Transportation Research Part B 86 (2016) 250–267

Contents lists available at ScienceDirect

Transportation Research Part B

journal homepage: www.elsevier.com/locate/trb

Reinforcement learning approach for train rescheduling on a single-track railway

D. Šemrov[a,*], R. Marsetič[a], M. Žura[a], L. Todorovski[b], A. Srdic[a]

[a] University of Ljubljana, Faculty of Civil and Geodetic Engineering, Jamova 2, 1000 Ljubljana, Slovenia
[b] University of Ljubljana, Faculty of Administration, Gosarjeva 5, 1000 Ljubljana, Slovenia

ARTICLE INFO

Article history:
Received 12 December 2014
Revised 23 December 2015
Accepted 7 January 2016

Keywords:
Train rescheduling
Artificial intelligence
Reinforcement learning
Q-learning

ABSTRACT

Optimal rail network infrastructure and rolling stock utilization can be achieved with use of different scheduling tools by extensive planning a long time before actual operations. The initial train timetable takes into account possible smaller disturbances, which can be compensated within the schedule. Bigger disruptions, such as accidents, rolling stock breakdown, prolonged passenger boarding, and changed speed limit cause delays that require train rescheduling. In this paper, we introduce a train rescheduling method based on reinforcement learning, and more specifically, Q-learning. We present here the Q-learning principles for train rescheduling, which consist of a learning agent and its actions, environment and its states, as well as rewards. The use of the proposed approach is first illustrated on a simple rescheduling problem comprising a single-lane track with three trains. The evaluation of the approach is performed on extensive set of experiments carried out on a real-world railway network in Slovenia. The empirical results show that Q-learning lead to rescheduling solutions that are at least equivalent and often superior to those of several basic rescheduling methods that do not rely on learning agents. The solutions are learned within reasonable computational time, a crucial factor for real-time applications.

1. Introduction

The effectiveness of a system is often measured by its reliability. This can be defined as the ability of the system to carry out its assigned function in accordance with the required constraints, and its objective function within a certain time. Railway system reliability is most commonly evaluated by train punctuality, which refers to the difference between the actual and expected arrival time of the train, specified in the timetable, at the final destination. The punctuality is one of the most important characteristics that users consider when selecting a transport mode.

In railway traffic, short delays due to variations in acceleration/deceleration profile of trains, their driving patterns and speed profiles are understandable and cannot be avoided. Therefore, running time supplements and buffer times are introduced into timetables (Kecman et al., 2012). These supplements can compensate only for small disruptions, such as a different speed profile due to weather conditions, but cannot compensate for larger delays. Disruptions and unexpected events, such as accidents, rolling stock breakdowns, infrastructure failures, prolonged stopping due to passenger boarding, and changed speed limits cause larger delays in the railway traffic that require more elaborate traffic management and timetable rescheduling. Dynamic train traffic management is indispensable for maintaining punctuality of the trains and

* Corresponding author. Tel.: +386 41 605104; fax +386 1 4250692.
E-mail addresses: darja.semrov@fgg.uni-lj.si (D. Šemrov), aleksander.srdic@fgg.uni-lj.si (A. Srdic).

http://dx.doi.org/10.1016/j.trb.2016.01.004
0191-2615/

minimizing the consequences of the delays. The reactive set of rescheduling actions performed by dispatchers in the control center must be feasible and effective. Train rescheduling is a real-time process where the dispatchers have only a few minutes to respond to the delay. Thus, they cannot examine all the feasible solutions to find the optimal one, and take their decisions using experience and intuition (D'Ariano et al., 2008; Hara et al., 2006).

This paper introduces a novel approach to the train rescheduling task that is based on reinforcement learning, more specifically Q-learning. Reinforcement learning has already been effectively used in traffic engineering, namely in traffic signal control (Abdulhai et al., 2003), signal plan optimization for straight traffic flow at intersection (Gregoire et al., 2007), in intersections (Shoufeng et al., 2008), road network (Prashanth et al., 2011) and (Arel et al., 2010), and motorway access (Veljanovska et al., 2010). Q-learning was also successfully used in railway engineering, e.g., for the tasks of light rail dispatching (Zou et al., 2006), railway access negotiation between the track owner and train service providers (Wong and Ho, 2010), and train marshalling (Hirashima, 2011, 2012). Note however, that none of the applications of reinforcement or Q-learning to railway traffic engineering addresses the particular problem of train rescheduling.

This paper is organized as follows. Section 2 reviews the previous studies and proposed approaches to the train rescheduling problem. Section 3 provides formal description of the reinforcement learning principles and introduces Q-learning. Section 4 introduces our approach to train rescheduling with Q-learning by systematically explaining the rationale behind the particular design choices. Section 5 reports on the empirical evaluation of the proposed approach on a simple artificial example and a complex one on the Slovenian railway infrastructure. The simple example aims at illustrating the utility of our approach, while the real-world example aims at evaluating its usability in a real-world environment. Section 6 discusses the results in the context of the related work. Finally, Section 7 summarizes the contribution of the paper and presents venues for future research.

2. Related work

The train rescheduling task is a complex and very challenging problem and therefore many different approaches to its solution have been proposed. We can classify the approaches to train rescheduling along the four dimensions of (1) the level of interruption (small disturbances vs. large disruptions), (2) the level of railway network detail (microscopic vs. macroscopic network models), (3) the optimization objective (e.g., number of delayed trains, total delay of all trains or total delay of passengers), and (4) the solution method.

Following the first classification dimension, we distinguish between two levels of interruption: disturbance and disruption. The "relatively small" perturbations (i.e., disturbances) influencing the rail operations can be handled only by rescheduling the timetable. When the perturbation is "relatively large" (i.e., a disruption), the resource duties must be rescheduled as well. In literature, the term "small delay" commonly refers to delays within the upper limit of 30 min (e.g., Acuna-Agost et al., 2011; Dotoli et al., 2013).

Recent studies of the train rescheduling problem differ also in the level of details considered in the railway models. We can classify the models into two large classes of microscopic and macroscopic models. The later commonly neglects the infrastructure capacity, such as the number of tracks at the station, and the traffic safety constraints. At the microscopic level the train location is known at the level of block sections, and the running time is computed accurately, due to detailed information of track layout and signals. The more detailed the model, the more computationally demanding is the calculation of the new timetable. Some papers considering macroscopic model are: Acuna-Agost et al. (2011), Min et al. (2011), Törnquist and Persson (2007), while papers considering the microscopic model are: Corman et al. (2010a), D'Ariano at al. (2007), Gély et al. (2006), Rodriguez (2007).

Regarding the optimization objective, we can identify two types of approaches: single-objective and multi-objective formulation. In both cases each objective can refer to the passengers, operators or combination of both. Majority of publications evaluate the effectiveness of proposed algorithm according to the ability to minimize the total delay (D'Ariano et al., 2008; Geske, 2006; Tazoniero et al., 2007). Researches also propose to evaluate the rescheduling for multi-objective functions, e.g. lowering the fuel consumption cost, and minimizing passenger travel time (Ghoseiri et al., 2004; Medanic and Dorfman, 2002), and minimizing passenger annoyance, such as the change of platforms and missed connections (Wegele and Schnieder, 2004).

Researchers have already proved the train rescheduling problem to be NP-hard (Sajedinejad et al., 2011; Ping et al., 2001) and NP-complete (Corman et al., 2011a; Ge, 2009). In both cases, the space of candidate solutions exponentially grows with the problem size. As a consequence, dynamic programming approach that would guarantee the optimality of the obtained solution, is often not applicable due to the complexity of the timetable and the railway infrastructure. That is why most of the approaches to train rescheduling employ heuristic methods.

Most commonly the problem is modeled as a job-shop problem (Cacchiani et al., 2014; Mascis and Pacciarelli, 2002) with different constraints. Besides the most known and widely used linear mixed programming (Acuna-Agost et al., 2011; Gély et al., 2006; Narayanaswami and Rangaraj, 2013); researchers have proposed and evaluated many other heuristic search approaches, such as: depth-first search (Törnquist Krasemann, 2012), branch and bound (D'Ariano et al., 2007), genetic algorithms (Dündar and Şahin, 2013), tabu search (Corman et al., 2010a), simulated annealing (Törnquist and Persson, 2005), and ant-colony optimization (Fan et al., 2012). The main criterion for comparative evaluation of different methods is the trade-off between the computation time and the solution quality.