrofit all gaps in the supplied train unit trajectories. This is because the manually produced plans may not respect all of the practical, railway specific requirements and thus can not be mapped onto the data model. (Recall that the data model is constructed so as not to allow violations of the requirements.) The infeasibilities of the manually produced plans may relate to depot track capacities, depot driver duties, depot driver time consumption and minimum turn-around times between consecutive train services.

Some of these infeasibilities arise in the automated, step-by-step planing system currently in use at DSB S-tog. However, it is also common for the planners to use a variety of "dirty tricks" to improve the economic attractiveness of a rolling stock plan manually, or to replace intolerable infeasibilities from the automated planning system with "less intolerable" infeasibilities. Strictly speaking, these tricks are violations of the railway specific requirements, but since they often save substantial costs they are accepted. The infeasibilities introduced by the planners are often used as a last resort and incorporate all the tacit knowledge of the planners. Therefore, it may indeed be very hard for any automated system to compete with plans having these mentioned infeasibilities in them.

Secondly, realistic data is also constructed for the grouping of anonymous train units in the plan. In the circulation planning phase of rolling stock planning, train units are anonymous since it may not be known which actual, physical train units may be available when the plan is to be commenced (some train units may e. g. be in unscheduled maintenance).

In the data instances used here, there is currently no information regarding whether two anonymous train units with non-overlapping train unit trajectories are actually intended to be supplied with rolling stock by assigning the same physical train unit in the train unit dispatching phase to both of them. This information is therefore generated artificially by coalescing train unit trajectories that do not overlap and that may be performed in sequence by connecting them at the intermediate depot.

This feature may make it harder for the integrated rolling stock planning model to find good solutions since it may have less train units (and thereby fewer degrees of freedom) to do so than the manual planners have. On the other hand, if the coalescing would not take place, there would be more anonymous train units than would fit in the actual depots, which would also mean missing retrofits and thus infeasibilities in the plan.

In the long-term circulation planning process of DSB S-tog, rolling stock plans are constructed for one week at a time, in which each day connects to the next. In this paper, however,

the scope is on each individual day, not consecutive days. For this reason it does not make sense to restrict the service distance each individual train unit may travel before it has to undergo scheduled maintenance or refilling of consumables, since these events occur at intervals far greater than that of a day. If arbitrary values for the service distance limit would be included, the bench marking results would not be comparable.

This feature makes it slightly easier for the integrated rolling stock planning model to find good solutions since it can make an individual train unit roll a bit longer than it may roll due to e. g. scheduled maintenance. Experiments have been conducted to show that the model also performs well with the service distance limits in place.

All of the data instances used for the experiments represent the circulation planning phase of the rolling stock planning process. For this reason the remaining railway specific requirements related to the short-term train unit dispatching phase are disregarded in the experiments.

These issues explained, the data used in the experiments represent a very close approximation to the real-life planning conditions.

7.2 Convergence Characteristics of the Proposed Heuristic

The main driver of the heuristic is the objective function, the net value. The net value is calculated as the benefits of providing seats to passengers that demand them minus the costs of doing so and minus penalties for unwanted features of the plan.

The penalties are estimated as real monetary values. A very high specific penalty has been set for those aspects definitely unwanted, e. g. uncovered train services. A moderate specific penalty has been set for those aspects that are just undesirable, e. g. train shuntings across the main line tracks. Simple tests have shown that these estimates make the model perform well.

Experiments have justified the current settings for those parameters that are in the objective function. Another parameter that governs the behaviour of the heuristic is the number of train unit trajectories to pick and remove, k.

More than 2,000 test runs have been performed to analyse the behaviour of the model when varying model parameters. Some results are shown in Figures 5 and 6. The observations from the test runs are given in the following and characterised in relation to *initial gain* (how well the model converges early in the process - more is better), *maximum gain* (how much gain does the model achieve before the convergence curve flattens out - more is better), and *variance* (how results from different test runs with the same parameters vary - less is better).

Low values of k yield high initial gain, little variance, but only moderate maximum gain. Higher values of k yield low initial gain, higher variance but better maximum gain. However, beyond a threshold value of k, the performance seems to deteriorate generally.

We believe this behaviour is related to two characteristics of the heuristic: Firstly, if $k = 1$, no swapping of resources between train units can occur. From $k = 2$ and onward, swapping can occur, but for higher values of k, the greediness of the algorithm takes over, in the sense that the formerly inserted trajectories use the resources at the expense of the latter trajectories inserted, this yielding a lower net value gain in total.

Based on the observations from the test runs, the best performance of the proposed heuristic is achieved by selecting a moderate value for k, e. g. $k = 3$.

7.3 Obtained Results

As seen in Table 5, the experiments show that the heuristic used in the integrated rolling stock planning model is able to make all instances of typically infeasible rolling stock plans feasible. In most cases, feasibility is reached with a processing time less than 1 minute. In the worst case

Table 5: Overview of data instances and experimental results for the improvement of rolling stock plans for $k = 3$. The results given are based on $n = 10$ individual test runs of each data instance. *Mean net value gain* is the mean of the relative difference of the net value between the original plan (with infeasibilities in it, if any) and the modified one. Stopping criterion of the individual test run was feasibility and no net value gain obtained in the last 5 minutes of the calculation.

Instance meta data		Train units			Original characteristics						Modified characteristics						Performance characteristics						
Date	Weekday	# of type 1	# of type $\frac{1}{2}$	Total	Costs [kDKK]	Penalties [kDKK]	Benefits [kDKK]	Net value [kDKK]	# of infeasible trajectories	# of uncov. train service arcs	Mean costs [kDKK]	Mean penalties [kDKK]	Mean benefits [kDKK]	Mean net value [kDKK]	Max # of infeas. trajectories	Max # of uncov. train srv. arcs	Min net value gain [%]	Mean net value gain [%]	Max net value gain [%]	Min processing time [h:m]	Mean processing time [h:m]	Max processing time [h:m]	Max proc. time to feasibility [h:m]
2012-10-19	Fri	89	22	111	1,571	26	3,772	2,175	0	0	1,412	5	3,765	2,348	0	0	7.8	8.0	8.2	0:29	0:41	0:59	0:00
2014-03-31	Mon	90	27	117	1,538	28	4,445	2,879	1	0	1,494	9	4,424	2,922	0	0	1.2	1.5	1.6	0:21	0:41	1:00	0:00
2014-04-01	Tue	90	27	117	1,537	27	4,445	2,881	3	0	1,490	9	4,416	2,917	0	0	1.0	1.3	1.9	0:23	0:38	1:15	0:00
2014-04-02	Wed	90	27	117	1,537	31	4,445	2,877	3	0	1,489	9	4,417	2,919	0	0	0.8	1.5	1.7	0:24	0:44	1:15	0:02
2014-04-03	Thu	90	27	117	1,537	28	4,445	2,879	2	0	1,490	9	4,423	2,924	0	0	0.9	1.6	1.9	0:19	0:44	1:05	0:00
2014-04-04	Fri	90	28	118	1,544	30	4,492	2,917	3	0	1,486	10	4,471	2,975	0	0	1.3	2.0	2.4	0:31	0:40	0:49	0:04
2014-04-05	Sat	53	7	60	984	21	2,780	1,775	2	0	982	5	2,779	1,792	0	0	0.8	0.9	1.0	0:10	0:14	0:17	0:03
2014-04-06	Sun	52	11	63	980	19	2,166	1,167	0	0	955	2	2,162	1,205	0	0	3.1	3.3	3.7	0:11	0:17	0:32	0:00
2014-04-07	Mon	90	27	117	1,538	27	4,445	2,880	1	0	1,498	9	4,430	2,922	0	0	1.3	1.5	1.9	0:21	0:36	0:51	0:00
2014-04-08	Tue	90	27	117	1,538	27	4,445	2,880	3	0	1,490	8	4,416	2,918	0	0	0.8	1.3	1.7	0:25	0:48	1:15	0:06
2014-04-09	Wed	90	27	117	1,537	31	4,445	2,877	3	0	1,488	9	4,416	2,919	0	0	1.2	1.5	1.9	0:16	0:39	1:19	0:06
2014-04-10	Thu	90	27	117	1,538	27	4,445	2,880	1	0	1,490	9	4,421	2,922	0	0	1.2	1.4	2.0	0:22	0:41	1:13	0:00
2014-04-11	Fri	90	28	118	1,545	30	4,492	2,917	4	0	1,482	10	4,471	2,980	0	0	1.9	2.1	2.4	0:19	0:38	0:56	0:00
2014-04-12	Sat	53	7	60	983	21	2,780	1,775	2	0	983	5	2,780	1,792	0	0	0.9	1.0	1.0	0:09	0:12	0:16	0:00
2014-04-13	Sun	52	7	59	977	19	2,166	1,169	0	0	968	2	2,166	1,196	0	0	2.2	2.3	2.5	0:09	0:13	0:18	0:00

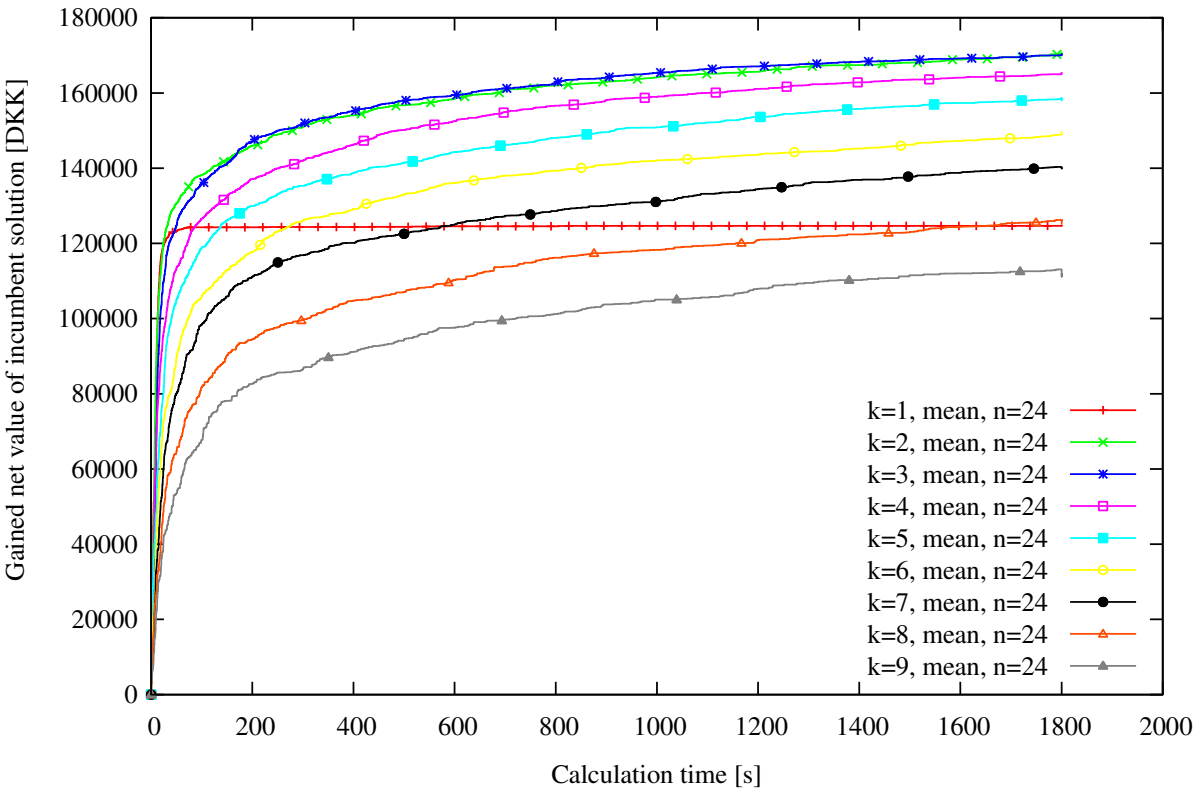

Figure 5: Convergence diagram for different values of k, i. e. the number of trajectories to remove and re-insert in one iteration. As may be seen, the best results are achieved with $k = 3$. Each line represents the mean of $n = 24$ test runs.

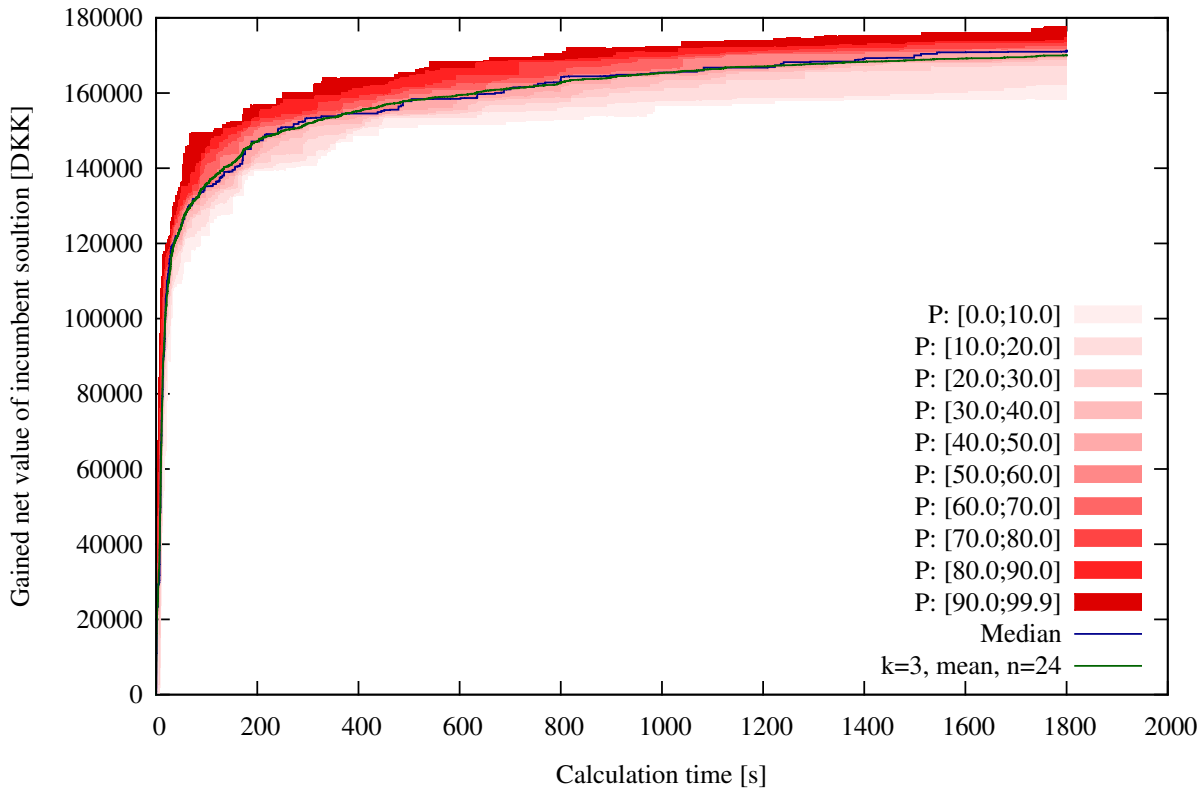

Figure 6: A convergence diagram showing the mean value and the median of the gain for 24 test runs with the given parameter settings. Also shown is the cumulative probability P in percent of getting a solution with a value less than a given gained net value.

of the 150 test runs reported in Table 5, feasibility is reached within 7 minutes of processing time.

Furthermore, the heuristic is able to improve the economic attractiveness of all tested rolling stock plan instances with an average economic gain of 2%. The highest net value gain is achieved for the instance 2012-10-19. As mentioned in section Section 7.1, this special instance represents an extraordinary plan providing (as it turns out) excessive, extra capacity on the one of the lines.

With a stopping criterion of no gain in 5 minutes the mean processing time is less than an hour for all instances. In no case is the processing time longer than 1 hour 20 minutes.

8 Discussion

In implementing the proposed, integrated model, we have demonstrated that it is possible to integrate into one model processes in rolling stock planning that are normally solved in a step by step manner. We have also shown that the proposed *ordered flow conservation* principle works in practice. The work presented here is a case study on the specific real-world data and constraints of DSB S-tog, however, the proposed principles may be applied in general to other types of problems.

Furthermore, we have observed that the proposed heuristic operating on the integrated rolling stock planning model has good convergence characteristics in that it will make typical infeasible plans feasible within minutes of computation time and with an additional gain in economic value of 2% on average.

These features make the integrated rolling stock planning model highly suited to simplify and improve present semi-automatic or manual rolling stock planning procedures at DSB S-tog.

The integrated rolling stock planning model proposed in this paper uses a shortest path algorithm to find new candidate train unit trajectories. This fact makes the algorithm highly greedy, a feature we believe makes the algorithm very suitable for fast "trimming and grooming" of an existing rolling stock plan with excess seating capacity in it. However, the same greediness property may make it less suited for constructing a rolling stock plan from scratch, since the former found shortest train unit trajectories will be good at the cost of latter found ones. Further research may devise methods to overcome this limitation.

Extensive experiments have been conducted to replace the heuristic proposed here with a metaheuristic using Simulated Annealing [1] with an exponential cooling scheme and reheating. However, none of these experiments produced better convergence. Future research into other metaheuristic frameworks may be conducted to improve convergence.

The model proposed is primarily intended for the tactical planning scope. However, with few modifications, the model may also be used for the short-term train unit dispatching phase of rolling stock planning. Schemes for reinstating cancelled train services are currently not handled, nor is the objective in a disruption recovery setting to recover with a minimum of changes to the non-disrupted plan. Future research should cover these aspects, if the model is to be used in the train unit dispatching context.

A further integration of train unit routing could be built into the model by also modelling the individual platform tracks instead of modelling them as a whole.

Moreover, a concept for creating relevant non-revenue train services should be considered in the future. The model proposed here only uses the non-revenue train services in the original plan (if appropriate), however these non-revenue train services may not fit the modified plan very well, since they are specifically created to fit the original plan.

Even though the practical, railway specific requirements that relate to the train unit dispatching phase have been implemented with the proposed, integrated rolling stock planning model, experiments with train unit dispatching plans have not been made yet. One major difference to the circulation plans used in the experiments is that the individual positions of the train units is given. This is because the individual, physical train units have a specific position at the time the plan is to be set in motion. In the long-term circulation planning phase, this position can be chosen. This difference somewhat limits the amount of candidate solutions for train unit dispatching.

Lastly, and perhaps most importantly, research should be conducted into the field of improving convergence by devising different schemes from which to choose train unit trajectories to remove as well as different schemes for inserting them into the plan.

Acknowledgements

The authors would like to thank for their contribution (in alphabetical order): Rasmus Dyhr Frederiksen, Rapidis; Julie Jespersen Groth, DSB; Jørgen T. Haahr, DTU; Jesper Hansen; Assistant Professor Richard Lusby, DTU; Professor Otto Anker Nielsen, DTU; Professor David Ryan, University of Auckland. The research leading to the results presented here has received funding from Innovation Fund Denmark, project # 0604-00716B and from the European Union Seventh Framework Programme (FP7-PEOPLE-2009-IRSES) under grant agreement #246647 (Project OptALI).

References

[1] Emile Aarts, Jan Korst, and Wil Michiels. Simulated Annealing. *Search Methodologies: Introductory Tutorials in Optimization and Decision Support Techniques, Search Methodologies: Introductory Tutorials in Optimization and Decis. Support Techniques*, pages 187–210, 2005.

[2] Erwin Abbink, Bianca van den Berg, Leo Kroon, and Marc Salomon. Allocation of Railway Rolling Stock for Passenger Trains. *Transportation Science*, 38(1):33–41, February 2004.

[3] Ravindra K. Ahuja and Claudio B. Cunha. Network Models in Railroad Planning and Scheduling. In *Tutorials in Operations Research, Informs*, pages 54–101. 2005.

[4] Ravindra K. Ahuja, James B. Orlin, Dushyant Sharma, and Larry A. Shughart. Solving Real-Life Locomotive-Scheduling Problems. *Transportation Science*, 39(4):503–517, 2005.

[5] Arianna Alfieri, Rutger Groot, Leo Kroon, and Alexander Schrijver. Efficient Circulation of Railway Rolling Stock. *Transportation Science*, 40(3):378–391, August 2006.

[6] Cynthia Barnhart, Peter Belobaba, and Amedeo R. Odoni. Applications of Operations Research in the Air Transport Industry. *Transportation Science*, 37(4):368–391, November 2003.

[7] Berkay Beygo. *Improvement of Decision Support System for Rail Depot Planning*. Master of science thesis, Chalmers University of Technology, Gothenburg, Sweden, 2011.

[8] Ralf Borndörfer, Markus Reuther, Thomas Schlechte, and Steffen Weider. A Hypergraph Model for Railway Vehicle Rotation. *OpenAccess Series in Informatics (OASIcs)*, 20:146–155, 2011.

[9] Ralf Borndörfer, Markus Reuther, Thomas Schlechte, and Steffen Weider. Integrated Optimization of Rolling Stock Rotations for Intercity Railways. Technical report, Konrad-Zuse-Zentrum für Informationstechnik, Berlin, 2012.

[10] Peter Brucker, Johann Hurink, and Thomas Rolfes. Routing of Railway Carriages. *Journal of Global Optimization*, pages 313–332, 2003.

[11] Valentina Cacchiani, Alberto Caprara, and Paolo Toth. A Fast Heuristic Algorithm for the Train Unit Assignment Problem. In Daniel Delling and Leo Liberti, editors, *12th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems*, volume 25 of *OpenAccess Series in Informatics (OASIcs)*, pages 1–9, Dagstuhl, Germany, 2012. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

[12] Luis Cadarso and Ángel Marín. Robust rolling stock in rapid transit networks. *Computers & Operations Research*, 38(8):1131–1142, August 2011.

[13] Alberto Caprara, Leo Kroon, Michele Monaci, Marc Peeters, and Paolo Toth. Passenger Railway Optimization. In *Handbook in OR &MS, Vol. 14*, volume 14. 2007.

[14] Jean-François Cordeau. A Benders decomposition approach for the locomotive and car assignment problem. *Transportation Science*, 34(March 2014):133–149, 2000.

[15] Jean-François Cordeau, François Soumis, and Jacques Desrosiers. Simultaneous Assignment of Locomotives and Cars to Passenger Trains. *Operations Research*, 49(March 2014):531–548, 2001.

[16] Jean-François Cordeau, Paolo Toth, and Daniele Vigo. A survey of optimization models for train routing and scheduling. *Transportation science*, (March 2014), 1998.

[17] Thomas H Cormen, Charles E Leiserson, and Ronald L Rivest. *Introduction to Algorithms, Second Edition*. The MIT Press, Cambridge Massachusetts, 2001.

[18] Jacob Dirksen. Parking and Litra Assignment for S-tog. 2010.

[19] Pieter-Jan Fioole, Leo Kroon, Gábor Maróti, and Alexander Schrijver. A rolling stock circulation model for combining and splitting of passenger trains. *European Journal of Operational Research*, 174(2):1281–1297, October 2006.

[20] Peter Foens. *Desicion Support for Depot Planning in the Railway Industry*. PhD thesis, Technical University of Denmark, 2006.

[21] Richard Freling, Ramon M. Lentink, Leo G. Kroon, and Dennis Huisman. Shunting of Passenger Train Units in a Railway Station. *Transportation Science*, 39(2):261–272, May 2005.

[22] Giovanni Luca Giacco, Andrea D'Ariano, and Dario Pacciarelli. Rolling Stock Rostering Optimization Under Maintenance Constraints. *Journal of Intelligent Transportation Systems*, 18(1):95–105, January 2014.

[23] Jørgen Thorlund Haahr, Richard Martin Lusby, Jesper Larsen, and David Pisinger. A Branch-and-Price Framework for Railway Rolling Stock Rescheduling During Disruptions. 2014.

[24] Frank Harary. *Graph Theory*. Addison Wesley, 1969.

[25] Dennis Huisman, Leo Kroon, Ramon M. Lentink, and Michiel J. C. M. Vromans. Operations Research in passenger railway transportation. *Statistica Neerlandica*, 59(4):467–497, November 2005.

[26] Stefan Irnich and Guy Desaulniers. Shortest Path Problems with Resource Constraints. In *Column Generation*, pages 33–65. Springer, 2005.

[27] Krishna C. Jha, Ravindra K. Ahuja, and Güvenç Sahin. New Approaches for Solving the Block-to-Train Assignment Problem. *Networks*, 51(1):48–62, 2008.

[28] Leo Kroon, Dennis Huisman, Erwin Abbink, Pieter-Jan Fioole, Matteo Fischetti, Gábor Maróti, Alexander Schrijver, Adri Steenbeek, and Roelof Ybema. The New Dutch Timetable: The OR Revolution. *Interfaces*, 39(1):6–17, February 2009.

[29] Leo Kroon, Ramon M. Lentink, and Alexander Schrijver. Shunting of Passenger Train Units: An Integrated Approach. *Transportation Science*, 42(4):436–449, November 2008.

[30] Sean Luke. *Essentials of Metaheuristics*. 2014.

[31] Etienne-Jules Marey. *La Méthode graphique dans les sciences expérimentales et principalement en physiologie et en médecine*. G. Masson, Paris, 1885.

[32] Gábor Maróti and Leo Kroon. Maintenance routing for train units: The interchange model. *Computers & Operations Research*, 34(4):1121–1140, April 2007.

[33] Otto Anker Nielsen. *Personal communication*. Technical University of Denmark, DTU Transport, Dept. of Transport, August 2013.

[34] Marc Peeters and Leo Kroon. Circulation of railway rolling stock: a branch-and-price approach. *Computers & Operations Research*, 35(2):538–556, February 2008.

[35] Matias S Rasmussen, Richard M Lusby, David M Ryan, and Jesper Larsen. Subsequence Generation for the Airline Crew Pairing Problem. Technical Report May, Technical University of Denmark, DTU Management Engineering, 2011.

[36] Andreas Razen. Einsatz-Management von Loks und Lokpersonal im Güterverkehr der SBB Cargo. 2004.

[37] David M Ryan. *The Solution of Massive Generalized Set Partitioning Problems in Aircrew Rostering*. University of Auckland, New Zealand, 1992.

[38] Mohammed Saddoune, Guy Desaulniers, Issmail Elhallaoui, and François Soumis. Integrated Airline Crew Pairing and Crew Assignment by Dynamic Constraint Aggregation. *Transportation Science*, 46(1):39–55, February 2012.

[39] Per Thorlacius. Rolling Stock Planning at DSB S-tog, Working paper #1: Processes, Cost Structures and Requirements, May 2014.

[40] Balachandran Vaidyanathan, Ravindra K. Ahuja, Jian Liu, and Larry A. Shughart. Real-life locomotive planning: New formulations and computational results. *Transportation Research Part B: Methodological*, 42(2):147–168, February 2008.

[41] Alexandre Sieira Vilar. What's Object Pooling All About? *Dr. Dobb's Journal: Software Tools for the Professional Programmer*, 25(8), 2000.