Journal of Rail Transport Planning & Management 26 (2023) 100389

Contents lists available at ScienceDirect

Journal of Rail Transport Planning & Management

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

An evaluation of the fairness of railway timetable rescheduling in the presence of competition between train operators

Edwin Reynolds [a], Matthias Ehrgott [b,*], Judith Y.T. Wang [c]

[a] STOR-i, Lancaster University, Bailrigg Lancaster, LA1 4YX, United Kingdom
[b] Department of Management Science, Lancaster University, Bailrigg, Lancaster, LA1 4YX, United Kingdom
[c] School of Civil Engineering & Institute for Transport Studies, University of Leeds, Leeds, LS2 9JT, United Kingdom

ARTICLE INFO

Keywords:
Railway optimisation
Fairness
Timetable rescheduling

ABSTRACT

Using the output of optimisation models to make real-time changes to railway timetables can be an effective way to reduce the propagation of delay. In this study, we develop a methodology for evaluating the fairness of such optimisation models with respect to competing train operators. Whilst both fairness and optimisation-based railway timetable rescheduling have both been widely studied, they have not previously been studied together. We propose definitions of fairness and efficiency for timetable rescheduling, and analyse the fairness of efficiency-maximising solutions for a case study with seven train operators. We also investigate the pairwise trade-offs between operators and show that the priority given to different train classes has an important impact on fairness.

1. Introduction

Effective real-time management of railway traffic is crucial to delivering good railway performance. In particular, making changes to the timetable in response to an initial delay can help to reduce the amount of additional delay caused to other trains as a result of the initial incident. This practice is known as *timetable rescheduling*. The Train Timetable Rescheduling Problem (TTRP) (Cacchiani et al., 2014) can be solved in order to determine the optimal way to reschedule the timetable. A large number of different TTRP problem variants, models, objective functions and solution methods have been studied.

However, the implications for TTRP models of economic competition between railway operators has not been considered. In recent decades, different forms of competition have been introduced in several European railway systems, such as those of Germany, Great Britain and Sweden (IBM, 2011). Where trains are operated by more than one different company over the same tracks, timetable rescheduling has the potential to impact these operators unequally. In order to be perceived as fair, a TTRP model must not systematically favour some operators over others. A perception of unfairness would be a serious barrier to the practical deployment of TTRP models in competitive railway systems. Therefore, it is essential that the fairness characteristics of such models are understood.

This study investigates the fairness of solutions obtained from solving the TTRP. We use a case study of Doncaster railway station in Great Britain during January 2017, during which time trains from seven different private passenger operators regularly used the station. The TTRP model proposed by Reynolds et al. (2020) is used to calculate the rescheduled timetables that are evaluated. First, appropriate definitions of efficiency and fairness are proposed for the TTRP. Second, the fairness of efficiency-maximising solutions is analysed. Third, the sources of unfairness in efficiency-maximising solutions are elucidated via an investigation of the competitive relationship between each pair of operators. Finally, the changes in efficiency and fairness that arise from varying the

* Corresponding author.
E-mail addresses: edwinsreynolds@gmail.com (E. Reynolds), m.ehrgott@lancaster.ac.uk (M. Ehrgott), j.y.t.wang@leeds.ac.uk (J.Y.T. Wang).

https://doi.org/10.1016/j.jrtpm.2023.100389
Received 7 February 2022; Received in revised form 17 February 2023; Accepted 14 April 2023
Available online 2 May 2023

priority given to class 1 (express passenger) trains over class 2 (ordinary passenger) trains is examined. Class 1 trains are often prioritised in current manual rescheduling practice because they typically carry more passengers and travel longer distances.

This research has been carried out in partnership with Network Rail, a public sector company that owns and operates railway infrastructure in Great Britain. The railway system in Great Britain is vertically separated, meaning that passenger and freight services are operated by companies that are separated from Network Rail. These companies are called Train Operating Companies (TOCs) and Freight Operating Companies (FOCs), respectively — we will simply call them *operators*. Most passenger operators run services under competitively tendered franchise agreements. These franchises each confer the right to operate a set of services grouped by type or region over a period of several years. Open-access operators, on the other hand, run services in purchased timetable slots. The organisational structure of the system, which is likely to change in the future, is described by Williams (2019).

Fairness to operators is taken into consideration by Network Rail during the timetable rescheduling process. Rescheduling is currently carried out in accordance with local *Route Regulating Policies*. One of the stated objectives of these policies is to "minimise delay to all [operators] … in the best interests of all" — unfair treatment of an operator would contradict this. Moreover, delay management teams in Network Rail's Rail Operating Centres (where timetable rescheduling decisions are made) include representatives from operators working collaboratively to achieve this goal. The introduction of TTRP models to these settings relies on gaining the acceptance of key stakeholders such as operators, which is unlikely unless any proposed model can be shown to be fair.

The railway system in Great Britain has a performance-based financial compensation mechanism called Schedule 8 (see Network Rail (2019)). Under this system, all delays are recorded and attributed to a particular cause. Every four weeks, this information is used to calculate the value of monetary payments to be made from Network Rail to operators and vice versa. This mechanism partially protects operators from losses resulting from events that are outside of their control. For example, if an incident caused by one operator causes delay to the train of a second operator, this second operator should receive financial compensation. However, the extent to which cash payments can compensate for delays is limited. This is partly because accurate valuation of the losses resulting from reputational damage to the operator is challenging. It is also unsatisfactory for the passengers of operators, who may derive social benefits from train punctuality that cannot be financially valued. It is still preferable, therefore, to carry out timetable rescheduling as fairly as possible, and reduce reliance on financial compensation.

This paper is organised as follows. In Section 2, the relevant literature is reviewed. Our definitions of fairness and efficiency are described in Section 3. The case-study is introduced in Section 4. In Subs Section 4.1, an analysis of the fairness of efficiency-maximising TTRP solutions is given. This is supplemented in Section 4.2 by an analysis of the interactions between pairs of operators. Finally, we consider the fairness-efficiency trade-off in Section 4.3. The concluding Section 5 offers conclusions and suggests some possible avenues of further research.

2. Literature review

The literature addressing the TTRP is already very well documented by a number of recent reviews (Cacchiani et al., 2014; Corman and Meng, 2015; Fang et al., 2015; Lamorgese et al., 2018). Our literature review will focus specifically on fairness, and objective functions. An overview of relevant literature on the topic of fairness is given in Section 2.1. This is followed up in Section 2.2 by a discussion of how these concepts apply to objective functions that have been used in the TTRP literature.

2.1. Fairness

The concept of fairness has been extensively studied by economists. It most often arises in the context of allocating or sharing limited benefits or resources between distinct entities. This is precisely the scenario faced during timetable rescheduling. We can think of the distinct entities as the different operators, and the resources to be allocated as the available track capacity. Track capacity consists of segments of track over time, and is limited by the available infrastructure.

Much of the classical literature considering fairness considers the social welfare problem — how a central planner should allocate goods in an economy. In the framework developed by Bergson (1938) and Samuelson (1947), the preferences of each entity $i \in \{1, \dots, n\}$ can be represented by a cardinal utility $u_i \geq 0$ that depends on the allocation that entity i receives. The set $\mathcal{U} \subset \mathbb{R}^n_+$ of vectors of utilities $(u_1, \dots, u_n)$ that arise from feasible allocations is called the *utility possibility set*. The problem for the central planner can be formulated as choosing a point from this set. One way of choosing a point is to find a point that maximises a *social welfare function* $f : \mathcal{U} \to \mathbb{R}^+$ that encodes value judgements about the size and distribution of the utilities of the entities.

Social welfare functions can be used to model different attitudes about inequality and fairness. For example, the utilitarian function $f(u) = \sum_{i=1}^n u_i$ is completely indifferent to inequality — it values improvements in the utility of each entity equally, regardless of how high the utility already is. At the other end of the scale, the minimax function $f(u) = \min_{i=1,\dots,n} u_i$ always seeks improvement in the utility of the worst-off entity, and is indifferent to improvements in the utility of any other entity. Atkinson (1970) proposed maximising the α-fairness welfare function

$$W_\alpha(u) = \begin{cases} \sum_{i=1}^n \frac{u_i^{1-\alpha}}{1-\alpha} & \alpha \geq 0, \alpha \neq 1 \\ \sum_{i=1}^n \log(u_i) & \alpha = 1, \end{cases}$$

which uses the parameter α to interpolate between these two extreme attitudes to inequality. When $\alpha = 0$, the function becomes the utilitarian function and when $\alpha \to \infty$, it converges to the minimax function. Another special case, when $\alpha = 1$, corresponds to the rule proposed by Nash (1950). For any value of α, the function W_α has the property of constant elasticity: given a fixed proportional increase in an entity's utility, the proportional increase in the welfare of that entity does not depend on the size of their utility. A detailed introduction to the theory of economic inequality is provided by Foster and Sen (1997).

Another framework in which fairness has been studied is the bargaining problem. In this problem, distinct entities must agree on how to share a jointly generated surplus. Although this is not exactly the situation that arises in timetable rescheduling, it is still an interesting way to understand fairness. For the two-player version, both Nash (1950) and Kalai and Smorodinsky (1975) have proposed axioms that must be satisfied by solutions to this problem before they can be considered fair. Each of these two studies additionally proposes solutions that satisfy these axioms. However, there is no general consensus on which set of axioms should be used. Moreover, axiomatic approaches are not useful in this application because they cannot be used to evaluate the fairness of a solution that does not satisfy the axioms.

In most applied decision making scenarios in which fairness is important, it is nevertheless not the only consideration. The fairness of any solution must usually be considered alongside its overall efficiency. The efficiency can be defined as the sum of the utilities of the entities (the utilitarian function), or in a problem specific way if this is more appropriate. Unsurprisingly, a trade-off has been observed between efficiency and fairness in many problems. In other words, greater fairness can often only be achieved by sacrificing some efficiency, whilst greater efficiency entails a reduction in fairness. This trade-off has been characterised theoretically by Bertsimas et al. (2011, 2012) using the concept of the price of fairness. This is defined as the proportion of total system efficiency lost by maximising a fairness objective compared with maximising efficiency. They argue that analysing the price of fairness is a useful way to navigate the fairness-efficiency trade-off.

2.2. Fairness and the TTRP

As far as the authors are aware, only few previous studies have analysed the fairness of a TTRP model with respect to different operators or other entities such as trains or passengers. One potential reason for this is that only a few railway systems are sufficiently liberalised to have multiple operators regularly running trains over the same track. Moreover, fairness and competition in railway markets are often less important to policy makers than overall system performance.

However, fairness has been studied for similar applications of optimisation to transport planning problems, including railway planning problems. For example, Gestrelius et al. (2020) consider the evaluation of railway timetables from a competition management perspective. They report that research literature and guidelines for facilitating competition between operators in timetable design are scarce. In a more quantitative study, Li et al. (2019) investigate the fairness-efficiency trade-off for a train timetabling application, in which α-fairness between passengers was considered. Other examples of similar applications in which fairness has been considered include railway crew scheduling (Jütte et al., 2017), railway crew rostering (Breugem et al., 2019), air traffic flow management (Bertsimas and Gupta, 2016) and airport slot allocation (Fairbrother et al., 2020).

Many different objective functions have been used for the TTRP. Almost all of these can be written in the form $\max f(u_1(x), \dots, u_n(x))$, where x is the vector of variables in the optimisation problem, each u_i is a function measuring the quality of the solution for train i, and f is a real-valued function. Examples of quantities measured by u_i (or $-u_i$ in the case of minimisation) include consecutive delay (delay experienced as a result of interaction with other trains during the rescheduled period) (Corman et al., 2010a; Pellegrini et al., 2014), deviation from planned schedule (Meng and Zhou, 2014; Lusby et al., 2013), a piecewise linear function of delay (Lamorgese and Mannino, 2015), and the cost of delays (Törnquist and Persson, 2007). In many studies, u_i aggregates several measures, either through development of a utility function (Harrod, 2011; Binder et al., 2017; Reynolds et al., 2020) or otherwise (Bettinelli et al., 2017; Caimi et al., 2012; Acuna-Agost et al., 2011). Luan et al. (2017) have considered non-discrimantory train dispatching. They formulate the equity of train delays as soft objectives or hard constraints when minimising average train delay costs. The inequity between operators is formalised as the maximal individual deviation from the average delay cost. They investigate the trade-off between train delays and train equity and test their model on a line of the Dutch railway network between Utrecht and Den Bosch. Shang et al. (2018) consider the optimal scheduling of skip-stopping solutions in an oversaturated urban rail transit network from the perspective of passenger equity. They consider multi-commodity optimisation using train and passenger space–time networks and demonstrate the model on a line of the Beijing subway network. The most recent paper in this regard is Sun et al. (2022), who explicitly consider the distribution of delay among entities such as trains, passengers and operators. They define equity by means of a Gini coefficient and demonstrate results on the same network (line) as Luan et al. (2017).

From a fairness perspective, the choice of function f is more important. This is because it can be interpreted as a social welfare function where the entities are the trains, and the welfare of each train is a proxy for the welfare of the passengers on the train. In other words, f can encode a particular level of aversion to inequality between the performance of individual trains. The most popular choice in the literature is to use a sum, which corresponds to the utilitarian function (Törnquist and Persson, 2007; Meng and Zhou, 2014; Pellegrini et al., 2014; Lamorgese and Mannino, 2015). Weighted sums have also been used to establish train priorities (Lusby et al., 2013; Pellegrini et al., 2014). Since this can lead to large inequalities in the delays experienced by different trains, some authors such as D'Ariano et al. (2007) and Corman et al. (2010b), have used a minimum function, corresponding to minimax fairness. A computational comparison between minimising the total and maximum consecutive delay, respectively, has been carried out by Pellegrini et al. (2014). Despite the variety of functions f used and their implications for fairness, there has been very little discussion of fairness in the literature. As a result, the effect of using objective functions with different levels of inequality aversion is not well understood.

Several researchers have used multi-criteria methods to solve instances of the TTRP. Whilst fairness has not been explicitly formulated, some of these studies use objectives that could help to achieve fairness. For example, minimising the number of cancelled trains or missed connections helps to achieve fairness because cancelled trains or those sufficiently late to break connections are likely to be the worst affected trains. A weighted-sum approach is used by Caimi et al. (2012) to optimise the number of scheduled trains, the weighted delay and the number of connections maintained. Corman et al. (2012) use a specialised metaheuristic to consider the trade-off between total delay and the number of missed connections. Samà et al. (2015) present a Data Envelopment Analysis methodology for evaluating solutions under a large number of different punctuality measures, including objectives of minimax type. Corman et al. (2011) propose a lexicographic objective, in which trains are divided into different priority classes. Other studies adopting multi-criteria approaches include Ginkel and Schöbel (2007), Cavone et al. (2017), Binder et al. (2017), Shakibayifar et al. (2018), Josyula (2019), Altazin et al. (2020). None of these studies analyse the effect of using multi-criteria methods on fairness.

In this paper, we make the following contributions:

1. We propose measures for both the efficiency and fairness of solutions to a set of TTRP instances.
2. We identify the relative weighting of class 1 and 2 trains in the objective function as a key driver of unfairness between operators.
3. We then investigate the trade-off between efficiency and fairness that is created by changing this parameter using a case study based on real data from Doncaster station in the UK.

3. Efficiency and fairness

In order to evaluate the fairness of railway timetable rescheduling with respect to efficiency, these two terms must first be defined. As far as the authors are aware, no definition of fairness has been proposed specifically for the TTRP.

3.1. Efficiency

Our measure of the overall system efficiency was developed with Network Rail (see Reynolds et al. (2020)). It is designed to model the utility of Network Rail, which can be seen as the central decision maker for rescheduling decisions. It is designed to take into account the overall quality of service provided to passengers.

Given a feasible solution x to a given instance, the efficiency is defined as

$$U(x) = \sum_{k\in\mathcal{K}_1} U_k(x) + w \sum_{k\in\mathcal{K}_2} U_k(x), \tag{1}$$

where $\mathcal{K}_1$ and $\mathcal{K}_2$ are sets containing the class 1 and class 2 trains, respectively, $U_k(x)$ is the utility accrued from train k, and $w = 0.4$ is a weight that controls the priority given to class 1 trains in comparison to class 2 trains. The priority given to class 1 trains by the value of w reflects the fact that class 1 trains typically carry more passengers. Furthermore, class 1 trains usually complete longer journeys and hence delays to class 1 trains can have a greater impact in terms of reactionary delay outside the geographical scope of the TTRP instance.

The utility $U_k(x)$ accrued from train k is calculated as a weighted average across the set J^k of timetabled events for train k within the area and time horizon modelled:

$$U_k(x) = \sum_{j\in J^k} \beta_k^j U_k^j(x). \tag{2}$$

Each event $j \in J^k$ in the timetable for train k corresponds to a particular part of the track r_j^k, and a time that train k is due to enter it. These can include arrival events at platforms and passing events at junctions or key points along a route. The values of β_k^j ensure that more important events, such as an arrival into Doncaster station or exiting the modelled area, are weighted more highly than events at minor stations. When alternative platforms are available (such as at Doncaster station), these are separate events $j \in J^k$.

The utility $U_k^j(x)$ accrued by train k at event j is equal to zero if the event j is not carried out. This could occur if the train passes through a timetabled platform without stopping to let passengers alight or depart (a cancelled stop) or if the route of a train is changed so that it does not visit the location specified by event j. Otherwise, $U_k^j(x)$ is equal to

$$\gamma(l) = \begin{cases} \phi^{-\omega|l|} & \text{if } |l| \le \Gamma \\ 0 & \text{if } |l| > \Gamma, \end{cases} \tag{3}$$

where l is the number of 15 s time intervals late that the event occurs. Note that the lateness l is negative if the train is early, and this is penalised equally to positive lateness to discourage station congestion. The parameter values used are $\phi = 1 + (1 \times 10^{-7})$, $\omega = 1 \times 10^5$ and $\Gamma = 240$. The weights β_j^k are set as follows: A train's final arrival or passing event in the time horizon is given weight 0.7 to reflect the importance of trains leaving the controlled area punctually. All other arrivals excluding those representing platform alternatives receive an equal proportion of the remaining weight 0.3. Non-final passing events receive a weight of zero. When r_j^k is a platform alternative, β_j^k has the same weight as the planned platform arrival but discounted by a factor of 0.9. This reflects the disruption to passengers as a result of changing platforms. Other considerations could also be taken into account, such as the service patterns and passenger interchange importance of individual stations. A preference is shown for events occurring as close to on-time as possible. This is achieved by making γ strictly increasing on $[-\Gamma, 0]$ and strictly decreasing on $[0, \Gamma]$. This preference reflects the fact that late events are bad for passengers and early events can cause unanticipated difficulties at stations. The efficiency $E(\mathbf{x})$ of a set of solutions $\mathbf{x} = \{x^i : i \in I\}$ to the whole set of instances I can be calculated by summing the individual efficiency of each solution. Denoting the efficiency function U when applied to each instance i as U_i, this can be written as

$$E(\mathbf{x}) = \sum_{i\in I} U_i(x^i). \tag{4}$$

3.2. Finding an optimal solution

Solving a TTRP instance consists in finding a set of rescheduling events that maximises the value of $E(x)$. In Reynolds et al. (2020), we describe this solution method, which solves a multi-commodity network flow model. The constraints in this model represent the railway track capacities in the modelled station area over a period of time, discretised to 15 s time intervals. We develop a branch and price algorithm for this problem. For each train, the algorithm eventually finds a path through the railway infrastructure time-space network that defines its optimal movement in the rescheduled timetable and maximises the utility or efficiency $U_i(x)$. For further details on the solution method we refer the reader to Reynolds et al. (2020).

3.3. Fairness

For a given instance, let the set of operators be O, and let $\mathcal{K}_o \subset \mathcal{K}$ be the set of trains that are operated by operator o. The efficiency function $U(x)$ can be rewritten as

$$U(x) = \sum_{o\in O} U_o(x), \tag{5}$$

where

$$U_o(x) = \sum_{k\in\mathcal{K}_1\cap\mathcal{K}_o} U_k(x) + w \sum_{k\in\mathcal{K}_2\cap\mathcal{K}_o} U_k(x) \tag{6}$$

is the part of the efficiency arising from trains operated by o. Since $U_o(x)$ includes only trains from operator o, it can be used to measure the utility of operator o.

The utilities $U_o(x)$ are difficult to compare because there might be different numbers of trains with different weights in the instance. For each operator $o \in O$, let x_o^* be an optimal solution when the objective is to maximise $U_o(x)$. Each solution x_o^* represents the best solution operator o can hope for, and $U_o(x_o^*)$ provides an upper bound for $U_o(x)$. This allows us to calculate a normalised utility for each operator

$$\hat{U}_o(x) = \frac{U_o(x)}{U_o(x_o^*)}. \tag{7}$$

These values can be compared between operators. A value of $\hat{U}_o(x) = 1$ indicates that operator o realises their maximum possible utility in the rescheduled solution x — all events for all trains are due to be carried out on time and as planned. Conversely, $\hat{U}_o(x) < 1$ indicates that one or more events have been cancelled or rescheduled to occur on a different platform, or late.

A social welfare function (such as α-fairness) could be applied to the set of utilities $\{\hat{U}_o(x) : o \in O\}$ to measure the fairness of the solution x for a single instance. This would allow fairness to be formulated as an objective function so that fairness-maximising solutions to individual instances could be computed to solve the TTRP. It would also open the possibility of using multi-criteria methods to balance the objectives of maximising fairness and maximising efficiency within each instance.

However, when considering operator fairness for timetable rescheduling, it is problematic to focus on single hour-long instances separately. This is because operators experience fairness and unfairness over a much longer period of time. The operation of a TTRP algorithm on a railway is likely to involve solving hundreds of different instances, involving repeated allocations of track capacity between the same sets of operators. Instead, it is much more appropriate to consider many consecutive instances of the problem as a single, combined allocation problem. That is the approach taken in this paper.

Considering fairness over a whole instance set rather than on an individual instance basis has important implications for fairness. It means that each individual instance need not be fair, provided any operators that loose out can be compensated in other instances. This is crucial when one considers that a typical TTRP instance considers changes to the timetable over a time horizon of only one hour. Many instances involve only a small number of decisions such as which train should go ahead of the other out of a pair of conflicting trains. It may be impossible to resolve such problems in a fair way, or doing so may require a large degradation in efficiency. Our approach of considering fairness over the whole instance set overcomes this issue.

Although fairness is measured over a whole set of instances, each instance must still be optimised individually, as and when delays arise on the railway. This makes it impossible to optimise our measure of fairness. As a result, we do not suggest the measure as a potential objective function, but rather as a tool for fairness evaluation. Unfortunately, this also precludes calculation of the price of fairness (Bertsimas et al., 2011).

Consider a set of solutions $\mathbf{x} = (x^i : i \in I)$ to a set of instances $i \in I$. We index the previous notation by i so that $O_i, U_{i,o}, \hat{U}_{i,o}$ and $x_o^{i,*}$ correspond to the notation $O, U_o, \hat{U}_o$ and x_o^*, respectively, when applied to each instance i. Note that the set of operators O_i can be different across instances.

The normalised aggregated utility for operator o over I can be calculated as

$$\hat{U}_o(\mathbf{x}) = \frac{\sum_{i\in I : o\in O_i} U_{o,i}(x^i)}{\sum_{i\in I : o\in O_i} U_{o,i}(x_o^{i,*})}. \tag{8}$$

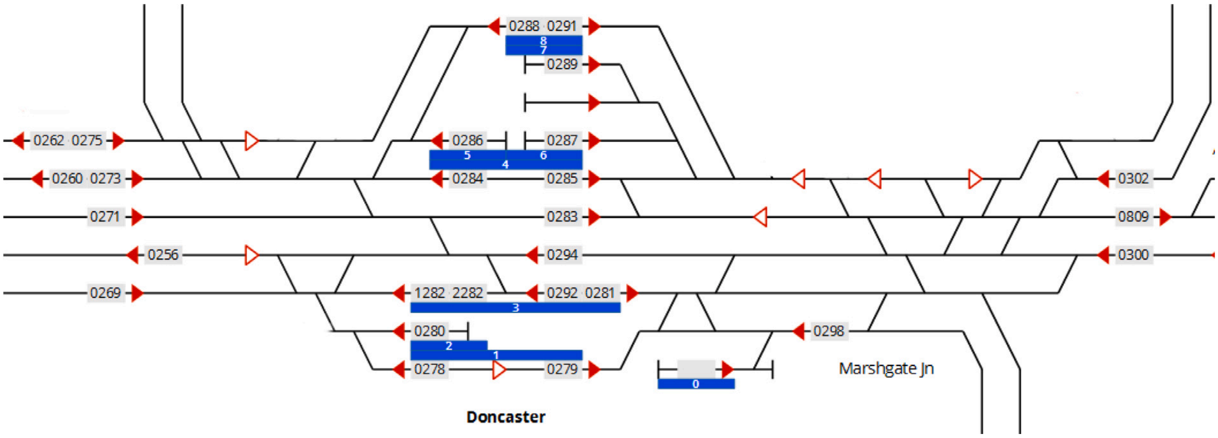

Fig. 1. A berth diagram of Doncaster Station. Retrieved from https://wiki.openraildata.com/index.php?title=DR. Accessed 14/04/20.

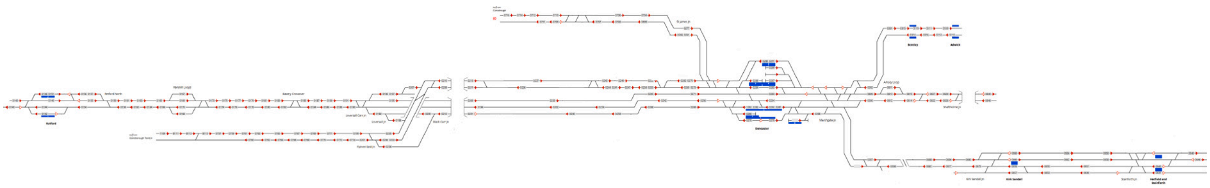

Fig. 2. A berth diagram of the area of track modelled. Retrieved from https://wiki.openraildata.com/index.php?title=DR. Accessed 14/04/20.

These values are then used in the α-fairness welfare function to produce our measure of fairness:

$$F_\alpha(\mathbf{x}) = \begin{cases} \sum_{o\in O} \frac{\hat{U}_o(\mathbf{x})^{1-\alpha}}{1-\alpha} & \alpha \geq 0, \alpha \neq 1 \\ \sum_{o\in O} \log \hat{U}_o(\mathbf{x}) & \alpha = 1. \end{cases} \tag{9}$$

3.4. The process for analysing the relationship between fairness and efficiency

In order to analyse this relationship we first compute the maximum utility for each operator by maximising $U_o(x)$. In order to make these numbers comparable (Table 1 shows that operators in our case study are very different), these values are then used to compute normalised aggregate utilities $\hat{U}_o(x)$ for each train operating company. From these, we deduce their α-fairness indicators (9). We further analyse the trade-offs between different operators by computing the additional utility accrued by an operator if trains operated by a operator are removed. This allows the computation of the proportional gain by potentially made by any train operating company in comparison to any other being removed from the dataset. These numbers in turn allow to quantify conflict between operators that compete for the same resources, i.e. track sections at specific times. We then explore the influence of parameter w on this conflict and thereby identify w as a main driver for unfairness. A similar analysis can be performed for any section of railway for which the problem of maximising the efficiency $U(x)$ as in (1) is possible with different sets of trains. We present a case study for Doncaster station in Section 4.

4. The case study

An area of railway around Doncaster station has been used as a case study for evaluating fairness. Definitions of any railway signalling terminology used below can be found in Reynolds et al. (2020). Doncaster station (see Fig. 1) lies on the East Coast Main Line, a busy railway corridor connecting London with Leeds, York, Newcastle and Edinburgh. The wider area covered (see Fig. 2) also contains portions of four double track lines that all begin at Doncaster and go towards Sheffield, Lincoln, Leeds and Hull, respectively. Figures 1 and 2 are published by the Open Rail Data Wiki under Creative Commons licence CC BY-SA 4.0, https://creativecommons.org/licenses/by-sa/4.0/. The area lies within a single area of signalling control, and contains 225 berths with 313 valid berth transitions. The station itself has 9 platforms and 85 track circuits. Doncaster station is an important interchange for a variety of inter-city and local services operated by seven different operators. It is also a busy bottleneck, with over 30 trains per hour at peak times. This makes it ideal for investigating the interactions between different operators.

Table 1
The passenger operators in the case study, including the total number and the number of class 1 and class 2 trains run through the modelled area during January 2017.

Operator	Abbr.	# Trains	# Class 1	# Class 2
Northern Trains	NT	4105	730	3375
London North Eastern Railway	LNER	3492	3492	0
Transpennine Express	TPE	825	825	0
Cross Country	XC	737	737	0
Grand Central	GC	496	496	0
Hull Trains	HT	354	354	0
East Midlands Railway	EMR	275	18	257

The data for the case study comes from January 2017. The seven different passenger operators running services through the area during this month are displayed in Table 1. Abbreviations will be used to refer to the operators throughout. Freight operators have been excluded from the analysis because freight services are often scheduled during less busy periods and are affected by rescheduling in different ways to passenger operators.

Table 1 shows that NT and LNER operated the most services during the period. Whilst most of the operators run inter-city (class 1) services, NT and East Midlands Railway run mostly local (class 2) services. NT, in particular, is responsible for the vast majority of class 2 trains. During this period, Grand Central and Hull Trains were open-access operators, while the remainder were operating franchises.

The month of January 2017 is split into 310 non-overlapping hour-long instances of the TTRP (between 8am and 6pm each day, for 31 days). These instances are created from real historical data about the timetable, and the traffic perturbations that actually occurred. The number of operators running trains in each instance ranges from two to six, with the most common number being five. By using instances that cover a whole month, we are able to understand fairness over the whole month, rather than on an instance-by-instance basis.

4.1. Evaluation of Fairness under Maximal Efficiency

The approach taken by Reynolds et al. (2020) and in many other TTRP studies is to maximise the efficiency of the system. It is therefore important to evaluate the fairness of these solutions. For each of the 310 instances, an efficiency-maximising solution was calculated using a solving time limit of 600 s. Only 12 instances were not solved to optimality within this time limit. For these instances, the best solution found during the time limit was selected, which was less than 1% away from optimality in all cases.

Fig. 3 shows three quantities for each operator: the normed aggregated utility $\hat{U}_o(\mathbf{x})$, the proportion of instances for which $\hat{U}_{o,i}(x^i) = 1$, and the number of instances $|\{i \in I : o \in O_i\}|$ in which the operator runs at least one train. This shows that NT had the smallest normed aggregated utility, followed by LNER. Other operators had higher figures, showing that there is inequality in the normed utility of operators. NT and LNER also run trains in the greatest number of instances and run the most trains in total. It can be seen that the normed utility of NT is only equal to 1 ($\hat{U}_{o,i}(x^i) = 1$) in 41.6% of instances, which is significantly lower than for other operators.

The overall figure for the α-Fairness over the full test set is −0.01689 (5 decimal places) with $\alpha = 1$. Fig. 4 shows the α-Fairness of each instance separately. From this it can be seen that the set of solutions as a whole is more fair than many of the individual instances in the test set. It is also apparent that fairer instances are more numerous than less fair instances.

The distribution of normed utility over the instances by operator is further visible in Fig. 5. For each operator, a boxplot is shown of $\hat{U}_{o,i}(x^i)$ for all of the instances $i \in I$ such that $\hat{U}_{o,i}(x^i) \neq 1$. These are instances in which the normed operator utility was less than the maximum possible amount available for that operator. It shows that while NT was affected a greater number of instances, the values of $\hat{U}_{o,i}(x^i)$ in these instances were on average closer to 1 than for most other operators. This might be because $\hat{U}_{o,i}(x^i)$ is measured as a proportion of the total possible utility, and therefore when an operator runs more trains in an instance, a delay to any one of them causes less proportional impact.

To summarise, the results in this section indicate that unfairness is present in efficiency-maximal timetable rescheduling solutions. NT achieves a smaller proportion of its total possible utility than any other operator, although this is due to small losses over many instances. The results also show that fairness varies between instances.

4.2. Identifying sources of unfairness

Values of $\hat{U}_{o,i}(x^i)$ that are less than 1 occur because TTRP solutions must allocate scarce track capacity to some operators in preference to others. TTRP instances contain trade-offs between the interests of different operators that must be resolved according to the objective function used. We examine these trade-offs in order to more fully understand the variation in normalised aggregated utility across different operators that leads to unfairness.

For a given instance, let $x^*_{\setminus o}$ be the optimal solution when it is solved with objective function

$$\sum_{o' \in O \setminus \{o\}} U_{o'}(x), \tag{10}$$

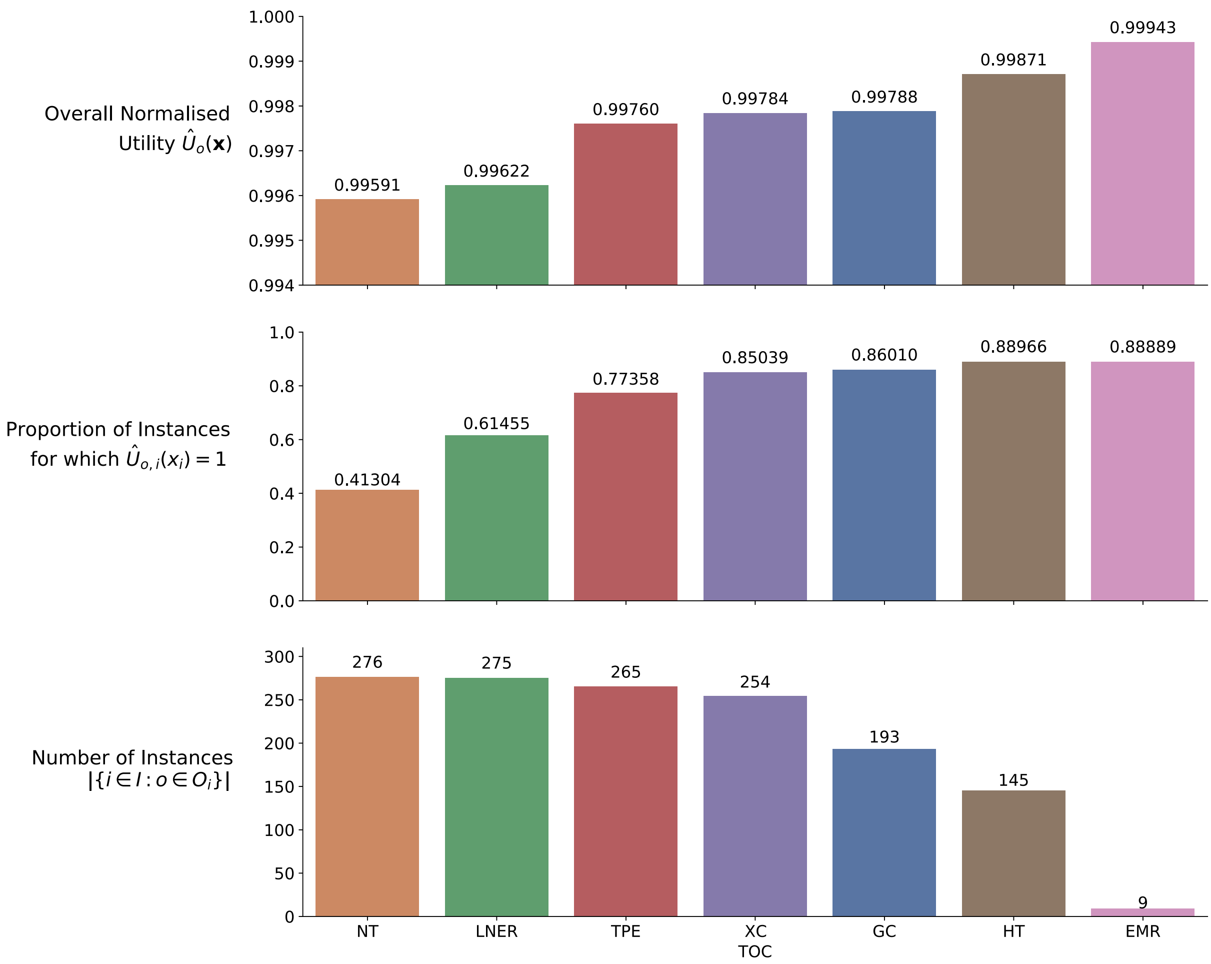

Fig. 3. Normed aggregated utility, number of instances with utility of 1, and number of instances for each operator.

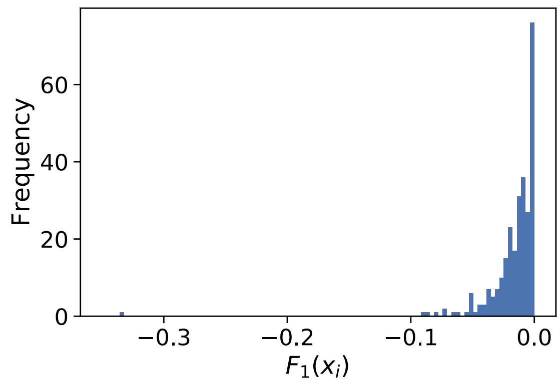

Fig. 4. Histogram of α-fairness scores of instances, with $\alpha = 1$.

for each operator $o \in O$. This objective function is obtained from the efficiency function by removing the utility accrued from any trains operated by o. For each remaining operator $o' \in O \setminus \{o\}$, the proportional gain in their utility is

$$\hat{U}_{o'\setminus o}(x) = \frac{U_{o'}(x^*_{\setminus o}) - U_{o'}(x)}{U_{o'}(x)}, \tag{11}$$

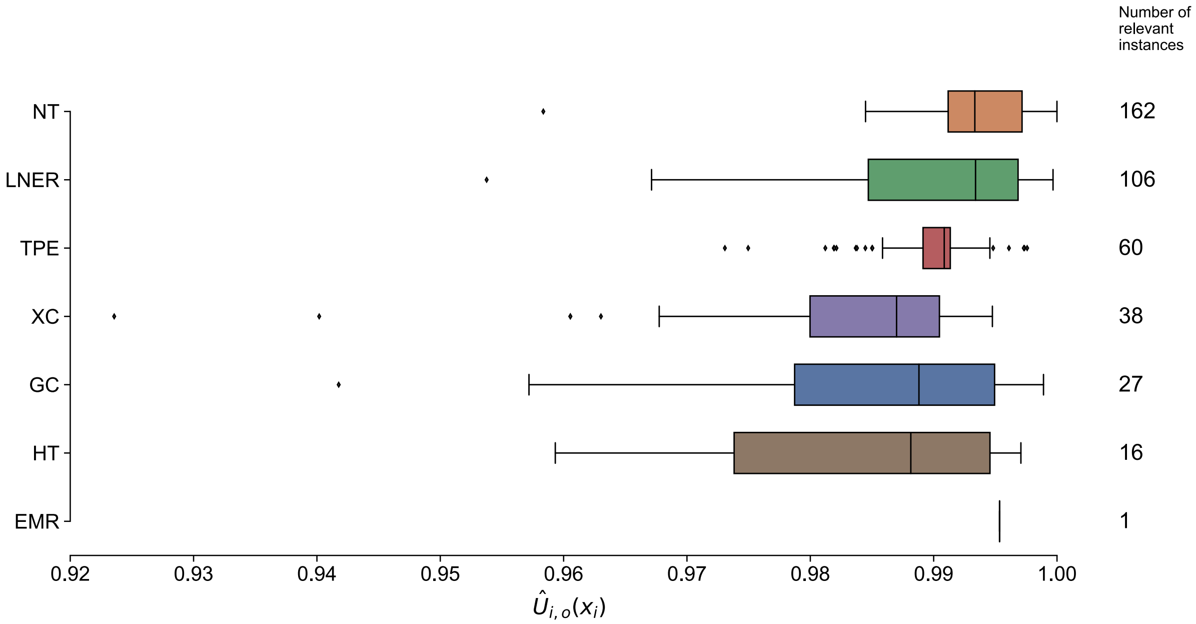

Fig. 5. A boxplot showing the distribution of $\hat{U}_{i,o}(x_i)$ over the set of instances $i \in I$, by operator. Instances i for which $\hat{U}_{i,o}(x_i) = 1$ are excluded. The number of instances that are included is shown on the right.

where x is an efficiency-maximising solution. Note that $o' \in O$ is necessary for this value to be defined, but that if $o \notin O$, we simply have that $\hat{U}_{o'\setminus o}(x) = 1$ (an operator that does not run trains in the instance is removed, so the utility of o' is not affected).

We can aggregate this over the whole instance set I. By letting $O_i, U_{i,o}, x^i$ and $x^{i,*}_{\setminus o}$ denote O, U_o, x and $x^*_{\setminus o}$ for instance $i \in I$ (as in Section 3.3), the proportional gain in aggregated utility can be written as

$$\hat{U}_{o'\setminus o}(\mathbf{x}) = \frac{\sum_{i\in I:o'\in O_i} U_{i,o'}(x^{i,*}_{\setminus o}) - \sum_{i\in I:o'\in O_i} U_{i,o'}(x^i)}{\sum_{i\in I:o'\in O_i} U_{i,o'}(x^i)}. \tag{12}$$

These values make it possible to identify the pairwise trade-offs between operators. A large value of $\hat{U}_{o'\setminus o}(\mathbf{x})$ means that operators o and o' are in greater competition for track capacity, whereas $\hat{U}_{o'\setminus o}(\mathbf{x}) = 0$ means that there is no competition.

The values of $\hat{U}_{o'\setminus o}(\mathbf{x})$ were calculated for each pair of operators using all 310 instances. They are displayed in Fig. 6 and represented visually in Fig. 8 as a directed graph in which the width of each directed arc (o', o) represents the magnitude of $\hat{U}_{o'\setminus o}(\mathbf{x})$. Fig. 7 shows the number of instances $i \in I$ for which $\hat{U}_{i,o'\setminus o}(x^i) > 0$ for each pair o, o' of operators.

Fig. 6 shows that LNER experienced a gain in utility of 0.2957% when NT was removed from every instance, the largest proportional increase of any pair of operators. Fig. 7 shows that when LNER was removed from the instances, NT improved its utility in 129 of the 310 instances, the largest number of instances of any pair of operators. From Fig. 8 it is apparent that LNER and NT both have significant trade-offs with a wide variety of different operators. Some pairs of operators that might have been expected to have significant trade-offs, because they both run large numbers of trains, did not have large trade-offs. TPE and XC are a good example of this, affecting the utility of each other only in 3 and 6 cases, respectively. This probably reflects characteristics of the timetable, in which their trains are rarely scheduled to pass through Doncaster Station at similar times. EMR experiences no utility trade-off, but this is not a significant finding since they operate trains in only 9 of the 310 instances.

An interesting aspect of the results displayed in Fig. 6 is that for some pairs of operators the value of $\hat{U}_{o'\setminus o}(\mathbf{x})$ is negative, although none of these are large in absolute value. Specifically, this is observed for (o', o) =(TPE, NT), (NT, TPE), (XC, HT) and (HT, TPE). Negative trade-offs over the whole instance set are observed because in 62 of the instances, there is some pair of operators for which $\hat{U}_{i,o'\setminus o}(x^i) < 0$. This means that when operator o was removed from the instance, the utility of operator o' decreased.

This phenomenon is known as *resource non-monotonicity*, and observing it shows that maximising the efficiency is a *resource non-monotonic* strategy. In this case, the resource in question is track capacity and the entities in question are the operators in the set $O \setminus \{o\}$. When operator o is removed from an instance, the collective availability of track capacity to the entities is increased since o is no longer competing for it. Formally, removing o from an instance weakly decreases the dual price of each *time-space resource* (see Reynolds et al. (2020) for an explanation of these terms). As a result, the entities are always collectively better off, i.e.

$$\sum_{o'\in O\setminus\{o\}} U_{o'}(x^{i,*}_{\setminus o}) \geq \sum_{o'\in O\setminus\{o\}} U_{o'}(x^i).$$

	NT	LNER	XC	TPE	GC	HT	EMR
NT	0	18.26	13.06	-0.49	0.63	0.2	0
LNER	29.57	0	4.33	5.83	3.57	1.37	0
XC	16.17	10.1	0	1.69	0.94	-0.08	0
TPE	-2.2	17.75	1.27	0	9.39	0.19	0
GC	9.58	17.89	0.73	3.66	0	0.32	0
HT	4.74	15.87	3.38	-1.66	0	0	0
EMR	5.74	0	0	0	0	0	0

Fig. 6. The values of $\hat{U}_{o'\setminus o}(\mathbf{x})$ multiplied by 100 (therefore given in 100$^{\text{ths}}$ of a percent) and rounded to 2 decimal places. Rows correspond to o'; columns to o.

	NT	LNER	XC	TPE	GC	HT	EMR
NT	0	129	115	12	16	3	0
LNER	91	0	37	40	34	13	0
XC	33	26	0	6	5	1	0
TPE	7	48	3	0	22	1	0
GC	12	27	2	7	0	2	0
HT	8	14	3	2	0	0	0
EMR	1	0	0	0	0	0	0

Fig. 7. The number of instances out of 310 for which $\hat{U}_{o'\setminus o}(x^i) > 0$. Rows correspond to o'; columns to o.

However, this collective improvement in utility is not shared equally under an allocation strategy of efficiency maximisation. Rather, we observe that whilst some operators benefit disproportionately, others experience a decrease in utility. This occurs when the removal of trains run by o from an instance causes an improvement for trains run by some operator $\hat{o}$ that is only possible with a degradation for trains run by o'.

4.3. The effect of train class weight on efficiency and fairness

The evidence from Sections 4.1 and 4.2 shows that in our case study, NT experiences the lowest normalised aggregated utility and that this arises as a result of trade-offs with the utility of LNER in particular. It is notable that NT also run the majority of class 2 trains in the instance set, and that these are down-weighted in the efficiency measure. Specifically, the train class weight is $w = 0.4$, meaning that the efficiency contribution of a class 2 train that runs exactly as planned is worth 0.4 of what it would if it

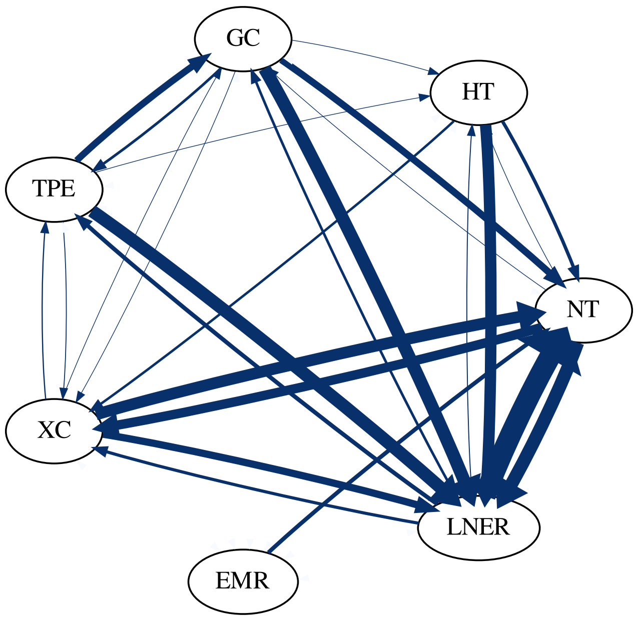

Fig. 8. A graph of the operators in which the width of arc (o', o) corresponds to $\hat{U}_{o'\backslash o}(\mathbf{x})$.

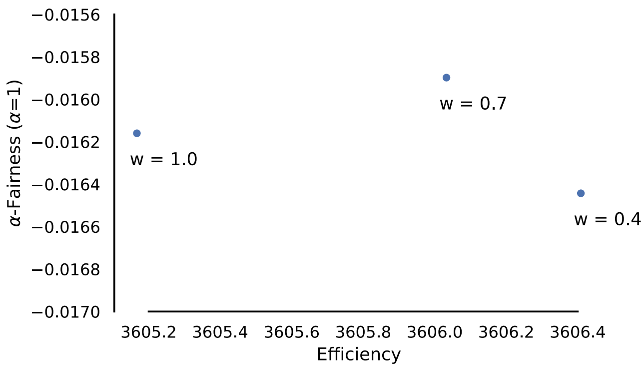

Fig. 9. Efficiency and fairness for different values of w.

were class 1. This suggests that an improvement in fairness might be achieved by using an objective function in which class 2 trains are down-weighted less severely.

To test this, all 310 instances were solved using the objective function U with three different values $w = 0.4, 0.7$ and 1.0. The scenario in which $w = 0.4$ is equivalent to maximising the efficiency, whereas equal weight is given to class 1 and class 2 trains in the $w = 1$ scenario. The final scenario, $w = 0.7$, was selected because it is halfway between these two extremes. For each value of w, the aggregated efficiency $E(\mathbf{x})$ and the aggregated fairness $F_\alpha(\mathbf{x})$ were calculated and are plotted in Fig. 9. Note that points to the right and up indicate higher efficiency and higher fairness and re therefore preferable. Fig. 10 compliments the information in Fig. 9 by showing the normalised aggregated utilities of the operators in each of the three scenarios.

These results show that as w is increased from 0.4 to 0.7, there is a trade-off between fairness and efficiency. The solutions obtained with $w = 0.7$ are more fair than $w = 0.4$, but less efficient. It is not surprising that the efficiency decreases, since function $U(x)$ with en $w = 0.4$ as shown in (1) is in fact the Network Rail definition of efficiency. Fig. 10 shows that the increase in fairness comes principally from an increase in utility for NT that is associated with a small cost for TPE and GC, with almost no difference for LNER. It is likely that this results in some of the conflicts between NT and TPE and between NT and GC from being decided in favour of NT as a result of the higher value of w.

A different pattern is observed as w is increased from 0.7 to 1.0. The efficiency continues to deteriorate as w deviates more from 0.4, and therefore the objective function deviates more from the efficiency measure. However, the rescheduling outcomes are less fair for $w = 1.0$ than for $w = 0.7$. Fig. 10 shows that the increase in w once again causes a substantial improvement to the utility of NT. However, this time a more significant decrease in the utilities of LNER and GC are observed. Since for $w = 0.7$ LNER already

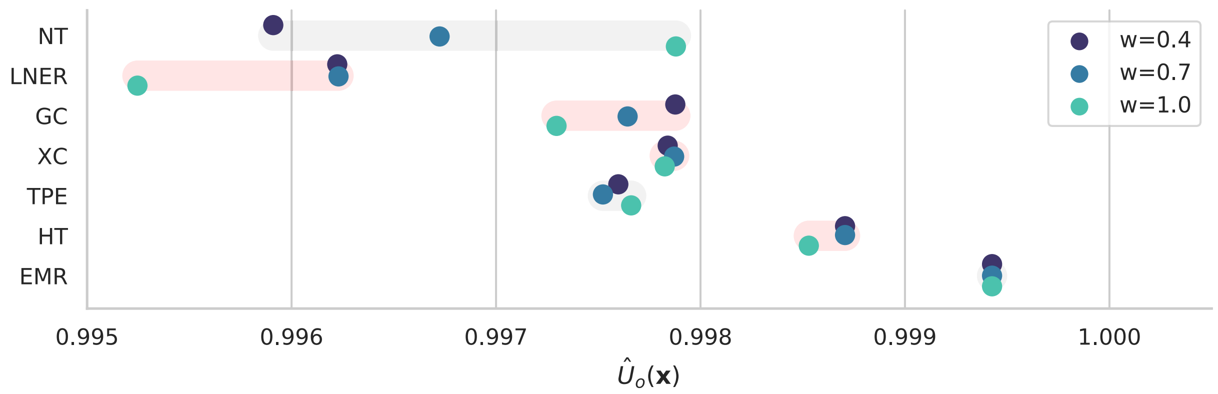

Fig. 10. The change in normalised aggregated operator utility for different values of w. Grey lines indicate that $\hat{U}_o(\mathbf{x})$ increases as w increases; red lines show the opposite.

has a lower normed aggregated utility than NT, this results in a decrease in fairness. The point for $w = 1.0$ is dominated by the point for $w = 0.7$, since it has both lower fairness and lower efficiency. This highlights that the points depicted in Fig. 9 are not (necessarily) Pareto optimal — bi-objective optimisation techniques would be required to find Pareto optimal points. However, as pointed out in Section 3.3, it is not possible to optimise fairness since this is measured over the whole instance set.

Fig. 10 shows that the largest changes in utility in response to changing w are experienced by NT, LNER and GC. These findings are consistent with the evidence from Section 4.2 that competition between NT and LNER is high. It is interesting to note that XC does not experience a significant fall in normalised aggregated utility, despite the fact that competition between NT and XC was also found to be high. This could indicate the importance of factors other than train weight in deciding how conflicts involving trains operated by NT and XC are resolved. Another possible explanation is that conflicts involving these two operators primarily relate to the small number of Class 1 trains operated by NT.

5. Conclusions

In this paper, we evaluate the fairness of a TTRP algorithm in the presence of competition between train operators. This is achieved by analysing results from a case study of Doncaster station during the month of January 2017. Measures for both efficiency and fairness are proposed and calculated both for each instance individually and over the whole month. Efficiency-maximising solutions are studied both to investigate fairness and to analyse pairwise trade-offs between operators. Finally, the effect of changing the train class weighting in the objective function is studied with respect to fairness and efficiency.

The results show that some unfairness is present in efficiency-maximising solutions. That unfairness is principally to the detriment of NT, the operator running most of the class 2 trains. Further, the largest trade-off in normalised aggregated utility is between NT and LNER. Increasing the train class weighting from 0.4 to 0.7 increases fairness by improving the utility of NT, but decreases efficiency by harming the utility of operators running class 1 trains. However, increasing w from 0.7 to 1.0 decreases both efficiency and fairness by harming the utility of LNER. This shows that the value of w needs to be carefully considered in any future deployment of optimisation-based timetable rescheduling.

Future research should focus on making direct comparisons between the fairness of TTRP solutions and the fairness of historical manually decided rescheduling actions. This would add to the evidence presented in this paper about the likely impact of deploying optimisation-based timetable rescheduling. Another direction of further study would be to identify how elements of the objective function other than w affect fairness. For example, it would be interesting to understand whether optimising the fairness of each individual instance leads to greater fairness over the whole instance set, and what the effect on efficiency would be. It might also be interesting to investigate the sensitivity of the solution to parameters β_j^k, even though we think that, because they refer to individual events for particular trains, their influence on the overall solution and fairness of the solution is small.

Recognising that optimising fairness of TTRP solutions over instances of one hour is of limited practical interest, it would certainly be of interest to develop solution methods that optimise fairness over a sequence of instances, such as the 310 instances we considered in the case study. This would require further research into solution methods for a set of instances. On the other hand it would allow a proper study of the trade-off between fairness and efficiency using multi-objective optimisation methods. Finally, it would be interesting to understand how operator fairness and passenger fairness interact, and how fairness for the TTRP ought to relate to the fairness of the original timetable construction process.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This paper is based on work completed while Edwin Reynolds was part of the EPSRC, United Kingdom funded STOR-i centre for doctoral training (EP/L015692/1). The authors also acknowledge the partial funding received from Network Rail, United Kingdom and their collaboration in providing data for this work.

References

Acuna-Agost, R., Michelon, P., Feillet, D., Gueye, S., 2011. SAPI: Statistical analysis of propagation of incidents. A new approach for rescheduling trains after disruptions. European J. Oper. Res. 215 (1), 227–243.
Altazin, E., Dauzère-Pérès, S., Ramond, F., Tréfond, S., 2020. A multi-objective optimization-simulation approach for real time rescheduling in dense railway systems. European J. Oper. Res. 286, 662–672.
Atkinson, A.B., 1970. On the measurement of inequality. J. Econom. Theory 2, 244–263.
Bergson, A., 1938. A reformulation of certain aspects of welfare economics. Q. J. Econ. 52 (2), 310–334.
Bertsimas, D., Farias, V.F., Trichakis, N., 2011. The price of fairness. Oper. Res. 59 (1), 17–31.
Bertsimas, D., Farias, V.F., Trichakis, N., 2012. On the efficiency-fairness trade-off. Manage. Sci. 58 (2020), 2234–2250.
Bertsimas, D., Gupta, S., 2016. Fairness and collaboration in network air traffic flow management: An optimization approach. Transp. Sci. 50 (1), 57–76.
Bettinelli, A., Santini, A., Vigo, D., 2017. A real-time conflict solution algorithm for the train rescheduling problem. Transp. Res. B 106, 237–265.
Binder, S., Maknoon, Y., Bierlaire, M., 2017. The multi-objective railway timetable rescheduling problem. Transp. Res. C 78, 78–94.
Breugem, T., Schulz, C., Schlechte, T., Borndörfer, R., 2019. A Three-Phase Heuristic for Cyclic Crew Rostering with Fairness Requirements. Technical Report, Zuse Institute Berlin.
Cacchiani, V., Huisman, D., Kidd, M.P., Kroon, L.G., Toth, P., Veelenturf, L.P., Wagenaar, J.C., 2014. An overview of recovery models and algorithms for real-time railway rescheduling. Transp. Res. B 63, 15–37.
Caimi, G., Fuchsberger, M., Laumanns, M., Lüthi, M., 2012. A model predictive control approach for discrete-time rescheduling in complex central railway station areas. Comput. Oper. Res. 39 (11), 2578–2593.
Cavone, G., Dotoli, M., Epicoco, N., Seatzu, C., 2017. A decision making procedure for robust train rescheduling based on mixed integer linear programming and data envelopment analysis. Appl. Math. Model. 52, 255–273.
Corman, F., D'Ariano, A., Hansen, I.A., Pacciarelli, D., 2011. Optimal multi-class rescheduling of railway traffic. J. Rail Transp. Plann. Manag. 1 (1), 14–24.
Corman, F., D'Ariano, A., Pacciarelli, D., Pranzo, M., 2010a. A tabu search algorithm for rerouting trains during rail operations. Transp. Res. Part B 44 (1), 175–192.
Corman, F., D'Ariano, A., Pacciarelli, D., Pranzo, M., 2010b. Centralized versus distributed systems to reschedule trains in two dispatching areas. Public Transport 2 (3), 219–247.
Corman, F., D'Ariano, A., Pacciarelli, D., Pranzo, M., 2012. Bi-objective conflict detection and resolution in railway traffic management. Transp. Res. C 20 (1), 79–94.
Corman, F., Meng, L., 2015. A review of online dynamic models and algorithms for railway traffic control. IEEE Trans. Intell. Transp. Syst. 16 (3), 1274–1284.
D'Ariano, A., Pacciarelli, D., Pranzo, M., 2007. A branch and bound algorithm for scheduling trains in a railway network. European J. Oper. Res. 183 (2), 643–657.
Fairbrother, J., Zografos, K.G., Glazebrook, K.D., 2020. A slot-scheduling mechanism at congested airports that incorporates efficiency, fairness, and airline preferences. Transp. Sci. 54 (1), 115–138.
Fang, W., Yang, S., Yao, X., 2015. A survey on problem models and solution approaches to rescheduling in railway networks. IEEE Trans. Intell. Transp. Syst. 16 (6), 2997–3016.
Foster, J., Sen, A., 1997. On Economic Inequality. Oxford University Press.
Gestrelius, S., Peterson, A., Aronsson, M., 2020. Timetable quality from the perspective of a railway infrastructure manager in a deregulated market: An interview study with Swedish practitioners. J. Rail Transp. Plann. Manag. 15, 100202.
Ginkel, A., Schöbel, A., 2007. To wait or not to wait? The bicriteria delay management problem in public transportation. Transp. Sci. 41 (4), 527–538.
Harrod, S., 2011. Modeling network transition constraints with hypergraphs. Transp. Sci. 45 (1), 81–97.
IBM, 2011. Rail liberalisation index. http://www.assorail.fr/wp-content/uploads/2015/08/Rail_Liberalisation_Index_2011.pdf.
Josyula, S.P., 2019. Parallel Algorithms for Real-Time Railway Rescheduling (Ph.D. thesis). Blekinge Institute of Technology, Department of Computer Science.
Jütte, S., Müller, D., Thonemann, U.W., 2017. Optimizing railway crew schedules with fairness preferences. J. Sched. 20 (1), 43–55.
Kalai, E., Smorodinsky, M., 1975. Other solutions to Nash's bargaining problem. Econometrica 43 (3), 513–518.
Lamorgese, L., Mannino, C., 2015. An exact decomposition approach for the real-time train dispatching problem. Oper. Res. 63 (1), 48–64.
Lamorgese, L., Mannino, C., Pacciarelli, D., Krasemann, J.T., 2018. Train dispatching. In: Borndörfer, R., Klug, T., Lamorgese, L., Mannino, C., Reuther, M., Schlechte, T. (Eds.), Handbook of Optimization in the Railway Industry. Springer, pp. 265–283.
Li, D., Zhang, T., Dong, X., Yin, Y., Cao, J., 2019. Trade-off between efficiency and fairness in timetabling on a single urban rail transit line under time-dependent demand condition. Transportmetrica B 7 (1), 1203–1231.
Luan, X., Corman, F., Meng, L., 2017. Non-discriminatory train dispatching in a rail transport market with multiple competing and collaborative train operating companies. Transp. Res. C 80, 148–174.
Lusby, R.M., Larsen, J., Ehrgott, M., Ryan, D.M., 2013. A set packing inspired method for real-time junction train routing. Comput. Oper. Res. 40 (3), 713–724.
Meng, L., Zhou, X., 2014. Simultaneous train rerouting and rescheduling on an N -track network : A model reformulation with network-based cumulative flow variables. Transp. Res. Part B 67, 208–234.
Nash, J.F., 1950. The bargaining problem. Econometrica 18 (2), 155–162.
Network Rail, 2019. Technical overview : Payments relating to disruption. https://www.networkrail.co.uk/wp-content/uploads/2019/03/Technical-overview-payments-relating-to-disruption.pdf.
Pellegrini, P., Marlière, G., Rodriguez, J., 2014. Optimal train routing and scheduling for managing traffic perturbations in complex junctions. Transp. Res. B 59, 58–80.
Reynolds, E., Ehrgott, M., Maher, S.J., Patman, A., Wang, J.Y.T., 2020. A multicommodity flow model for rerouting and retiming trains in real-time to reduce reactionary delay in complex station areas. (submitted for publication).

Samà, M., Meloni, C., D'Ariano, A., Corman, F., 2015. A multi-criteria decision support methodology for real-time train scheduling. J. Rail Transp. Plan. Manag. 5 (3), 146–162.
Samuelson, P.A., 1947. Foundations of Economic Analysis. Harvard University Press, Cambridge, MA.
Shakibayifar, M., Sheikholeslami, A., Jamili, A., 2018. A multi-objective decision support system for real-time train rescheduling. IEEE Intell. Transp. Syst. Mag. 10 (3), 94–109.
Shang, P., Li, R., Liu, Z., Yang, L., Wang, Y., 2018. Equity-oriented skip-stopping schedule optimization in an oversaturated urban rail transit network. Transp. Res. C 89, 321–343.
Sun, X., Luan, X., Corman, F., 2022. Inequity averse optimization of railway track management considering passenger route choice and Gini coefficient. SSRN.
Törnquist, J., Persson, J.A., 2007. N-tracked railway traffic re-scheduling during disturbances. Transp. Res. B 41 (3), 342–362.
Williams, K., 2019. Current railway models: Great Britain and overseas. https://www.gov.uk/government/collections/the-williams-rail-review.