Massachusetts Institute of Technology
Engineering Systems Division

ESD Working Paper Series

Train Timetable Design for Shared Railway Systems using a Linear Programming Approach to Approximate Dynamic Programming

Maite Peña-Alcaraz
(corresponding author)
Ph.D. Candidate
Engineering Systems Division
Massachusetts Institute of Technology
Cambridge, MA USA
maitepa@mit.edu

Andrés Ramos
Professor
Institute for Research in Technology
Comillas Pontifical University
Madrid, Spain
andres.ramos@iit.upcomillas.es

Mort Webster
Associate Professor of Energy and Mineral Engineering
Energy and Mineral Engineering
Pennsylvania State University
University Park, PA USA
mdw18@psu.edu

ESD-WP-2014-13 May 2014 esd.mit.edu/wps

Train Timetable Design for Shared Railway Systems using a Linear Programming Approach to Approximate Dynamic Programming

Maite Pena-Alcaraz[a] (maitepa@mit.edu), Mort Webster[b] (mdw18@psu.edu), Andrés Ramos[c] (andres.ramos@iit.upcomillas.es)

[a] Engineering Systems Division, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA

[b] Energy and Mineral Engineering, Pennsylvania State University, 123 Hosler Building, University Park, PA 16802, USA

[c] Institute for Research in Technology, Comillas Pontifical University, C/ Alberto Aguilera 23, 28015 Madrid, Spain

Corresponding author: Maite Pena-Alcaraz (maitepa@mit.edu), Room E40-246, 77 Massachusetts Avenue, Cambridge, MA 02139, USA.

Abstract— **In the last 15 years, the use of rail infrastructure by different train operating companies (shared railway system) has been proposed as a way to improve infrastructure utilization and to increase efficiency in the railway industry. Shared use requires coordination between the infrastructure manager and multiple train operators in a competitive framework, so that regulators must design appropriate capacity pricing and allocation mechanisms. However, the resulting capacity utilization from a given mechanism in the railway industry cannot be known in the absence of operations. Therefore assessment of capacity requires the determination of the train timetable, which eliminates any potential conflicts in bids from the operators. Although there is a broad literature that proposes train timetabling methods for railway systems with single operators, there are few models for shared competitive railway systems. This paper proposes a train timetabling model for shared railway systems that explicitly considers network effects and the existence of multiple operators requesting to operate several types of trains traveling along different routes in the network. The model is formulated and solved both as a mixed integer linear programming (MILP) problem (using a**

commercial solver) and as a dynamic programming (DP) problem. We solve the DP formulation with a novel algorithm based on a linear programming (LP) approach to approximate dynamic programming (ADP) that can solve much larger problems than are computationally intractable with commercial MILP solvers. The model simulates the optimal decisions by an infrastructure manager for a shared railway system with respect to a given objective function and safety constraints. This model can be used to evaluate alternative capacity pricing and allocation mechanism. We demonstrate the method for one possible capacity pricing and allocation mechanism, and show how the competing demands and the decisions of the infrastructure manager under this mechanism impact the operations on a shared railway system for all stakeholders.

***Keywords*— Shared railway systems capacity planning, train timetabling problem, linear programming approach to approximate dynamic programming, adaptive learning, q-factors**

1. INTRODUCTION

In the last 15 years, several countries have promoted the use of shared railway systems that allow independent operators to access the infrastructure. This enables higher levels of infrastructure utilization (Gomez-Ibanez and de Rus, 2006). However, shared railway corridors require coordination between the infrastructure manager – typically the owner of the infrastructure – and multiple train operators (Gomez-Ibanez, 2003). This coordination involves determining which trains can access the infrastructure at each time (capacity allocation) and the access price they need to pay (capacity pricing).

According to (Roth, 2002) and (Vazquez, 2003), ideal capacity pricing and allocation mechanisms should be reproducible and transparent, easy to understand, and non-discriminatory (especially when the operators compete in the same market). Simple capacity pricing and allocation rules are used in other network industries (such as electric power or telecommunication). However, in the railway industry, the characteristics of the network and operations critically affect capacity and safety. According to (Krueger et al., 1999) railway capacity is defined as "a measure of the ability to move a specific amount of traffic over a defined rail with a given set of resources under a specific service plan." As a result, the implications for the railway system of the capacity pricing and allocation mechanism remain unclear, even for simple mechanisms used in other contexts. In other words, to understand the implications of capacity regulation in the rail industry, we have to recognize that infrastructure capacity utilization is endogenous: it depends on the capacity allocation and must be calculated by determining the train timetable. The train timetable determines the arrival and departure time at every station of all trains scheduled. The resulting timetable will specify how competing demands for the infrastructure were prioritized to meet the infrastructure manager's objectives and constraints.

This paper presents a train timetabling model for shared railway systems that would allow regulators and decision makers to simulate the behavior of the infrastructure manager. This model will explicitly consider network effects to analyze the interdependencies between operation and infrastructure and to determine the optimal use of infrastructure capacity (train timetable) under different capacity pricing and allocation mechanisms. The model is not intended to replicate the actual behavior of infrastructure managers, but rather provides a tool for determining the optimal capacity allocation plan given the operator's demand for capacity and the specific capacity pricing and allocation mechanism.

There are several papers that propose mixed integer linear programming (MILP) approaches to solve the train timetabling problem. Castillo et al. (2009), Ghoseiri et al. (2004), Liebchen et al. (2004), and Zhou and Zhong (2005) present formulations to compute the train arrival and departure times with different objectives and constraints. Traditionally, these models have been called multi-mode resource constrained project scheduling models. Other studies are based on MILP approaches that formulate the train timetabling problem as a multi-commodity flow problem (Caimi et al., 2009; Caimi et al., 2011; Caprara et al., 2002; Caprara et al., 2007; Courdeau et al., 1998). This approach represents the final timetable as a collection of nodes and arcs that represent possible train arrival and departure times at stations. Infrastructure and operational constraints are imposed by determining subsets of compatible and incompatible arcs.

With the exception of (Caprara et al., 2007), the approaches above assume that there is a single operator trying to schedule trains on the infrastructure. Most of the papers also assume that all trains follow the same path. This paper proposes a multi-mode resource constrained project schedule formulation considering different operators using a common infrastructure for different types of services (commuter, intercity and freight services) with different routes. These additional complexities result in 1) the need to specify safety constraints that depend on the path of each train, and 2) the need for operators to specify their desired timetable with flexibility margins to ensure that their trains get scheduled with minimal adjustments in the case of a conflict with another operator's request. These additional considerations make the problem more difficult to solve than the traditional single-operator timetable: simplifications used in other train timetabling problems (such as exogenously fixing train order or classical decomposition approaches) cannot be used because the safety constraints are specific to each individual train and because the need for large flexibility margins impacts the execution time of the problem.

From a computational standpoint, the size of the model increases exponentially with the number of stations and the number of trains to schedule. We propose an alternative class of solution algorithms using approximate dynamic programming (ADP) techniques (Bertsekas and Tsitsiklis, 1996; Bertsekas, 2006; Powell, 2007). This paper develops a novel Q-factor Adaptive Relaxed Linear Programming (QARLP) algorithm that extends previous algorithms developed by (Farias and Van Roy, 2003a; Farias and Van Roy, 2003b). This algorithm allows us to decompose and solve larger problems that are intractable with MILP solvers with less computation while still converging to a solution within an optimality gap.

This paper makes both methodological and railway systems-specific contributions. In terms of methodology, we present a model that explicitly considers the relevant characteristics of shared railway systems, and a novel ADP algorithm for solving this train timetable problem for large system sizes. In terms of transportation science, the model simulates optimal decisions by an infrastructure manager for shared railway systems, which can be used to evaluate alternative capacity pricing and allocation mechanism. We propose this modeling and algorithmic framework to allow regulators and decision makers to quantify these impacts for alternative mechanism designs.

The rest of the paper is structured as follows: Section 2 describes and formulates the train timetabling problem in shared railway systems, and motivates the assumptions of the paper. Section 3 presents a dynamic programming (DP) formulation of the problem, and describes the linear programming (LP)-based ADP solution algorithm. Section 4 compares the computational performance of the ADP algorithm with the performance of commercial MILP solvers, presents the timetables obtained for several scenarios with traffic patterns similar to the traffic of the Northeast Corridor in the U.S., and discusses the implications of the results for different stakeholders. Section 5 presents concluding remarks.

2. TRAIN TIMETABLING PROBLEM IN SHARED RAILWAY SYSTEMS

In this section, we formulate the train timetabling problem for shared railway systems under capacity pricing and allocation mechanisms. There are two main types of mechanisms for capacity pricing and allocation in infrastructure capacity markets for shared railway systems: 1) mechanisms that determine the price at which capacity will be offered, and let operators decide whether they are willing to access the infrastructure or not (price-based, cost-allocation mechanisms); and 2) mechanisms that determine the amount of capacity that will be offered, and let the operators reveal the price that they are willing to pay to use that capacity (quantity-based, auction mechanisms) (Gibson, 2003). Cost-allocation mechanisms are typically complemented with priority rules that allow the infrastructure manager to decide which train to schedule when there are conflicts (multiple operators willing to pay the pre-determined access charges).

The model presented here determines the optimal set of trains that the infrastructure manager can accommodate, assuming that an auction mechanism is implemented. Under an auction, at some predetermined frequency, the operators will have the opportunity to submit bids. Each bid will consist of a list of the trains that the operator wants to schedule on the infrastructure, the desired timetable for each train, and the access charges they are willing to pay to schedule each train. The infrastructure manager will then determine the set of trains that can actually be scheduled, their timetable, and the access charges that the operators will pay. We assume that the infrastructure manager's objective is to maximize revenue, subject to infrastructure constraints (e.g., safety, infrastructure maintenance plans).

We also discuss below how to modify the model to determine the optimal set of trains that the infrastructure manager can accommodate under alternative cost-allocation mechanisms. The differences between the infrastructure manager models for each mechanism affect mainly the definition of the parameters and the choice of the objective function. The constraints however are related to the physical and safe operation of the trains and remain unchanged across mechanisms. The model formulation is discussed below.

2.1. Sets

�, � ∈{1,...,�} trains proposed by the operators in the bidding process

�, � ∈{1, 2,..., �} railway system stations

2.2. Parameters

We use lower-case letters to denote parameters. The information that the train operators provide in the bidding process for every train � is:

□□□□□, a Boolean matrix that indicates the initial station □ from which train □ departs

□□□□□, a Boolean matrix that indicates the final destination (station □) of train □

□□, the maximum access price (access fee) that the operator is willing to pay if train □ is scheduled. For price-based mechanisms the access price will be pre-determined (using for example a cost-allocation model) and fixed by the infrastructure manager depending on the characteristics of the service. It is important to note that the operator will only operate a train if that price is less than or equal to its willingness to pay.

□□□□□□, □□□□□□, the desired arrival and departure time of train □ at every station □ in the path of train □

Δ□□□, □□Δ□□, maximum acceptable translation, defined as the maximum acceptable difference between the desired timetable and the actual timetable at the initial station (see Figure 1) of train □ and per-unit penalty imposed by the operator if the infrastructure manager translates the train over the desired timetable. The penalty specifies the reduction in the access price that the operator is willing to pay

Δ□□□, □□Δ□□, maximum acceptable change in train □ total travel time (see Figure 1) and per-unit penalty imposed by the operator if the infrastructure manager increases the travel time of train □ at any station over the desired timetable

The information about the topology of the line and the type of service is represented by the following two matrices:

□□□□□□, a Boolean matrix that indicates whether train □ travels through station □ or not

□□□□□□□′, a Boolean matrix that indicates for each train □ the station □′ that train □ will visit immediately after station □. Train □ may not stop at station □′.

In addition, the topology of the track and the signaling system will determine the minimum safe headway (time elapsed) between consecutive maneuvers at every station:

h□□□□, h□□□□ minimum headway between consecutive arrivals/departures to/from station □

In some cases the minimum safe headway depends also on the type of service and on the characteristics of the rolling stock. If that is the case, the former parameters will have different values for each train pair. The infrastructure manager can set larger minimum headway to ensure the reliability of the timetable (including timeslack to recover delays in the system).

2.3. Variables

We use capital letters for variables. The endogenous decision variables of this problem are:

□□ binary variable that indicates whether train □ is scheduled

□□□□□□,□□□□□□ final arrival and departure time (timetable) of every train □ scheduled at every station □ in the path of the train

Δ□□□, Δ□□□□ final train □ translation and increment of travel time per station □. Note that these variables can be determined knowing □□□□□□,□□□□□□ and vice versa. This paper assumes Δ□□□□≥*0* to ensure that the resulting train timetable is feasible. Δ□□□ can either be positive or negative; so we define the auxiliary positive variable Δ□□□+ as the absolute value of Δ□□□.

□□□′□ is a binary disjunctive variable with value *1* if train □ departs before train □′ at station □ and value *0* otherwise.

2.4. Objective Function

As discussed before, the objective of the problem is to determine which trains should be scheduled and when to maximize the infrastructure manager's revenue:

□□□□□□□□− □□Δ□□Δ□□□+−□□Δ□□□□□□□□□ (1)

Alternative objective functions could be defined for different capacity pricing and allocation mechanisms. For example, the functions:

□□□□□□ (2)

□□□□□□□□□ (3)

could be used to maximize the number of trains scheduled or the number of priority trains scheduled respectively under a cost-allocation and priority rules mechanisms. In this case □□□ would be a parameter that indicates the priority level of each train □. This priority level can, for example, be proportional to the number of passengers times the miles of the service.

2.5. Constraints

The first set of constraints establishes the relation between the desired timetable and the final timetable of every train scheduled:

The departure time from the first station can be determined as:

□□□□□□=□□□□□□+Δ□□□, ∀□,□:□□ □□□ (4)

The travel time between intermediate stations can be determined as:

□□□′□□□−□□□□□□=□□□′□□□−□□□□□□+Δ□□□□, ∀□,□, □′: □□□ □□□□′ □□□ □□□□□′+□□ □□□′=0 (5)

At the final station, the travel time can be determined using:

□□□′□□□−□□□□□□=□□□′□□□−□□□□□□+Δ□□□□, ∀□,□, □′: □□□ □□□□′ □□□ □□ □□□′ (6)

Note that the arrival time at the initial station is not defined in the timetable, nor is the departure time from the last station.

To ensure that the timetable is feasible, the final stopping and travel times at each station must be greater than or equal to the stopping and travel time in the desired timetable:

□□□□□□−□□□□□□≥□□□□□□−□□□□□□, ∀□,□:□□□ □□□ □□□ □□ □□□+□□ □□□=0 (7)

□□□′□□□−□□□□□□≥□□□′□□□−□□□□□□, ∀□,□:□□□ □□□′□ □□□ □□ □□□+□□ □□□=0 (8)

The maximum translation and increment of travel time for each train scheduled are also constrained. The allowable translation of a train is bounded by a maximum translation defined by the operator:

−Δ□□□≤Δ□□□≤Δ□□□, ∀□ (9)

In addition, the absolute value of the translation (Δ□□□+=Δ□□□) is determined using the following linear constraints:

Δ□□□+≥Δ□□□,Δ□□□+≥−Δ□□□ ∀□ (10)

The maximum change on travel time is bounded by the maximum increment on travel time specified by the operator:

[illegible] (11)

The operator may impose additional conditions within the bid to define the acceptable changes with respect to the desired timetable. That happens when the operator is not interested in operating the train if the departure from or the arrival at one major station is changed. In this case, additional constraints are included to ensure that the timetable respects the operator's requests if the train is scheduled.

The final set of constraints ensures that the timetable proposed by the infrastructure manager can be accommodated by the existing infrastructure. The infrastructure manager must ensure first that the difference between the departure times of every pair of trains scheduled is greater than or equal to the minimum safe headway, so at least one of the following equations must hold:

[illegible] (12)

[illegible] (13)

These conditions can be expressed using the following constraints:

[illegible] (14)

[illegible] (15)

In these equations [illegible] is a "big enough" number for the disjunctive formulation. In this formulation we use [illegible], which is the smallest possible [illegible] that can be chosen for this problem. The binary disjunctive variable [illegible] is used to automatically activate only one of the constraints depending on the value of the other variables. [illegible] has value *1* if train [illegible] departs before train [illegible] at station □. This problem will have on the order of □(□2□) binary variables and will be very difficult to solve for large □ (number of trains).

Similar constraints are included for inter-arrival times to ensure that the order of the trains is preserved between stations.

□□□+1□□□−□□′□+1□□□≥ℎ□+1□□□−□ □□□′□+2−□□−□□′, ∀□,□′,□:□< □′ □□□ □□□ □□□+1 □□□ □□□ □□′□+1 (16)

□□′□+1□□□−□□□+1□□□≥ℎ□+1□□□−□ 3−□□□′□−□□−□□′, ∀□,□′,□:□< □′ □□□ □□□ □□□+1 □□□ □□□ □□′□+1 (17)

For these constraints, a value of □=ℎ□+1□□□+□□′□ □□□−□□□□□□+Δ□□□′+Δ□□□+max□□□□,□□ □□′ is used.

Note that these constraints are valid for any shared railway system. The information about the topology of the infrastructure, the route of the trains, the safe headways imposed by the signaling system, etc. is introduced in the model parameterization.

3. LINEAR PROGRAMMING APPROACH FOR APPROXIMATE DYNAMIC PROGRAMMING

As discussed above, the size of the MILP model proposed in Section 2 increases exponentially as a function of the number of stations and trains to schedule. We propose a novel solution algorithm using ADP techniques (Bertsekas and Tsitsiklis, 1996; Bertsekas, 2006; Powell, 2007) to tractably solve large timetabling problems in shared railway systems.

Specifically, we propose a Q-factor Adaptive Relaxed Linear Programming (QARLP) algorithm that extends the Approximate Linear Programming (ALP) and the Relaxed Linear Programming (RLP) algorithms developed by (Farias and Van Roy, 2003a; Farias and Van Roy, 2003b). QARLP introduces two main innovations with respect to ALP and RLP algorithms: 1) it incorporates the possibility of learning from previous solutions, allowing the algorithm to improve the solution obtained by refining the sampling strategy in subsequent iterations, and 2) it formulates the Bellman equation using Q-factors, enabling the identification of promising areas in the solution space without needing to explicitly compute the value function at the next stage, which is very difficult for large multidimensional problems. This approach decreases the solution time, using fewer samples and iterations than a MILP commercial solver while still ensuring convergence to the optimal solution within a specified optimality gap.

3.1. Dynamic Programming Formulation

The problem defined in Section 2 can be reformulated as follows.

3.1.1. Stages

There are □=1,...,□ stages or decision periods (one for each train proposed to be scheduled).

3.1.2. State

The Markovian state variable is the timetable of the trains scheduled so far:

□□=□□□□□□□□□□−1 (18)

3.1.3. Control

At every stage, the control variable indicates whether the infrastructure manager decides to schedule train □ or not, and, if scheduled, the specific timetable of train □ at all stations □ in the path.

□□={□□,□□□□□□,□□□□□□} (19)

A train can only be scheduled if it does not present any conflict with the trains already scheduled.

3.1.4. State Transition Function

Given the state and the control at one decision stage, the state in the following decision stage can be computed, which incorporates the timetable of the new train if it is scheduled.

□□+1=□(□□,□□)=□□□□□□□□□□−1 □□ □□=0□□□□□□□□□□−1;□□□□□□,□□□□□□ □□

□□=1 (20)

3.1.5. Cost Function

The cost associated with a state-control pair is the sum of the penalties minus the revenue obtained if train □ is finally scheduled. The sign of the cost function has been chosen to formulate a minimization problem. The cost associated to each state-action pair is evaluated using:

□□□,□□=□ □□=−□□+Δ□□□□□Δ□□+Δ□□□□□Δ□□ · □□ (21)

3.1.6. Bellman Equation

The policy that minimizes the sum of current and future costs at every decision stage can be determined by solving the Bellman equation and calculating the cost-to-go or value function:

□□∗□□=min□□□□□+□□+1∗□□□,□□ □□□ □=1,...,□−1 (22)

□□∗□□=min□□□□□ (23)

This equation can be reformulated using Q-factors, which represent the cost-to-go for every feasible state-control pair:

□□∗□□, □□=□ □□+min□□+1□□+1∗□□□,□□,□□+1 □□□ □=1,...,□−1 (24)

□□∗□□,□□=□ □□ (25)

The relation between the cost-to-go function and the Q-factor is:

□□∗□□=min□□□□∗□□,□□ (26)

The optimal policy (timetable) can be determined solving the Bellman equation or the Q-factor Bellman equation using backward induction. However, when the dimension of the state space and the dimension of the control space increase, solution of the exact DP program becomes impracticable because the size of the problem grows exponentially. The benefit of reformulating the MILP model as a DP problem is that we can apply efficient solution algorithms such as the one proposed next.

3.2. Linear Programming Algorithm

Borkar (1988), De Ghellinck (1960), and Manne (1960) show that solving the Bellman equation (22) is equivalent to solving the LP problem proposed in equation (27) for any positive vector □ because the inequality □≤□∗ holds for every feasible solution □ of the problem. The vector □ is called the state-relevance weight vector

max□J, s.t. □□□□+□ J□+1□□□□,□□≥J□□□, ∀□□,□□ (27)

This LP problem has as many variables as possible states (value of the cost-to-go function at each state) and as many constraints as possible state-control pairs. When the state and control space of the problem are large, this results in a very large number of variables and constraints.

Schweitzer and Seidman (1985) and de Farias and Van Roy (2003a) proposed a modification of the previous formulation called the Approximate Linear Problem (ALP):

max□Φ□, s.t. □□□□+□ Φ□+1□□□□,□□□□+1≥Φ□□□□□, ∀(□□,□□) (28)

where the real value function □□∗ is approximated by a linear combination of basis functions □□□□=□=1,...,□ Φ□□□□□□□=Φi□□□□. In this approximation, there are only □ · □ variables (number of basis functions and number of stages). However, the number of constraints remains the same as in algorithm 1 (one constraint for each state-control pair).

To reduce the number of constraints in this problem, De Farias and Van Roy (2003b) proposed a Relaxed Linear Problem (RLP) formulation. RLP proposes a strategy which samples the constraints from the ALP formulation to include. De Farias and Van Roy showed that for an appropriate probability distribution function Ψ over the set of state-control pairs, the number of constraints that must be sampled does not depend on the number of state-control pairs. In particular, to obtain a solution close enough to the optimal solution obtained using the ALP formulation with 1−□ confidence level (Pr/□∗−Φ□□□□1,□ −□∗−Φ□□□□1,□ /<□ ≥1−□), the number of samples required is on the order of a polynomial in the number of state variables, 1/□ , and log1/□ . Note that these convergence results are computed over the basis-function approximation, that is, the RLP formulation converges to the best approximation over the basis functions chosen with confidence level 1−□ within a number of samples that does not depend on the number of state-action pairs. The RLP formulation is:

max□Φ□, s.t. □□□□+□ Φ□+1□□□□,□□□□+1≥Φ□□□□□, ∀□□,□□∈□ (29)

where □ is the set of state-action pairs sampled.

The main drawback of the RLP formulation presented in equation (29) is that the convergence results proved in (de Farias and Van Roy, 2003b) are based on an idealized choice of the probability distribution used to sample the constraints. In particular, the choice assumes knowledge of an optimal policy. Although it is unrealistic to assume that the optimal policy is known a priori, it is possible to obtain a reasonable approximation of the optimal policy by solving the RLP. Applying this idea, we propose the following Adaptive Relaxed Linear Programming (ARLP) algorithm:

Step 0: Set □=0, and sample □0, giving each state-control pair equal probability to be sampled (Ψ0 uniform distribution).

Step 1: Solve the problem max□Φ□,

s.t. □□□□+□ Φ□+1□□□□,□□□□+1≥Φ□□□□□, ∀□□, □□∈□□

Step 2: Set □=□+1 . Determine the optimal policy according to the last problem solved. Choose the next set of constraints sampled using a probability distribution function Ψ□ that assigns higher probabilities to those solutions in the optimal policy in the last iteration. In general, the variance of the probability distribution Ψ□ will decrease with □.

Step 3: It □>□0 or the difference between the objective function is smaller than □, stop. Otherwise go to Step 1.

Algorithm 1. Adaptive Relaxed Linear Programming (ARLP) algorithm

The ARLP algorithm iteratively solves a sequence of RLP problems, each with a manageable number of variables and constraints. This approach takes advantage of the reduced dimensionality of the RLP formulation while incorporating a mechanism to refine the sampling strategy Ψ□ using the best approximation of the optimal solution obtained so far. As a consequence, the convergence of the algorithm would not require the knowledge of the appropriate probability distribution function Ψ a priori.

However, because the basis function approximation used reduces the dimensionality of the problem, finding the state-control pair (timetable) that corresponds to known basis function values becomes a challenge. In other words, although it is very easy to determine the value of each basis function for a given state-control pair, solving the inverse problem (determining a state-control pair associated with a given basis function value) is extremely difficult in these cases. This is because the basis functions are a projection from the higher-order state-action space to a reduced dimensionality space, and the mapping from the low-dimensional projection back to the higher dimensional space is underdetermined. Therefore, there is no straightforward way to define Ψ□ based on low-cost regions in the basis function space, and to sample state-action pairs from it.

This problem can be solved by 1) reformulating the algorithm using Q-factors instead of the cost-to-go function, and 2) determining Ψ□ implicitly using a Metropolis Hasting algorithm (Rubinstein and Kroese, 2008) that accepts or rejects a state-control sample based on how promising the sample is using the Q-factor best guess.

To do that, we sample uniformly across the state-action space to determine a candidate state-action pair (□□′,□□′), we compute the value of its associated Q-factor (Φ□□□′,□□′□□), and compute the ratio of this Q-factor with the range of possible Q-factor values in the latest iteration (Φ□□□,□□□□,Φ□□□,□□□□):

ϕ=Φ□□□′,□□′□□−Φ□□□,□□□□Φ□□□,□□□□−Φ□□□,□□□□ (30)

We then draw a sample □ from a probability distribution Ξ□. We accept the state action pair if ϕ≥□. To ensure convergence to the optimal solution, the variance of the probability distribution Ξ□ must decrease with □, the probability of accepting any sample must be strictly positive, and the probability of accepting any sample with associated ϕ≥*1* must be 1. To ensure that these conditions hold, we determine the probability of accepting a proposed sample based on a sample drawn from Ξ□~□*(− 1*□*,1)*. In other words, we accept a sample with probability □=max*1*□,min*1*,ϕ. The performance of the algorithms improves when the state-control pairs associated with binding constraints in the previous iteration are retained in future iterations.

We call this algorithm a Q-factor Adaptive Relaxed Linear Problem (QARLP) algorithm.

Step 0: Set □*=0*, and sample □*0*, giving each state-control pair equal probability to be sampled (Ψ*0* uniform distribution).

Step 1: Solve the problem max□Φ□*,*

s.t. □□□□+□ Φ□+*1*□□□□,□□,□□+*1*□□+*1*≥Φ□□□,□□□□, ∀□□,□□,□□+*1*∈□□

Step 2: Set □=□*+1* . Determine the optimal policy according to the last problem solved. Repeatedly sample state-control pairs until one pair is accepted if the ratio of its associated Q-factor (Φ□□□,□□□□) and the range of possible Q-factor values in the latest iteration (□) is greater or equal than a draw of a probability distribution Ξ□ *(*□*)*.

Step 3: It □>□*0* or the difference between the objective function is smaller than □, stop. Otherwise save binding constraints from previous iterations and go to Step 1.

Algorithm 2. Q-factor Adaptive Relaxed Linear Programming (QARLP) algorithm

Note that in this algorithm □,Φ,□ have slightly different meanings than in Algorithm 1: □ is a positive constant for every state-control pair at every decision stage and Φ,□ are functions of the state-control pair (not only of the state: Φ,□= Φ,□ □(□□,□□)).

The choice of basis functions that capture relevant information about the state and the action while at the same time decreasing the amount redundant information (and hence the dimensionality of the problem) is a critical design choice of these types of ADP algorithms. In this paper, we use basis functions that capture: 1) the total number of trains scheduled, as well as the total changes in the operator's desired timetable (state variable), 2) whether train □ is scheduled or not, and the total changes in the desired timetable, 3) the number of conflicts of the trains scheduled so far with the following trains to be scheduled, and 4) and a constant. This reduces the dimensionality of the approximate cost-to-function to □=*7*. That is,

Φ□□□,□□=(□′<□ □□′,□′<□ Δ□□□,□′<□ Δ□□□,□□,Δ□□□,Δ□□□,□□,1) (31)

where □□ is the number of conflicts of the trains scheduled so far with the following trains: □□=□′>□ (□□′). The variable □□′ has value *1* if the desired timetable of train conflicts with the timetable of any train scheduled so far and *0* otherwise. We chose these basis functions after trying different options (with and without □□; with higher and lower level of aggregation of the timetable information) and picked these ones because the error of the approximation around the optimal solution is low and because they allow differentiating promising solutions from solutions that are not as promising.

4. Results

In this section, we present the results for the train timetabling problem from both the MILP and the ADP formulations. The first subsection presents the computational results of the paper, comparing the solution times between the commercial MILP solver and the QARLP algorithm. The second subsection presents the timetables obtained for several scenarios of demands from operators, and presents several insights for the context of a shared infrastructure and multiple agents. The third subsection summarizes the main implications for each stakeholder group.

4.1. Computational results

We begin by presenting the results obtained from solving the timetable problem for the representative infrastructure (Figure 2). It consists of a double-track corridor with 12 stations. The system presented includes the critical characteristics required to represent a corridor such as the Northeast Corridor in the U.S., for which the Federal Railroad Administration is currently developing a new capacity pricing and allocation mechanism to foster rail efficiency.

Stations 1, 2 and 12 represent main stations in the same metropolitan area (e.g., Boston), stations 3, 4, 5, 9, 10 and 11 are all in another metropolitan area (e.g., New York), and stations 6, 7 and 8 are in yet a third distinct metropolitan area (e.g., Washington D.C.). Five types of services are considered: Boston commuter trains travelling around the Boston metropolitan area (stations 1, 2, and 12); New York commuter trains; D.C. commuter trains; and intercity and freight trains travelling between Boston and Washington D.C. Intercity and freight trains may not stop at every station. Freight trains will travel the line at speeds much lower than commuter and intercity trains. Intercity trains travel at higher speeds than commuter trains.

At present, around 2,000 commuter trains, 150 intercity trains and 70 freight trains travel around the Northeast Corridor every day. In practice, most of the conflicts to schedule trains occur around peak hours; where the infrastructure manager would have to control for conflicts within sets of around 100-250 trains to make changes in the timetable.

We assume that the commuter operators (one in each city) request to schedule commuter trains in their respective urban areas every 30 minutes, and that one intercity operator requests to operate a train every hour. The number of trains requested by operators depends on the total time horizon considered.

The MILP formulation from Section 2 is implemented in GAMS 24.1.2 and solved using CPLEX 12.5 on a PC at 2.40 GHz, 4GB, intel core i7, under Microsoft Windows 7 64 bits. To reduce the size of the problem, when the desired arrival and departure times of two trains are very far apart, the value of the binary variable ⍰⍰⍰′⍰ is fixed a priori (since the relative order in which they pass through the station cannot change). We run CPLEX with options `CHEAT=0.05`, `⍰INSHeur=50`, `Threads=−1` for faster solution times, and use a 5% optimality gap. A smaller optimality gap may be required if the timetable problem has multiple quasi-optimal solutions with very different implications for different operators in terms of which trains are scheduled to ensure that the infrastructure manager choice of the trains to schedule is not arbitrary. In practice, for the cases solved for this paper, the difference in the objective function between scheduling or not one additional train is large. As a consequence any solution within a 5% optimality gap of the optimal solution ensures that the set of trains scheduled is the same than the set of trains scheduled in the optimal solution unless there are twin trains (operators willing to pay the same to operate trains with the exact same timetable). In that case neither CPLEX nor the QARLP algorithm would be able to distinguish those trains in the solution and the choice of one solution over other would be random.

We then solve the identical problem using the QARLP algorithm proposed in the previous section. Although theoretically the relative order of the trains does not change the solutions obtained or the convergence speed of the algorithm to the optimal solution, in practice the relative order of the trains may speed up the process of finding the optimal solution. The results presented in this paper correspond to cases in which the relative train order (trains considered at each stage) was randomly assigned.

Table 1 shows the number of equations, variables and discrete variables for problems with several different numbers of requested commuter and intercity trains. Figure 3 compares the execution time and the number of iterations required to convergence (within 5% integrality gap) of the MILP approach and the QARLP algorithm.

Note that the QARLP execution time increases as a polynomial function of the size of the problem (number of stations and number of trains to schedule). In contrast, the MILP solution times increase exponentially with the size of the problem. In particular, solving the MILP problem with CPLEX for 75, 90, or 120 trains within a 5% convergence gap is computationally intractable. Extrapolating from a regression estimate (Figure 3), the solution time using CPLEX for 120 trains would be approximately 46 days. The solutions obtained with the QARLP algorithm for 90 and 120 trains in approximately 20 minutes are better than those obtained for the MILP formulation with CPLEX after 20 hours and 35 hours respectively. In the cases with 15, 30, and 60 trains the solutions obtained with both methods are almost identical.

4.2. Design of Timetables for Corridors with Traffic Patterns similar to the Northeast Corridor of the U.S.'s Traffic

Figures 4, 5 and 6 show time-space diagrams for timetables designed by the infrastructure manager model for different capacity use demand scenarios. The y-axes represent distance in miles from station 1 and the x-axes represent time in minutes at which different trains are scheduled to pass through each point of the line (vs. desired scheduled in dashed line). The horizontal segments represent the stopping times at stations. We assume no interaction between trains travelling in different directions.

Figure 4 shows the timetable for a scenario with demand to schedule an intercity train in the system when commuter trains around the three urban areas operate every 30 minutes. Figure 5 shows a scenario in which two competing intercity operators request scheduling intercity trains when commuter trains around the three urban areas operate every 30 minutes. Figure 6 shows a scenario with demand to schedule a freight train in the system when commuter trains around the three urban areas operate every 1 hour. The infrastructure manager model proposes the final timetable analyzing the trade-off between eliminating trains and readjusting the desired schedules, according to the objective function in (1). For clarity purposes, only the schedules of conflicting trains are shown in the figures.

For this example, we assume that each commuter train operator pays 1 unit to schedule a commuter service and gets a 5% discount from the original access fee for every minute that one of their trains is rescheduled. When the infrastructure manager tries to schedule an intercity train, it will initially conflict with 14 commuter trains (see Figure 4). Rescheduling the commuter trains to accommodate the intercity service requires that the commuter operators receive a discount of 2.1 units on their total access fees. As a result, the infrastructure manager would only schedule the intercity train if it represents more than 2.1 units of revenue.

If the frequency of commuter trains increases, for example to one commuter train every 15 minutes instead of every 30 minutes, the intercity train will initially conflict with 22 commuter trains and will only be scheduled if it represents more than 3.6 units of revenue for the infrastructure manager. Conversely, if the frequency of commuter trains decreases to one train every 60 minutes, the intercity train will be scheduled if it represents at least 1.5 units of revenue for the infrastructure manager. The model can be used to quantify the trade-off between commuter and intercity trains for any other frequency of service (see Figure 7). The exact value of the trade-off for low frequencies of commuter services depends on whether there are conflicts among the desired timetables of the trains or not.

Scheduling two intercity trains in Figure 5, assuming the same willingness-to-pay to schedule commuter trains, results in a 4.0 units of total discount for commuter operators (the intercity service will initially conflict with 14 commuter trains). That means that the infrastructure manager would only schedule the two intercity trains if they represent more than 4.0 units of revenue. If the revenue from scheduling the intercity trains represent between 2.1 and 4.0 units, at most one of the intercity trains would be scheduled.

Furthermore, note that although both trains would like to depart station 1 at minute 0, one of them will depart at minute 3 and the other one at minute 8. In some cases, none of the operators may be interested in operating a second intercity service just 5 minutes after another one. In this example, intercity operators may avoid getting their train scheduled just after other intercity by controlling 1) how flexible their schedule is, 2) how much discount in the access fee they obtain if the schedule of the train is changed, and 3) how much they are willing to pay to access the infrastructure.

Figure 6 assumes that the willingness to pay to access the infrastructure for a freight operator is lower than the willingness to pay of other operators. We also assume that the freight operator's flexibility (total allowed translation and increment of travel time) is higher than other operators' flexibility.

For the same commuter frequency, a freight train will initially conflict with more commuter trains than an intercity train since it travels at a lower speed. As a result, independent of how much each commuter service operator pays, the freight train will be scheduled as long as the net access fee paid to the infrastructure manager is positive. The minimum access fee that a freight operator must pay when the line is more congested will depend on how many trains have to be rescheduled to eliminate conflicts. If the commuter operator wants to increase the frequency of commuter service from one train per hour to one train every 30 minutes, the freight train will only be scheduled if the net access fee that the freight operator is willing to pay represents more than 3 units of revenue for the infrastructure manager (since three commuter services could not be operated). In general, the relative speeds among different types of services have a major impact on the capacity utilization of the system.

4.3. Implications for Stakeholders

Here we present the main implications of these results for the operators, the infrastructure manager, the users, and the regulators.

4.3.1. The capacity pricing and allocation mechanism determines how different operators are able to compete to access capacity

This effect is important when there are different types of operators (such as commuters and long-distance operators) competing to access the infrastructure. Although the results presented above were calculated assuming that the capacity pricing and allocation mechanism in place is an auction, similar effects occur when other mechanisms are implemented. For instance, the results of the model demonstrate that intercity and freight services would have a disadvantage to compete with commuter trains under capacity allocation methods that maximize how many trains are scheduled in the system. This problem could be solved by designing mechanisms that account for the miles or passenger miles, by assigning higher priority to long distance services, or by assigning a societal utility to each service. The impact of these measures could be quantified using the model proposed for the specific system.

Competing companies may be concerned about getting their train in the schedule just after another competing train. The results for this problem suggest that when the infrastructure is congested, scheduling two competitive trains for the same service at similar times represents a high opportunity cost for the infrastructure manager. Operators may be able to avoid having their train scheduled at the same time as another competing train by controlling the flexibility of their desired timetable and their willingness to pay to access the infrastructure.

4.3.2. The characteristics of the system and of the capacity pricing and allocation mechanism affect the ability of the infrastructure manager to recover infrastructure cost

The results of the model also show that depending on the operators' infrastructure access demand (e.g., frequency of commuter services), the price that an intercity or freight operator will have to bid to be able to schedule a train can vary considerably. This price reflects the congestion rent. The infrastructure manager's ability to recover costs using an auction mechanism, therefore, depends on the level of service (defined by the number of train services offered and the mix between different types of train services) of the line. In congested infrastructure, greater cost recovery is expected.

Although cost-allocation mechanisms are usually presented as capacity pricing and allocation mechanisms that allow the infrastructure manager to recover infrastructure costs, the operators' demand to access the infrastructure will certainly depend on the infrastructure access fee. For instance, the operators would require high average ridership levels to provide a service when they pay high access charges and vice-versa. By this relation the infrastructure access fee is related to the consumer utility not necessary related to the congestion rent of the infrastructure manager.

4.3.3. The characteristics of the system and of the capacity pricing and allocation mechanism also condition the level of service experienced by the users of the railway system (passengers and freight shippers)

This implication is a consequence of the previous ones. The capacity pricing and allocation mechanism determines how different operators can compete to get access to the infrastructure and the access charges they will have to pay. This will ultimately affect the operators' infrastructure access demand, which translates into how many services each operator will schedule. This will have a direct impact on the level of service.

Furthermore, the equilibrium between different types of services will also impact the level of service that the users will experience: whether it is possible to have frequent intercity services, whether it is possible to operate higher-speed services in the system, which long-distance services can be operated during peak-hours, etc.

4.3.4. Regulators' design of capacity pricing and allocation mechanisms should consider not only overarching regulation goals but also the specific characteristics of the railway system

This paper shows that the implications of the capacity pricing and allocation mechanism for the operators, the infrastructure manager, and the users strongly depend on the characteristics of the system. As a consequence, regulators should consider these characteristics to analyze whether overarching regulation goals such as mix of services, level of infrastructure cost recovered, and level of service are met.

5. CONCLUSIONS

This paper proposes a train timetabling model for shared railway system. The model is formulated as both a MILP problem and as a DP problem. The MILP is solved using commercial software and the DP is solved using a novel algorithm for ADP. The timetables designed with the model are used to evaluate how capacity pricing and allocation may impact different railway system stakeholders. As a result, the contributions of this paper are both methodological and domain specific.

On the methodological side, the main contributions of the paper include:

1) The formulation of a train timetabling model for shared railway systems that would allow regulators and decision makers to simulate the behavior of the infrastructure manager. This model explicitly considers network effects to analyze the interdependencies between operation and infrastructure and the existence of different operators willing to operate trains with different paths within the infrastructure under various capacity pricing and allocation mechanisms.

2) The development of a novel algorithm for rapidly solving the train timetable problem in shared railway systems, which requires fewer samples and iterations and ensures convergence to the optimal solution within a specified optimality gap. We obtain solutions within 5% of the optimal solution for problem sizes that cannot be solved within a 5% convergence gap using commercial MILP software.

3) The algorithm developed, a Q-factor Adaptive Relaxed Linear Programming (QARLP) algorithm, extends the Approximate Linear Programming (ALP) and the Relaxed Linear Programming (RLP) algorithms developed by (Farias and Van Roy, 2003a; Farias and Van Roy, 2003b). QARLP introduces two main innovations with respect to ALP and RLP algorithms: 1) it incorporates the possibility of learning from previous solutions, allowing the algorithm to improve the solutions obtained by refining the sampling strategy in the next iterations, and 2) it formulates the Bellman equation using Q-factors, allowing the identification of promising areas in the solution space without computing the value function at the next state (very difficult in large multidimensional problems). These ideas can be generalized for other instances.

Moreover, the results of the train timetabling model can be used to simulate and evaluate the best possible behavior of the infrastructure manager in shared railway systems under different capacity pricing and allocation mechanisms. The domain-specific contributions of this paper are:

4) This paper shows that the capacity pricing and allocation mechanisms critically impact the operations on a shared railway system for all the stakeholders: it determines whether the infrastructure manager will be able to recover costs; it determines how operators offering different or competing services will be able to compete to access the infrastructure; and hence it ultimately determines the level of service that will be offered to rail users.

5) This paper also shows that the extent of this impact depends on the characteristics of the system and the traffic level. We propose the use of this model as a tool to allow regulators and decision makers to quantify the impact of different capacity pricing and allocation mechanisms.

This paper considers the operators' infrastructure access demand (characterized both as the demand to schedule trains and the revealed willingness to pay to access the infrastructure) as exogenous to the problem. However, the operator's infrastructure access demand depends on the capacity pricing and allocation mechanism. Future research will integrate the infrastructure manager model proposed with models of operator bidding behavior to better quantify the trade-offs between utilization and level of service on the one hand, and infrastructure cost recovered under different capacity pricing and allocation mechanisms. These results will be valuable to design and evaluate appropriate capacity pricing and allocation mechanisms aimed at the particular characteristics of a specific shared railway system.

ACKNOWLEDGEMENT

The authors gratefully acknowledge support by the US Department of Energy Office of Science, Biological and Environmental Research Program, Integrated Assessment Research Program, Grant No. DE-SC0003906 and by the Rafael del Pino Foundation.

REFERENCES

Bertsekas D., (2006). *Dynamic Programming and Stochastic Control*, Ed. Athena Scientific, Belmont, MA.

Bertsekas D., J.N. Tsitsiklis, (1996). *Neuro Dynamic Programming*, Ed. Athena Scientific, Belmont, MA.

Borkar V., (1988). A convex analytic approach to Markov decision processes. *Probability Theory Related Fields, Vol. 78, pp. 583-602.*

Caimi, G., D. Burkolter, T. Herrmann, F. Chudak, M. Laumanns, (2009). Design of a railway scheduling model for dense services. *Networks and Spatial Economics, 9, 25-46.*

Caimi, G., F. Chudak, M. Fuchsberger, M. Laumanns, R. Zenklusen, (2011). A new resource-constrained multicommodity flow model for conflict-free train routing and scheduling. *Transportation Science, 45(2), 212-227.*

Caprara, A., M. Fischetti, P. Toth, (2002). Modeling and solving the train timetabling problem. *Operations Research, 50(5), 851-861.*

Caprara, A., L. Kroon, M. Monaci, M. Peeters, P. Toth, (2007). Passenger railway optimization. In Barnhart, C. and Laporte, G. (Ed.), *Handbooks in operations research and management science* (pp. 129-187) Elsevier.

Castillo, E., I. Gallego, J. Ureña, J. Coronado, (2009). Timetabling optimization of a single railway track line with sensitivity analysis. *TOP, 17(2), 256-287.*

Cordeau, J.F., P. Toth, D. Vigo, (1998). A survey of optimization models for train routing and scheduling. *Transportation Science, 32(4), 380-404.*

De Ghellinck G., (1960). Les problemes de decisions sequentielles. *Cahiers du Centre d'Etudes de Recherche Operationnelle, Vol. 2, pp. 161-179.*

Farias D.P., B. Van Roy, (2003a). The linear programming approach to Approximate Dynamic Programming. *Operations Research, Vol. 51, No. 6, pp. 850-865.*

Farias D.P., B. Van Roy, (2003b). On constraint sampling in the Linear Programming approach to Approximate Dynamic Programming, IEEE Conference on Decision and Control.

Ghoseiri, K., F. Szidarovszky, M.J. Asgharpour, (2004). A multi-objective train scheduling model and solution. *Transportation Research, Part B: Methodological, 38(10), 927-952.*

Gibson S., (2003). Allocation of capacity in the railway industry. *Utilities Policy, Vol. 11, pp 39-42.*

Gomez-Ibanez J.A., (2003). *Regulating infrastructure: monopoly, contracts, and discretion.* Harvard University Press, Cambridge, MA, USA.

Gomez-Ibanez J.A., G. de Rus, (2006). *Competition in the railway industry: an international comparative analysis*. Edward Elgar Editors, Northampton, MA, USA.

Krueger, H., Farrington, P. A., Black Nembhard, H., Sturrock, D. T., Evans, G.W., (1999). Parametric modeling in rail capacity planning. WSC'99. 1999 Winter Simulation Conference Proceedings. `Simulation - A Bridge to the Future', 2.

Liebchen, C., M. Proksch, F.H. Wagner, (2004). Performance of algorithms for periodic timetable optimization. San Diego, California, USA.

Manne A.S., (1960). Linear programming and sequential decisions. *Management Science, Vol. 6, No. 3, pp. 259-267.*

Powell W., (2007). *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, Ed. Willey, NJ.

Roth A.E., (2002). The economist as engineer: Game theory, experimentation, and computation as tools for design economics*. Econometrica, Vol. 70, No. 4, pp. 1341-1378.*

Rubinstein, R.Y., and D.P. Kroese (2008). *Simulation and the Monte Carlo Method*, Second Edition. John Wiley & Sons, Hoboken, N.J., chapter 6.

Vazquez C., (2003). Modelos de casacion de ofertas en mercados electricos. Ph.D. Thesis, Comillas Pontifical University (in Spanish).

Zhou, X., M. Zhong, (2005). Bicriteria train scheduling for high-speed passenger railroad planning applications. *European Journal of Operational Research, 167(3), 752-771.*

LIST OF TABLES

Number of Trains	Equations	Continuous/Total Variables	Discrete Variables
15	970	510	91
30	3,715	1,607	292
60	14,533	5,565	919
120	57,481	20,537	3,145

Table 1. Train timetabling problem size for traffic patterns representative of the Northeast Corridor traffic

LIST OF FIGURES

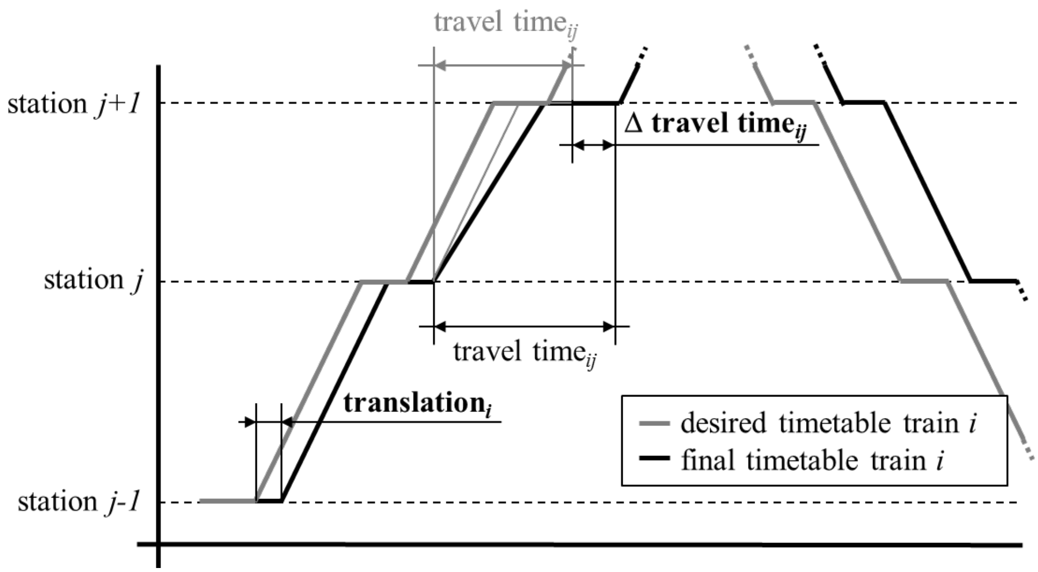

Figure 1. Time-space diagram representation of possible changes with respect to desired timetable

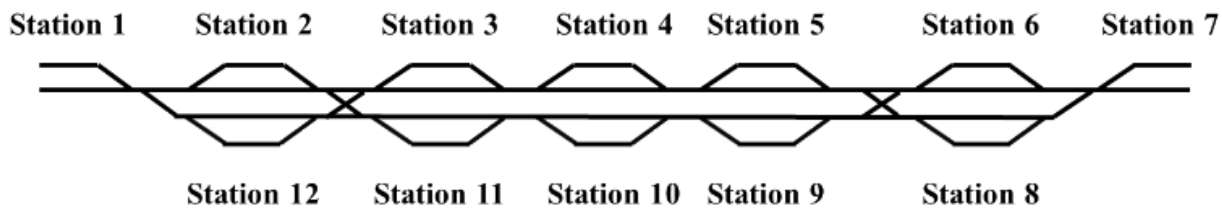

Figure 2. Detailed corridor infrastructure

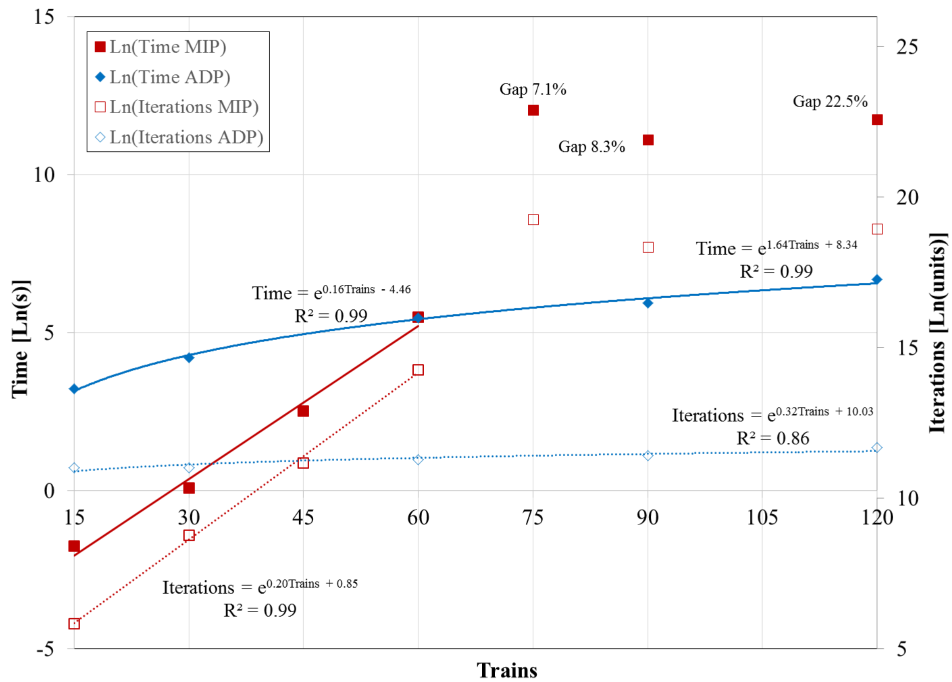

Figure 3. Comparison of the execution time and the iterations required to convergence (within 5% integrality gap) of the MILP approach and the QARLP algorithm.

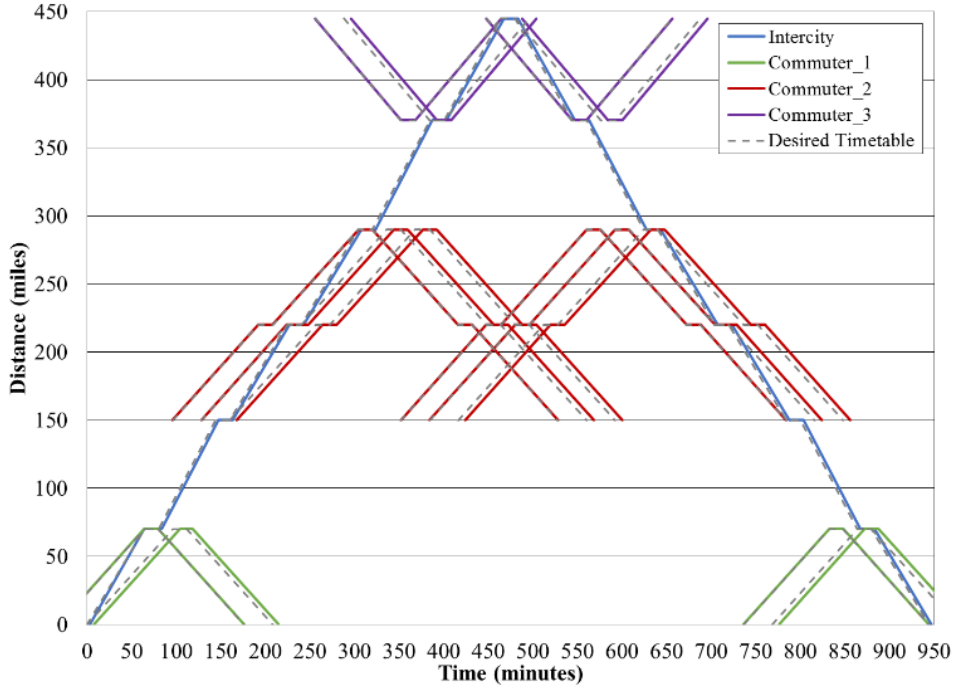

Figure 4. Timetable proposed by infrastructure manager to schedule an intercity train in a system with commuter trains operating every 30 minutes.

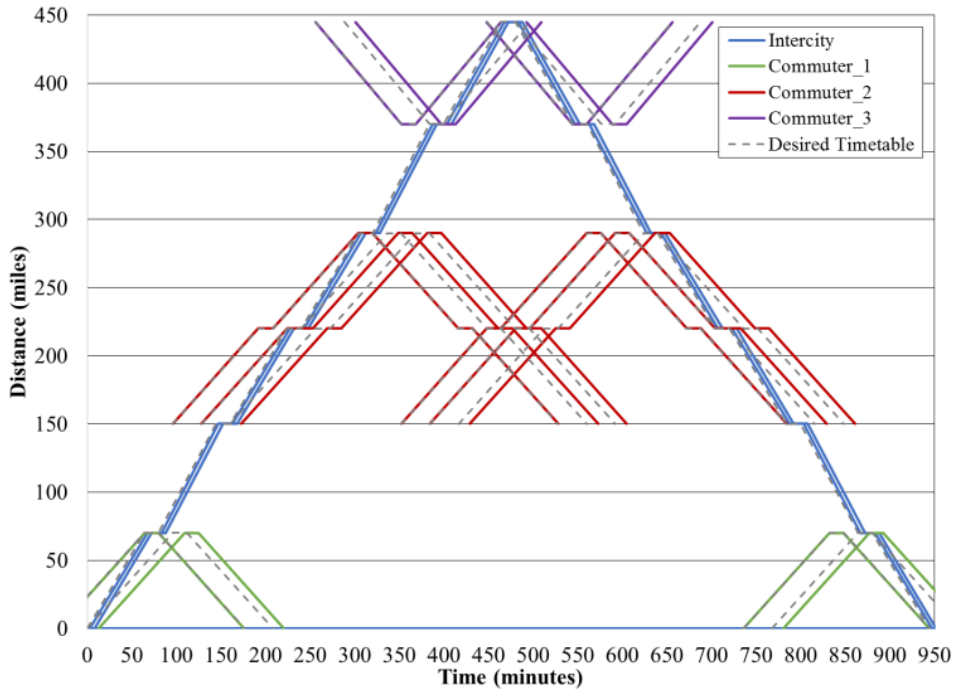

Figure 5. Timetable proposed by infrastructure manager to schedule two intercity trains in a system with commuter trains operating every 30 minutes.

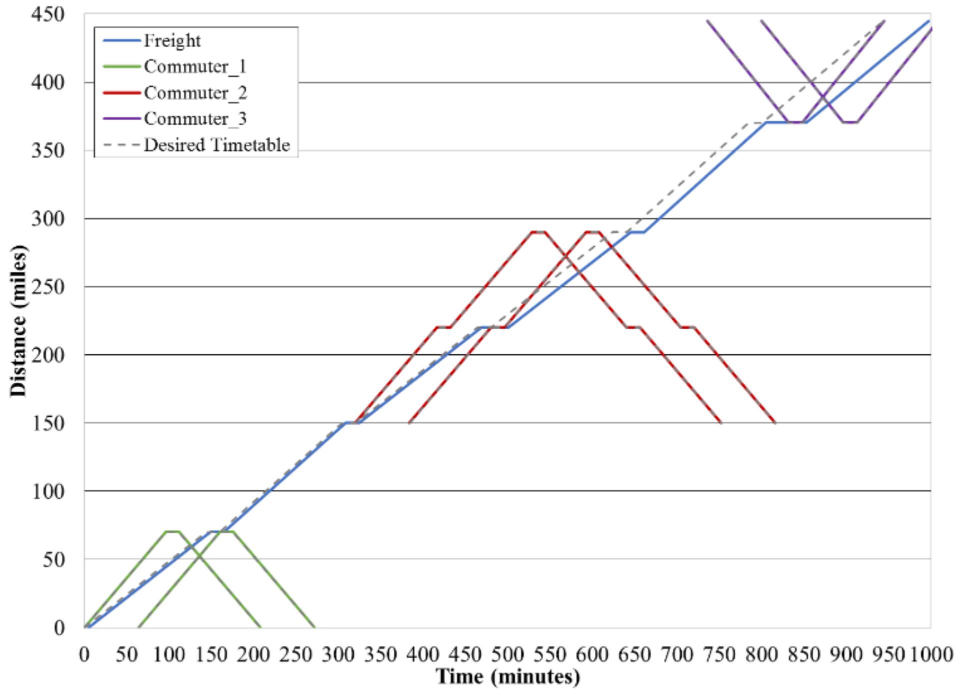

Figure 6. Timetable proposed by infrastructure manager to schedule a freight train in a system with commuter trains operating every 60 minutes.

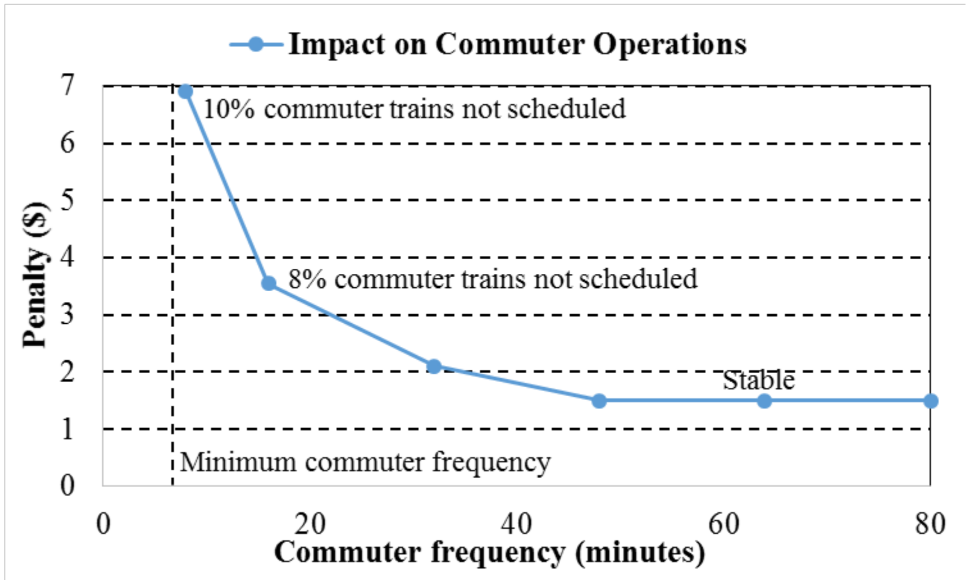

Figure 7. Intercity train schedule impact on commuter traffic as a function of the commuter frequency