

A Data-Driven Linear Programming Model for Energy-Optimal Metro Timetables

Qiuyu Fan

Department of Computational Applied Mathematics and Operations Research,
Rice University, Houston, TX, USA
qf3@rice.edu

J. Kevin Tobin

Hitachi Rail Canada, Toronto, ON, Canada
kevin.tobin@hitachirail.com

Bart P.G. Van Parys

Centrum Wiskunde & Informatica (CWI), Amsterdam, The Netherlands
bart.van.parys@cwi.nl

Shuvomoy Das Gupta

Department of Computational Applied Mathematics and Operations Research,
Rice University, Houston, TX, USA
sd158@rice.edu

Abstract

We propose a data-driven linear programming model to compute energy-optimal timetables for networks operating under communications-based train control (CBTC); the model was developed in collaboration with Hitachi Rail Canada, the largest provider of CBTC systems worldwide. Our model minimizes the network’s *effective energy consumption*, defined as the total traction energy consumed by the trains minus the regenerative braking energy transferred from braking trains to nearby accelerating trains. Our main modeling contribution is to show that the signed overlap between a braking and an accelerating phase is concave in the event times, so maximizing the transferred energy is a convex problem that can be reformulated into a linear program via hypograph constraints. In contrast, prior work either captures this synchronization structure with binary variables or other nonconvex constraints, or relies on multiple stages interleaved with simulation. Our model computes an energy-optimal timetable subject to the operational constraints of the railway network in under a second on a laptop for every instance studied and predicts the effective energy consumption of the network without requiring time-consuming simulations. We demonstrate the effectiveness of our model on two real-world CBTC networks. First, on 11 full-day operational instances of Shanghai Metro Line 8 with 1,000–1,332 active trains, the model computes energy-optimal timetables in under 0.7 seconds per instance and predicts 19.5%–28.2% reductions in effective energy consumption relative to the existing operational timetables; the industrial physics-based simulator SPSIM independently corroborates these reductions. Second, on 17 real-world instances of the Docklands Light Railway with 108–336 active trains, the model predicts reductions of 19.3%–28.8%, solving each instance in under 0.11 seconds. Our model is being integrated into Hitachi Rail Canada’s industrial timetable

compiler.

Keywords. Metro railway networks, Communications-based train control, Train scheduling, Energy efficiency, Sustainability

1 Introduction

Background. The United Nations projects that the share of the world’s population living in urban areas, based on national definitions, will rise from approximately 58% in 2025 to 67% in 2050 (United Nations, 2025), placing sustained pressure on existing transportation infrastructure. Rail transit, with its high passenger-carrying capacity and strong safety record (Abril et al., 2008), is well positioned to absorb much of this growth. Communications-based train control (CBTC) systems play a central role in the operation of modern metro networks (Pascoe and Eichorn, 2009). Used by some of the world’s busiest metros, they provide high-resolution train location determination and continuous data communication, which improve both safety and efficiency and allow operational schedules to be adjusted dynamically in response to congestion and fluctuating demand. Because CBTC systems can implement automatic driving strategies and execute a prescribed timetable precisely (IRSE International Technical Committee, 2009), they are particularly well suited to energy-efficient operation. These capabilities make fast and energy-efficient timetable optimization a natural way to improve both the efficiency and the sustainability of CBTC-enabled metro networks.

Motivation. In CBTC-enabled metro systems, trains are electrically powered. During a trip from one platform to the next, the train follows the four-phase speed pattern shown in Figure 1: acceleration, speed holding, coasting, and braking (Howlett and Pudney, 1995). Most of the energy consumption occurs in acceleration phase. During braking, the traction motors operate as generators and produce *regenerative braking energy*. Because CBTC systems can implement a given railway timetable with high degree of accuracy to the trains operating in the railway network, this energy can be transferred from braking trains to nearby trains that are accelerating. For perspective, the New York City transit railway system consumes more than 1,600 GWh of electricity per year, and all trains installed since 2018 are capable of producing and transferring regenerative braking energy; under favorable conditions, the recovered energy may amount to as much as 50% of the consumed energy (Mohamed et al., 2018).

Energy-optimal timetable. An energy-optimal timetable offers substantial energy savings without requiring infrastructural alterations to a CBTC-enabled railway system. Formally, a railway timetable is a data structure containing the arrival and departure times of each train at every platform visited during a designated service period, typically spanning 18 to 24 hours in most operational networks. An energy-optimal timetable chooses these temporal decision variables to minimize the network’s *effective energy consumption*: the *total energy consumption of all trains during their accelerating phases* minus the *total regenerative braking energy successfully transferred to the accelerating trains from eligible braking counterparts*.

Related work. Energy-efficient railway operation is studied at two levels: energy-efficient train control, which shapes a single train’s speed profile between two stops, and energy-efficient timetabling, which coordinates the event times of many trains. Scheepmaker et al. (2017) review both levels and their interaction, Yang et al. (2016) survey the urban-rail line of work, and González-Gil et al.




Figure 1: Illustration of a train's energy-use cycle under CBTC operation. For a trip between two platforms, the qualitative speed profile consists of acceleration, speed holding, coasting, and braking. The lower plot shows the corresponding power profile: the acceleration interval (red) accounts for most traction energy demand, whereas the braking interval (green) produces regenerative energy. With suitable scheduling, that recovered energy can be supplied to nearby accelerating trains.

(2013) review the technologies for recovering regenerative braking energy. One of the technologies is the train-to-train transfer exploited in this paper, which requires no dedicated storage hardware.

Within the timetabling level, work on regenerative-energy utilization spans several modeling paradigms. Peña-Alcaraz et al. (2012) study Madrid Metro Line 3 using a mixed-integer programming formulation tailored to a single line. Yang et al. (2013) initiated a cooperative-scheduling line that maximizes the overlap time between the braking and accelerating phases of successive trains (the objective that our energy-weighted overlaps refine), and Wang et al. (2022) design a mixed-integer linear programming model to increase the synchronization between accelerating and braking trains at the same station; Su et al. (2013) couple the timetable with energy-efficient driving regimes. Li and Lo (2014a) use a genetic algorithm to search for timetables that balance regenerative-energy utilization with train tractive energy; Le et al. (2015) and Liu et al. (2018) develop a nonlinear integer program solved by simulated annealing and by a hybrid algorithm, respectively. Das Gupta et al. (2015) give a more computationally tractable mixed-integer program for coordinating regenerative-energy exchange among suitable train pairs and demonstrate it on the Docklands Light Railway, but their model does not optimize actual energy savings directly. Yin et al. (2017) embed regenerative-energy utilization and dynamic passenger demand in mixed-integer linear programs solved by Lagrangian relaxation, and Canca and Zarzo (2017) design energy-efficient timetables for two-way rapid transit lines. The bi-objective model of Yang et al. (2020) combines timetable optimization with energy allocation and passenger assignment through a non-dominated sorting genetic algorithm (see also Yang et al., 2021); Gemander et al. (2023) optimize underground timetables under uncertain dwell and running times through stochastic programming; and Bärman et al. (2021) design exact Benders decomposition schemes to optimize train timetables for peak power reduction and regenerative-energy use improvement. A common pattern in these mixed-integer formulations above is a set of binary variables encoding whether the braking phase of one train overlaps the accelerating phase of another; in Section 4.2 we show that this overlap can be modeled as a concave function of the event times, so no integer variables are actually needed.

A complementary line of work combines timetabling with train control. The model in Li and Lo (2014b) jointly optimize dynamic train scheduling and speed control, Su et al. (2019) and Zhang et al. (2024) both use a two-level optimization approach: speed profile is optimized at the first level, while the optimal timetable is solved given the optimal speed profile at the second level. Luan et al. (2018) extend integrated real-time traffic management to energy-efficient train operations, and Wang and Goverde (2017) address multi-train trajectory optimization through a multiple-phase optimal-control formulation solved by a pseudospectral method, later adapting it to retime a fixed timetable for improved energy efficiency (Wang and Goverde, 2019). Shi et al. (2026) develop a decomposition approach to solve the integrated mixed-integer linear programming model for train rescheduling and speed control during disruption, and Howlett et al. (2023) formulate an unconstrained convex model to solve for optimal timetable and driving strategies for each train in a fleet. In contrast, we focus on building our optimization model based on how CBTC system works in practice: the on-board controller computes the energy-optimal speed profile for whatever trip times are assigned, and the optimization chooses only the event times, which is what keeps the network-wide model a single linear program.

From the standpoint of deployed industrial tools, the two-stage framework of Das Gupta et al. (2016) appears to be, to our knowledge, the only earlier model embedded in a CBTC timetable compiler. That work computes an energy-efficient timetable through two consecutive linear programs: the first stage minimizes network-wide energy use over all trains subject to the railway constraints, and the second stage fixes the resulting trip times and then maximizes regenerative-energy transfer

between suitable train pairs. However, this model in Das Gupta et al. (2016) has some limitations that limit its real-world applicability. First, although each stage’s linear program solves quickly, the handoff between the stages requires an intermediate simulation step that can take hours, so the total runtime is much longer. Second, the two-stage objective does not model actual network energy consumption directly; instead, it optimizes proxy quantities whose practical effect on energy usage must be assessed afterward by extensive physics-based simulation, which limits managerial usefulness when explicit energy predictions are desired. Third, the second stage relies on the exact locations of power-consumption peaks, information that is often unavailable in practice. These limitations in Das Gupta et al. (2016) have motivated us to design the new optimization model for CBTC systems that we describe in this paper.

Relative to that model, the present one keeps the operational constraint families of Section 3 and the dwell-midpoint pairing of braking with accelerating trains, and redesigns the optimization itself: transferred regenerative energy is now fitted directly in energy units as an affine function of phase overlap, the effective phases vary affinely with trip time, the overlaps are captured exactly through linear hypograph constraints in place of the earlier ℓ_1 alignment heuristic, and a single linear program minimizes (and predicts, without physics-based simulation) the network’s effective energy consumption.

Contribution. We develop a single-stage linear optimization model that constructs energy-optimal timetables for all trains in a CBTC-enabled metro railway network over an entire service period. Our contributions are as follows.

- *A single-stage linear model.* The model minimizes effective energy consumption (traction energy minus transferred regenerative energy) in one linear program, in contrast with earlier approaches that rely on NP-hard mixed-integer formulations or on multiple stages interleaved with time-consuming simulation.
- *Data-driven, simulation-free energy prediction.* Traction energy, transferred regenerative energy, and their interaction are modeled from operational or simulated data, so the model predicts the effective energy consumption of a timetable without invoking a physics-based simulator. This makes the model suitable for managerial settings in which explicit energy predictions are needed quickly.
- *A convex-analytic model of regenerative-energy transfer.* We show that the signed overlap between the effective braking and accelerating phases of a train pair is concave in the timetable, being a pointwise minimum of affine functions of the event times. Maximizing the transferred regenerative energy is therefore a convex problem, which a hypograph reformulation expresses through linear constraints without loss of optimality. Prior work either captures this synchronization structure with binary variables or other nonconvex constraints, or relies on multiple stages interleaved with expensive simulation.
- *Validation on two real-world networks and industrial integration.* On 11 full-day operational instances of Shanghai Metro Line 8 (one of the busiest metro services in the world) with 1,000–1,332 active trains, the model solves in under 0.7 seconds per instance on a laptop and predicts effective-energy reductions of 19.52%–28.17% relative to the existing operational timetables; these predictions are corroborated by Hitachi Rail’s physics-based simulator SPSIM. On 17 instances of the Docklands Light Railway benchmark with 108–336 active trains, the model predicts reductions of 19.33%–28.82%. The model is being integrated into Hitachi Rail Canada’s industrial timetable compiler.

This work arose from a collaboration with Hitachi Rail Canada, the largest provider of CBTC

systems worldwide. The collaboration began in 2023 with Thales Ground Transportation Systems (GTS). On May 31, 2024, Hitachi Rail completed its acquisition of Thales’s worldwide GTS business for €1.66 billion (Hitachi Rail, 2024). Due to this transaction, the established Canadian GTS organization became Hitachi Rail Canada, and our collaboration continued with the Toronto-based team under Hitachi Rail Canada. We therefore use the current name *Hitachi Rail Canada* throughout this paper.

1.1 Organization

The remainder of the paper is structured as follows. Section 2 introduces notation and preliminary concepts. Section 3 presents the operational constraints required for a metro railway network. In Section 4, we develop the effective-energy-consumption objective. Section 5 assembles the full optimization model. Section 6 reports experiments on Shanghai Metro Line 8 and the Docklands Light Railway benchmark, and Section 7 describes the architecture used for industrial deployment of the model.

2 Notation and preliminaries

Unless stated otherwise, every set appearing in this paper is finite and equipped with a strict order. For such a set S , we write $|S|$ for its cardinality and $S(i)$ for its i -th element.

Table 1: List of symbols in the order they appear in the paper.

Symbol	Description
$\mathcal{N}$	All platforms in the railway network
$\mathcal{A}$	All tracks in the railway network
$\mathcal{T}$	Collection of trains under consideration
$\mathcal{N}^t$	Platforms visited by train t , listed in chronological order
$\mathcal{A}^t$	Tracks traversed by train t , listed in chronological order
a_i^t	Decision variable for the arrival time of train t at platform i
d_i^t	Decision variable for the departure time of train t from platform i
$[\underline{\tau}_{ij}^t, \bar{\tau}_{ij}^t]$	Admissible trip-time window for train t from platform i to platform j on a track
φ	Set of crossing-overs in the network
$\mathcal{B}_{ij}$	Train pairs associated with turn-around movements on crossing-over (i, j)
$[\underline{\kappa}_{ij}^{tt'}, \bar{\kappa}_{ij}^{tt'}]$	Admissible trip-time window for the train pair (t, t') on crossing-over (i, j)
$[\underline{\delta}_i^t, \bar{\delta}_i^t]$	Allowable dwell-time window for train t at platform i
χ	Platform pairs located at the same interchange station
$\mathcal{C}_{ij}$	Connecting train pairs associated with platform pair $(i, j) \in \chi$
$[\underline{\chi}_{ij}^{tt'}, \bar{\chi}_{ij}^{tt'}]$	Transfer-time window between train t at platform i and train t' at platform j
$\mathcal{H}_{ij}$	Ordered train pairs that successively traverse track (i, j)
$h_i^{tt'}$	Required headway between the departures of trains t and t' from platform i
$[\underline{\tau}_{\mathcal{P}}^t, \bar{\tau}_{\mathcal{P}}^t]$	Allowable total travel-time window for train t along its train path
$E_{i,j,t}^{\text{con,tr}}$	Traction energy consumed by train t while accelerating on the trip from platform i to platform j
$E_{i,j,t,t'}^{\text{con,cr}}$	Traction energy consumed by train t while accelerating through crossing-over $(i, j) \in \varphi$
Ω	Feasible platform pairs for regenerative-energy transfer
$\mathcal{T}_i$	Trains that arrive at, dwell at, and depart from platform i
μ_i^t	Dwell midpoint $(\bar{a}_i^t + \bar{d}_i^t)/2$ of train t at platform i in the reference timetable
$\overrightarrow{t}$	Train temporally close to the right of t
$\overleftarrow{t}$	Train temporally close to the left of t

Table 1 (continued)

Symbol	Description
$\tilde{t}$	Train temporally close to t
$\mathcal{E}$	Synchronization events between suitable train pairs
$\vec{\mathcal{E}}$	Subset of $\mathcal{E}$ containing tuples of the form $(i, j, t, \vec{t})$
$\overleftarrow{\mathcal{E}}$	Subset of $\mathcal{E}$ containing tuples of the form $(i, j, t, \overleftarrow{t})$
$E_{i,j,t,\vec{t}}^{\text{reg}}$	Regenerative energy transferred to train t (accelerating at i) by train $\vec{t}$ (braking at j)
$\sigma_{ij}^{t\vec{t}}$	Overlap of the accelerating phase of t at i with the braking phase of $\vec{t}$ at j
$E_{i,j,t,\overleftarrow{t}}^{\text{reg}}$	Regenerative energy transferred to train $\overleftarrow{t}$ (accelerating at j) by train t (braking at i)
$\sigma_{ij}^{\overleftarrow{t}t}$	Overlap of the accelerating phase of $\overleftarrow{t}$ at j with the braking phase of t at i
$\underline{\beta}_i^t, \overline{\beta}_i^t$	Start and end offsets of the effective braking phase of train t before its arrival at platform i ($\underline{\beta}_i^t \geq \overline{\beta}_i^t$)
$\underline{\alpha}_i^t, \overline{\alpha}_i^t$	Start and end offsets of the effective accelerating phase of train t after its departure from platform i
$\varpi_{ij}^{t\vec{t}}, \psi_{ij}^{t\overleftarrow{t}}$	Auxiliary variables that linearize the overlap $\sigma_{ij}^{t\vec{t}}$ in the final model (OPT_MODEL)

Platforms are elements of the set $\mathcal{N}$. We say that two platforms are *opposite* if they belong to the same station but trains on those platforms travel in opposite directions. Two tracks are likewise opposite if they join the same pair of stations in opposite directions. We use the term *track* for a directed arc joining two distinct, non-opposite platforms, and collect all such arcs in $\mathcal{A}$.

The railway network is represented by the directed graph $\mathcal{G} = (\mathcal{N}, \mathcal{A})$. A *train-line* is a directed path whose nodes correspond to non-opposite platforms and whose arcs correspond to non-opposite tracks. Next, a *crossing-over* is a directed arc linking two train-lines at their terminal platforms. In CBTC terminology, when a physical train reaches the terminal platform of one line, traverses a crossing-over, and then begins service on another line, the CBTC management treats the movement *operationally* involving two distinct trains (even though physically there is one physical train). The set of trains considered in the model is $\mathcal{T}$. For any $t \in \mathcal{T}$, the platforms and tracks visited by t , listed in temporal order, are denoted by $\mathcal{N}^t \subseteq \mathcal{N}$ and $\mathcal{A}^t \subseteq \mathcal{A}$, respectively. The directed path formed by these ordered platforms and tracks is called the *train-path* of t .

The optimization variables are the arrival and departure times of trains at the platforms they visit. We write a_i^t for the arrival time of train $t \in \mathcal{T}$ at platform $i \in \mathcal{N}^t$, and d_j^t for the departure time of train t from platform $j \in \mathcal{N}^t$.

3 Modeling the operational constraints

The operational constraints relate the arrival and departure events of the trains and are necessary for the proper functioning of the railway network; we present them in Sections 3.1–3.5. First, we describe how the model protects against uncertainty in the constraint data.

Incorporating robust constraints through box uncertainty. In practice, there can be some uncertainty in the data associated with the operational constraints, which can cause them to diverge from their intended values. To guard against such uncertainties, we integrate box uncertainty into our model. Abstractly, every operational constraint below can be represented as $\ell \leq x - y \leq u$, where x and y stand for decision variables, and ℓ and u refer to the problem data. We model the uncertainty in the data by the boxes $\ell \in [\ell_{\text{lb}}, \ell_{\text{ub}}]$ and $u \in [u_{\text{lb}}, u_{\text{ub}}]$. Requiring $\ell \leq x - y \leq u$ to hold for every ℓ and u in these boxes amounts to imposing the worst case in each box, so the robust counterpart of the constraint is $\ell_{\text{ub}} \leq x - y \leq u_{\text{lb}}$. For all our instances, $\ell_{\text{ub}} \leq u_{\text{lb}}$ holds; therefore,




Figure 2: Graphical summary of the operational constraints that any feasible timetable in a metro railway network must satisfy.

each robust window is nonempty. For brevity, we present the robust versions of our constraints directly in the representations that follow. While all the constraints share this abstract form, they represent very different physical requirements, which Figure 2 illustrates.

3.1 Trip-time constraint

Trip-time limits are central because they govern both traction energy use and also regenerative-energy production. As we discussed earlier, nearly all traction energy for a train's trip from one platform to the next is spent during acceleration, while regenerative energy is generated during braking. We therefore distinguish the following two trip-time constraints.

3.1.1 Track-based trip-time constraint

Consider train $t \in \mathcal{T}$ traveling from platform i to platform j on track $(i, j) \in \mathcal{A}^t$. Let d_i^t and a_j^t denote its departure and arrival times. The resulting travel time must lie in the admissible window $[\underline{\tau}_{ij}^t, \bar{\tau}_{ij}^t]$. Therefore,

$$\underline{\tau}_{ij}^t \leq a_j^t - d_i^t \leq \bar{\tau}_{ij}^t, \quad \text{for } (i, j) \in \mathcal{A}^t, t \in \mathcal{T}. \quad (\text{TRACK})$$

This constraint is shown in Figure 2(a).

3.1.2 Crossing-over-based trip-time constraint

A crossing-over is a directed arc connecting two train-lines at their terminal platforms. Consider $(i, j) \in \varphi$, where platforms i and j belong to different lines, and let $\mathcal{B}_{ij}$ denote the train pairs associated with turn-around movements on this crossing-over. For $(t, t') \in \mathcal{B}_{ij}$, train t departs platform i , traverses (i, j) , and then continues from platform j under the new operational label t' . The elapsed time between departure from i as train t and arrival at j as train t' must remain in the admissible window $[\underline{\kappa}_{ij}^{tt'}, \bar{\kappa}_{ij}^{tt'}]$. Hence,

$$\underline{\kappa}_{ij}^{tt'} \leq a_j^{t'} - d_i^t \leq \bar{\kappa}_{ij}^{tt'}, \quad \text{for } (t, t') \in \mathcal{B}_{ij}, (i, j) \in \varphi. \quad (\text{CROSS})$$

This constraint is shown in Figure 2(b).

3.2 Dwell-time constraint

When train $t \in \mathcal{T}$ reaches platform $i \in \mathcal{N}^t$, it remains there for a dwell interval in $[\underline{\delta}_i^t, \bar{\delta}_i^t]$ so that old passengers can get off and new passengers can board. Accordingly, the difference between departure and arrival times at that platform must satisfy

$$\underline{\delta}_i^t \leq d_i^t - a_i^t \leq \bar{\delta}_i^t, \quad \text{for } i \in \mathcal{N}^t, t \in \mathcal{T}. \quad (\text{DWELL})$$

Figure 2(c) depicts this constraint. The index i ranges over all of $\mathcal{N}^t$ because every train dwells at each platform on its path, including the first and last ones: the train reaches $\mathcal{N}^t(1)$ either from the depot or via a turn-around movement, and after departing $\mathcal{N}^t(|\mathcal{N}^t|)$ it either returns to the depot or begins service on another line through a turn-around.

3.3 Connection constraint

Passengers may need to transfer at interchange stations when no direct train serves their desired origin–destination pair. Let $\chi \subseteq \mathcal{N} \times \mathcal{N}$ collect the platform pairs at which such transfers are possible. For $(i, j) \in \chi$, the two platforms belong to the same station. Let $\mathcal{C}_{ij}$ denote the train pairs (t, t') for which passengers arrive on train t at platform i and continue on train t' departing from platform j . In a feasible timetable, the time available for this transfer has to remain inside the prescribed window. Accordingly, the connection constraint is

$$\underline{\chi}_{ij}^{tt'} \leq d_j^{t'} - a_i^t \leq \bar{\chi}_{ij}^{tt'}, \quad \text{for } (t, t') \in \mathcal{C}_{ij}, (i, j) \in \chi, \quad (\text{CONNECT})$$

where $\underline{\chi}_{ij}^{tt'}$ and $\bar{\chi}_{ij}^{tt'}$ are the lower and upper limits of the transfer window. Figure 2(d) illustrates this constraint.

3.4 Headway constraint

Railway operations also require a minimum temporal separation between consecutive trains, called the *headway*. For track $(i, j) \in \mathcal{A}$, let $\mathcal{H}_{ij}$ denote the ordered train pairs that traverse the track successively. If $(t, t') \in \mathcal{H}_{ij}$, then train t' follows train t on the same track. The headway requirements at platforms i and j are $h_i^{tt'}$ and $h_j^{tt'}$, which leads to the following constraint:

$$h_i^{tt'} \leq d_i^{t'} - d_i^t \quad \text{and} \quad h_j^{tt'} \leq d_j^{t'} - d_j^t, \quad \text{for } (t, t') \in \mathcal{H}_{ij}, (i, j) \in \mathcal{A}. \quad (\text{HEADWAY})$$

The headway constraint is shown in Figure 2(e).

The headway constraints also restrict regenerative-energy synchronization. When a train accelerates away from a platform, that platform cannot be occupied immediately by another train, so pairing two consecutive trains on the same track (one accelerating and one braking at the same platform) is never feasible from a safety standpoint. We account for this restriction when constructing suitable train pairs in Section 4.2.

3.5 Total travel-time constraint

Reliable service requires the total travel time of each train $t \in \mathcal{T}$ to lie within a specified window $[\underline{\tau}_{\mathcal{P}}^t, \bar{\tau}_{\mathcal{P}}^t]$, where the subscript $\mathcal{P}$ abbreviates *train-path*: the time elapsed between the train's departure from its first platform $\mathcal{N}^t(1)$ and its arrival at its last platform $\mathcal{N}^t(|\mathcal{N}^t|)$ must satisfy

$$\underline{\tau}_{\mathcal{P}}^t \leq a_{\mathcal{N}^t(|\mathcal{N}^t|)}^t - d_{\mathcal{N}^t(1)}^t \leq \bar{\tau}_{\mathcal{P}}^t, \quad \text{for } t \in \mathcal{T}. \quad (\text{TRAVEL})$$

Keeping total travel times within these windows protects connections and preserves network capacity. This constraint is shown in Figure 2(f).

Domain of the event times

Bounds on event times are set relative to the service horizon. We normalize the beginning of the service period to time zero, which corresponds to the day's first departure from some platform. An upper bound on the final event can then be obtained by assigning every trip time and dwell time its maximal allowed value; call this bound $m > 0$. The arrival and departure variables are therefore restricted by

$$0 \leq a_i^t \leq m, \quad 0 \leq d_i^t \leq m, \quad \text{for } i \in \mathcal{N}^t, t \in \mathcal{T}. \quad (\text{DOMAIN})$$

4 Modeling the effective energy consumption

This section develops the objective representing *effective energy consumption*. We define this quantity as the total traction energy consumed during train acceleration, E^{con} , minus the regenerative braking energy transferred to eligible accelerating trains, E^{reg} . Hence the objective is $E^{\text{con}} - E^{\text{reg}}$. The data-driven model for E^{con} is presented in Section 4.1, followed by the model for E^{reg} in Section 4.2.

4.1 Modeling total consumed energy E^{con}

Relationship between energy consumption and the trip-time constraint. Most electrical energy used by a train is drawn during acceleration between successive platforms, so trip-time decisions are a primary determinant of both traction consumption and regenerative-energy generation. Once a trip time is specified, CBTC tools can compute the corresponding energy-optimal speed profile. In our study we use KDC (Hitachi Rail, 2026), a proprietary physics-based train-performance simulator used within Hitachi Rail's CBTC toolchain, and it computes speed profiles that follow the acceleration, speed-holding, coasting, and braking sequence. This profile is consistent with the theoretical optimum described by Howlett and Pudney (1995); see Cheng and Howlett (1993); Howlett (2000); Howlett et al. (2009); Khmel'nitsky (2000); Liu and Golovitcher (2003) for foundational treatments and Albrecht et al. (2016) for a broader review. The speed profile determines the train's electrical power draw and regeneration over time, and therefore its associated power-versus-time curve, illustrated qualitatively in Figure 1.




Figure 3: Qualitative behavior of the traction-energy function for train t on track (i, j) as the trip time $a_j^t - d_i^t$ varies. The wide horizontal range is included only to display the nonlinear shape observed in long inter-city runs. In metro operation, the feasible interval $[\underline{\tau}_{ij}^t, \bar{\tau}_{ij}^t]$ is typically measured in seconds, and over that narrow range an affine approximation of consumed energy is reasonable.

Energy consumption of all trains. We represent the network-wide traction energy used by all trains as

$$E^{\text{con}} \triangleq \sum_{t \in \mathcal{T}} \sum_{(i,j) \in \mathcal{A}^t} \underbrace{E_{i,j,t}^{\text{con,tr}}}_{\text{depends on trip time } (a_j^t - d_i^t)} + \sum_{(i,j) \in \varphi, (t,t') \in \mathcal{B}_{ij}} \underbrace{E_{i,j,t,t'}^{\text{con,cr}}}_{\text{depends on trip time } (a_j^{t'} - d_i^t)},$$

where $E_{i,j,t}^{\text{con,tr}} : \mathbb{R}_{++} \rightarrow \mathbb{R}_{++}$ denotes the traction energy used by train $t \in \mathcal{T}$ while accelerating on the trip from platform i to platform j for $(i, j) \in \mathcal{A}^t$; in a CBTC setting its argument is the trip time $a_j^t - d_i^t$. Likewise, $E_{i,j,t,t'}^{\text{con,cr}} : \mathbb{R}_{++} \rightarrow \mathbb{R}_{++}$ represents the traction energy used by train t when it traverses crossing-over $(i, j) \in \varphi$ and then continues from platform j under label t' , with $(t, t') \in \mathcal{B}_{ij}$. Its argument is the corresponding crossing-over travel time $a_j^{t'} - d_i^t$.

Data-driven approach to model the energy consumption. The functions $E_{i,j,t}^{\text{con,tr}}$ and $E_{i,j,t,t'}^{\text{con,cr}}$ generally do not admit a tractable closed form (Howlett et al., 2009; Das Gupta et al., 2016), although their qualitative behavior is well understood. Under an optimal speed profile, each is monotone nonincreasing in trip time (Milroy, 1980). Empirical studies indicate that the same monotonic trend persists on average even under manual or otherwise suboptimal driving; see Peña-Alcaraz et al. (2011, Figure 1). Moreover, the relevant energy data are accessible in practice, either from historical operations or from physics-based simulation. A key feature of CBTC metro systems is that the permissible variation in the trip-time constraints (TRACK) and (CROSS) is usually only a few seconds. We therefore adopt the following assumption:

Assumption 4.1. *For every track-based or crossing-over trip-time constraint, the width of the admissible time window is on the order of seconds; equivalently, $\bar{\tau}_{ij}^t - \underline{\tau}_{ij}^t$ in (TRACK) and $\bar{\kappa}_{ij}^{tt'} - \underline{\kappa}_{ij}^{tt'}$ in (CROSS) are on the order of seconds.*

Assumption 4.1, together with the monotone dependence on trip time, justifies affine approxima-

tions of the energy-consumption functions. Figure 3 illustrates this idea for a track-based trip; the crossing-over case behaves analogously. To obtain the best affine fit, we first apply least-squares to energy-versus-trip-time data, using either historical observations or outputs from physics-based simulators such as the KDC simulator employed in this work (Hitachi Rail, 2026).

In other words, we approximate consumed energy as

$$\begin{aligned} E_{i,j,t}^{\text{con,tr}} &\triangleq c_{i,j,t}^{\text{con,tr}}(a_j^t - d_i^t) + b_{i,j,t}^{\text{con,tr}}, \quad \text{for } (i,j) \in \mathcal{A}^t, t \in \mathcal{T}, \\ E_{i,j,t,t'}^{\text{con,cr}} &\triangleq c_{i,j,t,t'}^{\text{con,cr}}(a_j^{t'} - d_i^t) + b_{i,j,t,t'}^{\text{con,cr}}, \quad \text{for } (i,j) \in \varphi, (t,t') \in \mathcal{B}_{ij}, \end{aligned} \quad (1)$$

where $(c_{i,j,t}^{\text{con,tr}}, b_{i,j,t}^{\text{con,tr}})$ is computed by solving the following least-squares problem:

$$(c_{i,j,t}^{\text{con,tr}}, b_{i,j,t}^{\text{con,tr}}) = \underset{(\tilde{c}_{i,j,t}, \tilde{b}_{i,j,t})}{\operatorname{argmin}} \sum_{k=1}^p \left(\tilde{c}_{i,j,t}(\bar{a}_j^t - \bar{d}_i^t)^{(k)} + \tilde{b}_{i,j,t} - (\bar{E}_{i,j,t}^{\text{con,tr}})^{(k)} \right)^2,$$

where $(i,j) \in \mathcal{A}^t, t \in \mathcal{T}$; here and in what follows, bars distinguish observed data from decision variables, so $(\bar{E}_{i,j,t}^{\text{con,tr}})^{(k)}$ denotes the measured traction energy associated with the k -th observed track-based trip time $(\bar{a}_j^t - \bar{d}_i^t)^{(k)}$, for $k = 1, 2, \dots, p$. Likewise, $(c_{i,j,t,t'}^{\text{con,cr}}, b_{i,j,t,t'}^{\text{con,cr}})$ is computed by solving the following least-squares problem:

$$(c_{i,j,t,t'}^{\text{con,cr}}, b_{i,j,t,t'}^{\text{con,cr}}) = \underset{(\tilde{c}_{i,j,t,t'}, \tilde{b}_{i,j,t,t'})}{\operatorname{argmin}} \sum_{k=1}^q \left(\tilde{c}_{i,j,t,t'}(\bar{a}_j^{t'} - \bar{d}_i^t)^{(k)} + \tilde{b}_{i,j,t,t'} - (\bar{E}_{i,j,t,t'}^{\text{con,cr}})^{(k)} \right)^2,$$

where $(i,j) \in \varphi, (t,t') \in \mathcal{B}_{ij}$ with $(\bar{E}_{i,j,t,t'}^{\text{con,cr}})^{(k)}$ denoting the measured traction energy associated with the k -th observation of the crossing-over trip time $(\bar{a}_j^{t'} - \bar{d}_i^t)^{(k)}$, for $k = 1, 2, \dots, q$.

4.2 Modeling transferred regenerative energy E^{reg}

Robust approximations of the power-versus-time graph. Maximizing regenerative-energy exchange over suitable train pairs is equivalent to maximizing the overlap between their consumption and regeneration power profiles. In CBTC systems, these power-versus-time curves depend on trip time and are highly nonlinear, with no convenient analytical form. Empirically, however, most of the power is concentrated near a peak, which motivates a robust lumped approximation via the FWHM (Full Width at Half Maximum) rule. This idea, popularized in early quantum mechanics (Gamow, 1985) and now standard in areas such as numerical analysis, spectroscopy, signal processing, and statistics (Mahajan, 2010, pp. 35–37), replaces each power profile by a rectangle whose height equals the peak value and whose width equals the full width at half maximum. In our setting, these rectangles are associated with the effective accelerating and braking intervals containing most of the energy as shown in Figure 4.

Consider a train t arriving at (i.e., braking), then dwelling, and finally departing (i.e., accelerating) from platform i . As shown in Figure 4, the effective braking and accelerating phases of the train can be compactly contained in the FWHM rectangles. We denote the beginning and end of the effective braking phase of t by $a_i^t - \underline{\beta}_i^t$ and $a_i^t - \bar{\beta}_i^t$, respectively, and the beginning and end of the effective accelerating phase of train t by $d_i^t + \underline{\alpha}_i^t$ and $d_i^t + \bar{\alpha}_i^t$, respectively. Note the resulting ordering: $\underline{\beta}_i^t$ and $\bar{\beta}_i^t$ are both offsets measured backward from the arrival time, so the effective braking phase is the interval $[a_i^t - \underline{\beta}_i^t, a_i^t - \bar{\beta}_i^t]$ with $\underline{\beta}_i^t \geq \bar{\beta}_i^t$, whereas the accelerating-phase offsets satisfy $\underline{\alpha}_i^t \leq \bar{\alpha}_i^t$.




Figure 4: Effective braking and accelerating phases of a train t at platform i , computed with the FWHM (Full Width at Half Maximum) rule. Each power profile is replaced by a rectangle whose height equals the peak power and whose width equals the full width at half maximum. The rectangles determine the endpoints $a_i^t - \underline{\beta}_i^t$ and $a_i^t - \overline{\beta}_i^t$ of the effective braking phase, and $d_i^t + \underline{\alpha}_i^t$ and $d_i^t + \overline{\alpha}_i^t$ of the effective accelerating phase.

Once these effective phase endpoints are known, the *effective overlapping time* between two trains correspond to the duration for which the effective accelerating interval of one train overlaps the effective braking interval of the other, provided that the two platforms are supplied by the same electrical substation so that transmission losses can be neglected. Increasing this overlap increases the amount of regenerative energy that can be transferred between suitable train pairs, which we model next.

Constructing suitable train pairs for transfer of regenerative braking energy. We now identify train pairs that can realistically exchange regenerative energy. Let Ω collect the platform pairs between which such a transfer is physically possible: pairs supplied by the same electrical substation, so that transmission losses are negligible. For any $(i, j) \in \Omega$, let $\mathcal{T}_i \subseteq \mathcal{T}$ be the trains that arrive at, dwell at, and depart from platform i . To avoid duplicate platform pairs, we index the elements of Ω lexicographically with $i < j$. Fix $t \in \mathcal{T}_i$. We seek a train $\tilde{t} \in \mathcal{T}_j$ that, together with t , forms a suitable pair for regenerative-energy transfer. Our selection is based on an initial operating feasible timetable; in existing railway networks operated by CBTC systems such an operational timetable is always available. This initial feasible timetable provides feasible arrival and departure times $\bar{a}_i^t$ and $\bar{d}_i^t$ for every platform $i \in \mathcal{N}^t$; we write $\mu_i^t \triangleq (\bar{a}_i^t + \bar{d}_i^t)/2$ for the midpoint of the corresponding dwell interval. Intuitively, among the trains that use platform j , the best candidate is the one whose dwell midpoint lies closest in time to that of t in the reference timetable. This closeness can occur on either side of t , leading to the following definitions.

Definition 4.1. Fix $(i, j) \in \Omega$ and a train $t \in \mathcal{T}_i$. Consider the trains $s \in \mathcal{T}_j$ with $0 \leq \mu_j^s - \mu_i^t \leq r$, where r is an empirical parameter chosen by the timetable designer and much smaller than the full timetable horizon. If such trains exist, the one minimizing $\mu_j^s - \mu_i^t$, with ties broken by the trains'

index order, is called temporally close to the right of t and is denoted by $\vec{t}$; otherwise t has no right event on (i, j) .

Definition 4.2. Fix $(i, j) \in \Omega$ and a train $t \in \mathcal{T}_i$. Consider the trains $s \in \mathcal{T}_j$ with $0 < \mu_i^t - \mu_j^s \leq r$. If such trains exist, the one minimizing $\mu_i^t - \mu_j^s$, with ties broken by the trains' index order, is called temporally close to the left of t and is denoted by $\overleftarrow{t}$; otherwise t has no left event on (i, j) . The left inequality is strict so that a pair of trains with coincident midpoints generates only a right event.

Definition 4.3. Fix $(i, j) \in \Omega$. For any train $t \in \mathcal{T}_i$, a train $\tilde{t} \in \mathcal{T}_j$ is temporally close to t if it is temporally close to the right or to the left of t .

Thus, each potential synchronization event is identified by a tuple $(i, j, t, \tilde{t})$ satisfying the definitions above; because the platform pairs in Ω satisfy $i < j$, no event is enumerated twice. We collect these tuples in the set $\mathcal{E}$ and partition $\mathcal{E}$ into the subsets $\vec{\mathcal{E}}$ and $\overleftarrow{\mathcal{E}}$, which contain tuples of the forms $(i, j, t, \vec{t})$ and $(i, j, t, \overleftarrow{t})$, respectively. For a right event $(i, j, t, \vec{t}) \in \vec{\mathcal{E}}$, we seek to align the effective accelerating phase of t with the effective braking phase of $\vec{t}$. For a left event $(i, j, t, \overleftarrow{t}) \in \overleftarrow{\mathcal{E}}$, the corresponding alignment is between the effective accelerating phase of $\overleftarrow{t}$ and the effective braking phase of t . The associated effective overlapping times are denoted by $\sigma_{ij}^{t\vec{t}}$ and $\sigma_{ij}^{\overleftarrow{t}t}$, respectively. An overlap may be positive when the two effective phases intersect, zero when they just meet, or negative when they are separated in time; these cases will be illustrated graphically later in this section.

The amount of regenerative energy that can be transferred depends on the total effective overlap, so our objective is to maximize that transfer. We write it as

$$E^{\text{reg}} \triangleq \sum_{(i,j,t,\vec{t}) \in \vec{\mathcal{E}}} \underbrace{E_{i,j,t,\vec{t}}^{\text{reg}}}_{\substack{\text{depends on} \\ \text{right event} \\ \text{overlapping time } \sigma_{ij}^{t\vec{t}}}} + \sum_{(i,j,t,\overleftarrow{t}) \in \overleftarrow{\mathcal{E}}} \underbrace{E_{i,j,t,\overleftarrow{t}}^{\text{reg}}}_{\substack{\text{depends on} \\ \text{left event} \\ \text{overlapping time } \sigma_{ij}^{\overleftarrow{t}t}}}, \quad (2)$$

where (i) $E_{i,j,t,\vec{t}}^{\text{reg}} : \mathbb{R} \rightarrow \mathbb{R}$ maps the overlap between the accelerating phase of t and the braking phase of $\vec{t}$ to the corresponding transferred regenerative energy, with argument $\sigma_{ij}^{t\vec{t}}$, and (ii) $E_{i,j,t,\overleftarrow{t}}^{\text{reg}} : \mathbb{R} \rightarrow \mathbb{R}$ maps the overlap between the accelerating phase of $\overleftarrow{t}$ and the braking phase of t to transferred regenerative energy, with argument $\sigma_{ij}^{\overleftarrow{t}t}$. To model (2) tractably, we proceed in three steps: we fit the transferred-energy functions (Step 1), fit the effective accelerating and braking phases (Step 2), and express the overlaps in the decision variables (Step 3). We now describe each step in turn.

Step 1. Fitting the transferred-energy functions. Using a data-driven fit, we model $E_{i,j,t,\vec{t}}^{\text{reg}}$ and $E_{i,j,t,\overleftarrow{t}}^{\text{reg}}$ as affine functions of the effective overlapping times $\sigma_{ij}^{t\vec{t}}$ and $\sigma_{ij}^{\overleftarrow{t}t}$, respectively. This affine dependence is natural under the FWHM approximation: within the rectangles, power is nearly constant at its peak value, so the transferred energy scales with the overlap duration. Assumption 4.1 keeps the overlaps within the narrow range where such fits are accurate. Thus, we write

$$\begin{aligned}
E_{i,j,t,t}^{\text{reg} \rightarrow} &\triangleq c_{i,j,t,t}^{\text{reg} \rightarrow} \sigma_{ij}^{\overrightarrow{t}t} + b_{i,j,t,t}^{\text{reg} \rightarrow}, \\
E_{i,j,t,t}^{\text{reg} \leftarrow} &\triangleq c_{i,j,t,t}^{\text{reg} \leftarrow} \sigma_{ij}^{\overleftarrow{t}t} + b_{i,j,t,t}^{\text{reg} \leftarrow},
\end{aligned} \tag{3}$$

where $(c_{i,j,t,t}^{\text{reg} \rightarrow}, b_{i,j,t,t}^{\text{reg} \rightarrow})$ and $(c_{i,j,t,t}^{\text{reg} \leftarrow}, b_{i,j,t,t}^{\text{reg} \leftarrow})$ are nonnegative coefficients obtained from a least-squares fit constrained to a nonnegative slope and intercept, whose inputs are effective overlapping times and whose outputs are transferred regenerative-energy measurements. These input data may come either from historical operations or from physics-based simulation; in our numerical study they are generated by the KDC simulator. Positive intercepts $b_{i,j,t,t}^{\text{reg} \rightarrow}$ and $b_{i,j,t,t}^{\text{reg} \leftarrow}$ represent residual regenerative transfer, which can occur because some energy lies outside the FWHM rectangles even when the overlap is zero. Hence, (3) gives a practical interpretation of the objective in (2).

Since the fitted coefficients are nonnegative, positive overlap yields positive transferred regenerative energy. A nonpositive overlap, however, may still correspond to a small amount of transfer. When $\sigma_{ij}^{\overrightarrow{t}t} \leq -b_{i,j,t,t}^{\text{reg} \rightarrow} / c_{i,j,t,t}^{\text{reg} \rightarrow}$ (or $\sigma_{ij}^{\overleftarrow{t}t} \leq -b_{i,j,t,t}^{\text{reg} \leftarrow} / c_{i,j,t,t}^{\text{reg} \leftarrow}$), the transfer is negligible or zero. In those cases, the model allows negative values of $E_{i,j,t,t}^{\text{reg} \rightarrow}$ (or $E_{i,j,t,t}^{\text{reg} \leftarrow}$) to act as a penalty in the maximization objective. In our model, this penalty is an intentional modeling choice in contrast with truncating the transfer at zero inside the model because, even when the overlap is negative, we still want the separation as small as possible, so that some residual synchronization can be obtained through the optimization process.

Step 2. Fitting the effective accelerating and braking phases. To evaluate effective overlap, we need the start and end times of the effective accelerating and braking intervals. Consider train $t \in \mathcal{T}$ traveling from platform i to platform j on track $(i, j) \in \mathcal{A}^t$. The effective accelerating phase after departure from platform i begins at $d_i^t + \underline{\alpha}_i^t$ and ends at $d_i^t + \overline{\alpha}_i^t$. The effective braking phase before arrival at platform j begins at $a_j^t - \underline{\beta}_j^t$ and ends at $a_j^t - \overline{\beta}_j^t$. The quantities $\underline{\alpha}_i^t, \overline{\alpha}_i^t, \underline{\beta}_j^t, \overline{\beta}_j^t$ depend on the train's power profile, and hence on the trip time $a_j^t - d_i^t$. Exact analytical expressions for these quantities as functions of trip time are not available, but Assumption 4.1 again motivates affine approximations. Therefore, for $t \in \mathcal{T}$ and $(i, j) \in \mathcal{A}^t$, we define $\underline{\alpha}_i^t, \overline{\alpha}_i^t, \underline{\beta}_j^t, \overline{\beta}_j^t$ as follows:

$$\begin{aligned}
\underline{\alpha}_i^t &\triangleq c_{i,t}^{\underline{\alpha}}(a_j^t - d_i^t) + b_{i,t}^{\underline{\alpha}}, & \overline{\alpha}_i^t &\triangleq c_{i,t}^{\overline{\alpha}}(a_j^t - d_i^t) + b_{i,t}^{\overline{\alpha}}, \\
\underline{\beta}_j^t &\triangleq c_{i,t}^{\underline{\beta}}(a_j^t - d_i^t) + b_{i,t}^{\underline{\beta}}, & \overline{\beta}_j^t &\triangleq c_{i,t}^{\overline{\beta}}(a_j^t - d_i^t) + b_{i,t}^{\overline{\beta}},
\end{aligned} \tag{EFF_ACCEL_&_BRK_PHASES}$$

where the coefficient pairs $(c_{i,t}^{\underline{\alpha}}, b_{i,t}^{\underline{\alpha}})$, $(c_{i,t}^{\overline{\alpha}}, b_{i,t}^{\overline{\alpha}})$, $(c_{i,t}^{\underline{\beta}}, b_{i,t}^{\underline{\beta}})$, and $(c_{i,t}^{\overline{\beta}}, b_{i,t}^{\overline{\beta}})$ are obtained by least squares, using trip time as the input measurement and the corresponding phase start/end times as the output measurements. For each train t and platform j , the incoming track $(i, j) \in \mathcal{A}^t$ is unique, so indexing these coefficients by (i, t) is unambiguous.

Step 3. Expressing the overlaps in the decision variables. Now, we describe how to model the effective overlapping times $\sigma_{ij}^{\overrightarrow{t}t}$ and $\sigma_{ij}^{\overleftarrow{t}t}$ in terms of the decision variables.

First, we discuss the modeling of the right-event overlapping time $\sigma_{ij}^{\overrightarrow{t}t}$. By Allen's interval algebra (Allen, 1983), exactly thirteen temporal relations are possible between the accelerating phase of train t and the braking phase of train $\overrightarrow{t}$; they are shown in Figure 5. In the first four cases,




Figure 5: The thirteen possible overlapping times $\sigma_{ij}^{\bar{t}, t}$ between the accelerating phase of train t and the braking phase of train $\bar{t}$. In the first four cases, $\sigma_{ij}^{\bar{t}, t}$ is nonpositive (red-shaded region); in the remaining nine cases, it is positive (green-shaded region). All thirteen cases obey the closed form $\sigma_{ij}^{\bar{t}, t} = \min \left\{ a_j^t - \bar{\beta}_j^t, d_i^t + \bar{\alpha}_i^t \right\} + \min \left\{ -d_i^t - \underline{\alpha}_i^t, -a_j^t + \underline{\beta}_j^t \right\}$.

the overlap is nonpositive: the two intervals are disjoint (red-shaded); in the remaining nine, it is positive (green-shaded). Fortunately, we can write all thirteen possible overlapping times in closed form:

$$\begin{aligned}\sigma_{ij}^{\vec{t}\vec{t}} &= \min \left\{ \vec{a}_j^{\vec{t}} - \vec{\beta}_j^{\vec{t}}, d_i^t + \vec{\alpha}_i^t \right\} - \max \left\{ d_i^t + \underline{\alpha}_i^t, \vec{a}_j^{\vec{t}} - \underline{\beta}_j^{\vec{t}} \right\} \\ &= \min \left\{ \vec{a}_j^{\vec{t}} - \vec{\beta}_j^{\vec{t}}, d_i^t + \vec{\alpha}_i^t \right\} + \min \left\{ -d_i^t - \underline{\alpha}_i^t, -\vec{a}_j^{\vec{t}} + \underline{\beta}_j^{\vec{t}} \right\},\end{aligned}\quad (\text{RGHT_EVNT_OVLP_TIME})$$

which is concave in the decision variables, being a pointwise minimum of affine functions (Boyd and Vandenberghe, 2004, Section 3.2). For example, suppose train t departs platform i at $d_i^t = 100$ s with $\underline{\alpha}_i^t = 4$ s and $\vec{\alpha}_i^t = 18$ s, while train $\vec{t}$ arrives at platform j at $\vec{a}_j^{\vec{t}} = 120$ s with $\underline{\beta}_j^{\vec{t}} = 15$ s and $\vec{\beta}_j^{\vec{t}} = 3$ s. Then the effective accelerating phase is $[104, 118]$, the effective braking phase is $[105, 117]$, and (RGHT_EVNT_OVLP_TIME) gives $\sigma_{ij}^{\vec{t}\vec{t}} = \min\{117, 118\} - \max\{104, 105\} = 12$ s, exactly the length of time the two phases share. Next, we discuss the modeling of the left-event overlapping time $\sigma_{ij}^{\overleftarrow{t}\overleftarrow{t}}$. Replacing i, j, t , and $\vec{t}$ in (RGHT_EVNT_OVLP_TIME) with $j, i, \overleftarrow{t}$, and t , respectively, gives

$$\begin{aligned}\sigma_{ij}^{\overleftarrow{t}\overleftarrow{t}} &= \min \left\{ a_i^t - \vec{\beta}_i^{\overleftarrow{t}}, d_j^{\overleftarrow{t}} + \vec{\alpha}_j^{\overleftarrow{t}} \right\} - \max \left\{ d_j^{\overleftarrow{t}} + \underline{\alpha}_j^{\overleftarrow{t}}, a_i^t - \underline{\beta}_i^{\overleftarrow{t}} \right\} \\ &= \min \left\{ a_i^t - \vec{\beta}_i^{\overleftarrow{t}}, d_j^{\overleftarrow{t}} + \vec{\alpha}_j^{\overleftarrow{t}} \right\} + \min \left\{ -d_j^{\overleftarrow{t}} - \underline{\alpha}_j^{\overleftarrow{t}}, -a_i^t + \underline{\beta}_i^{\overleftarrow{t}} \right\},\end{aligned}\quad (\text{LEFT_EVNT_OVLP_TIME})$$

which is again a concave function of our decision variables. Because each overlap is concave in the event times and each Step 1 fit is affine with a nonnegative slope, the total transferred regenerative energy is a concave function of the timetable; Section 5 maximizes it through linear hypograph constraints.

5 Final optimization model

Combining the ingredients developed above yields the following linear optimization model for minimizing the railway network's effective energy consumption:

minimize

$$\begin{aligned} & \sum_{t \in \mathcal{T}, (i,j) \in \mathcal{A}^t} \left(c_{i,j,t}^{\text{con,tr}} (a_j^t - d_i^t) + b_{i,j,t}^{\text{con,tr}} \right) + \sum_{(i,j) \in \varphi, (t,t') \in \mathcal{B}_{ij}} \left(c_{i,j,t,t'}^{\text{con,cr}} (a_j^{t'} - d_i^t) + b_{i,j,t,t'}^{\text{con,cr}} \right) \triangleright \text{consumption} \\ & - \sum_{(i,j,t,\vec{t}) \in \vec{\mathcal{E}}} \left(c_{i,j,t,\vec{t}}^{\text{reg}} \sigma_{ij}^{\vec{t}\vec{t}} + b_{i,j,t,\vec{t}}^{\text{reg}} \right) - \sum_{(i,j,t,\overleftarrow{t}) \in \overleftarrow{\mathcal{E}}} \left(c_{i,j,t,\overleftarrow{t}}^{\text{reg}} \sigma_{ij}^{\overleftarrow{t}\overleftarrow{t}} + b_{i,j,t,\overleftarrow{t}}^{\text{reg}} \right) \triangleright \text{transferred regen. energy} \end{aligned}$$

subject to

$$\left. \begin{aligned} \sigma_{ij}^{\vec{t}\vec{t}} &\leq \varpi_{ij}^{\vec{t}\vec{t}} + \psi_{ij}^{\vec{t}\vec{t}} \\ \varpi_{ij}^{\vec{t}\vec{t}} &\leq a_j^{\vec{t}} - \bar{\beta}_j^{\vec{t}} \\ \varpi_{ij}^{\vec{t}\vec{t}} &\leq d_i^{\vec{t}} + \bar{\alpha}_i^{\vec{t}} \\ \psi_{ij}^{\vec{t}\vec{t}} &\leq -d_i^{\vec{t}} - \underline{\alpha}_i^{\vec{t}} \\ \psi_{ij}^{\vec{t}\vec{t}} &\leq -a_j^{\vec{t}} + \underline{\beta}_j^{\vec{t}} \end{aligned} \right\} \text{ for } (i,j,t,\vec{t}) \in \vec{\mathcal{E}} \quad \begin{array}{l} \triangleright \text{hypograph form} \\ \text{of (RGHT_EVNT_OVLP_TIME)} \end{array}$$

$$\left. \begin{aligned} \sigma_{ij}^{\overleftarrow{t}\overleftarrow{t}} &\leq \varpi_{ij}^{\overleftarrow{t}\overleftarrow{t}} + \psi_{ij}^{\overleftarrow{t}\overleftarrow{t}} \\ \varpi_{ij}^{\overleftarrow{t}\overleftarrow{t}} &\leq a_i^{\overleftarrow{t}} - \bar{\beta}_i^{\overleftarrow{t}} \\ \varpi_{ij}^{\overleftarrow{t}\overleftarrow{t}} &\leq d_j^{\overleftarrow{t}} + \bar{\alpha}_j^{\overleftarrow{t}} \\ \psi_{ij}^{\overleftarrow{t}\overleftarrow{t}} &\leq -d_j^{\overleftarrow{t}} - \underline{\alpha}_j^{\overleftarrow{t}} \\ \psi_{ij}^{\overleftarrow{t}\overleftarrow{t}} &\leq -a_i^{\overleftarrow{t}} + \underline{\beta}_i^{\overleftarrow{t}} \end{aligned} \right\} \text{ for } (i,j,t,\overleftarrow{t}) \in \overleftarrow{\mathcal{E}} \quad \begin{array}{l} \triangleright \text{hypograph form} \\ \text{of (LEFT_EVNT_OVLP_TIME)} \end{array}$$

Eq. (EFF_ACCEL_&_BRK_PHASES), $\triangleright$ effective-phase endpoints

Eq. (TRACK), (CROSS), (DWELL), (CONNECT), $\triangleright$ operational constraints

Eq. (HEADWAY), (TRAVEL), (DOMAIN), $\triangleright$ headway, total travel time, and domain

(OPT_MODEL)

where the decision variables are a_i^t , d_i^t , $\sigma_{ij}^{\vec{t}\vec{t}}$, $\varpi_{ij}^{\vec{t}\vec{t}}$, $\psi_{ij}^{\vec{t}\vec{t}}$, $\sigma_{ij}^{\overleftarrow{t}\overleftarrow{t}}$, $\varpi_{ij}^{\overleftarrow{t}\overleftarrow{t}}$, and $\psi_{ij}^{\overleftarrow{t}\overleftarrow{t}}$. Note that the first ten constraints recast (RGHT_EVNT_OVLP_TIME) and (LEFT_EVNT_OVLP_TIME) as linear constraints using the hypograph approach (Boyd and Vandenberghe, 2004, Section 4.1.3). Because the objective coefficients c^{reg} are nonnegative, the optimization drives each auxiliary variable up to its bound. Hence, at an optimal solution, every overlap variable with a positive coefficient equals its overlap value in (RGHT_EVNT_OVLP_TIME) or (LEFT_EVNT_OVLP_TIME). In words, (OPT_MODEL) chooses the arrival and departure times, subject to all the operational constraints, to make the predicted traction energy minus the transferred regenerative energy as small as possible.

Predicting effective energy consumption from the solution. After solving (OPT_MODEL), we obtain an energy-optimal timetable. Its predicted effective energy consumption is

$$\begin{aligned}
& \sum_{t \in \mathcal{T}, (i,j) \in \mathcal{A}^t} \left(c_{i,j,t}^{\text{con,tr}} (a_j^t - d_i^t) + b_{i,j,t}^{\text{con,tr}} \right) + \sum_{(i,j) \in \varphi, (t,t') \in \mathcal{B}_{ij}} \left(c_{i,j,t,t'}^{\text{con,cr}} (a_j^{t'} - d_i^t) + b_{i,j,t,t'}^{\text{con,cr}} \right) \\
& - \sum_{(i,j,t,\vec{t}) \in \vec{\mathcal{E}}} \max \left\{ c_{i,j,t,\vec{t}}^{\text{reg}} \sigma_{ij}^{\vec{t}\vec{t}} + b_{i,j,t,\vec{t}}^{\text{reg}}, 0 \right\} - \sum_{(i,j,t,\overleftarrow{t}) \in \overleftarrow{\mathcal{E}}} \max \left\{ c_{i,j,t,\overleftarrow{t}}^{\text{reg}} \sigma_{ij}^{\overleftarrow{t}\overleftarrow{t}} + b_{i,j,t,\overleftarrow{t}}^{\text{reg}}, 0 \right\}.
\end{aligned}
\tag{PRED_EFF_ENG_CNSM}$$

In (PRED_EFF_ENG_CNSM), the first line represents traction energy consumed during acceleration, while the second line subtracts the regenerative energy that is actually transferred. The full expression therefore gives the predicted effective energy consumption. When $\sigma_{ij}^{\vec{t}\vec{t}} \leq 0$ and $\sigma_{ij}^{\overleftarrow{t}\overleftarrow{t}} \leq 0$, the affine terms $c_{i,j,t,\vec{t}}^{\text{reg}} \sigma_{ij}^{\vec{t}\vec{t}} + b_{i,j,t,\vec{t}}^{\text{reg}}$ and $c_{i,j,t,\overleftarrow{t}}^{\text{reg}} \sigma_{ij}^{\overleftarrow{t}\overleftarrow{t}} + b_{i,j,t,\overleftarrow{t}}^{\text{reg}}$ act as penalties in the objective of (OPT_MODEL); accordingly, the clipped quantities $\max\{c_{i,j,t,\vec{t}}^{\text{reg}} \sigma_{ij}^{\vec{t}\vec{t}} + b_{i,j,t,\vec{t}}^{\text{reg}}, 0\}$ and $\max\{c_{i,j,t,\overleftarrow{t}}^{\text{reg}} \sigma_{ij}^{\overleftarrow{t}\overleftarrow{t}} + b_{i,j,t,\overleftarrow{t}}^{\text{reg}}, 0\}$ recover the actual regenerative energy transferred in the individual right and left events, respectively. Moreover, each affine transfer term is bounded above by its clipped counterpart, so the objective value of (OPT_MODEL) is an upper bound on (PRED_EFF_ENG_CNSM): minimizing the former minimizes an upper bound on the predicted effective energy consumption.

6 Numerical experiments

We evaluate the model on two real-world railway networks. Section 6.1 describes Metro Line 8 of the Shanghai Railway Network and presents the corresponding numerical results; Section 6.2 applies the same analysis to the Docklands Light Railway.

We ran all experiments in this section on a 16-inch MacBook Pro with an 18-core Apple M5 Max processor, 128 GB of unified memory, and macOS 26.5.2. We used Julia 1.12.6, JuMP 1.30.1 (Lubin et al., 2023), and Gurobi Optimizer 13.0.2 (Gurobi Optimization, LLC, 2026) for solving each linear program. The solution times in Tables 2 and 4 are the medians of three Gurobi-reported solver times obtained from three independent runs; each minimization solve was warm-started at the corresponding baseline timetable. The synchronization events, the platform-pair set Ω , and the fitted energy coefficients are supplied with the instance data. In the Shanghai instances, the synchronization events were constructed with the closeness parameter $r = 120$ seconds; in the Docklands Light Railway instances, the paired dwell midpoints lie within 10 minutes of each other in the reference timetables. In every instance of both networks, no braking train’s effective phase enters more than one synchronization event, so each transfer draws on a distinct braking phase.

6.1 Metro Line 8 in Shanghai Railway Network

6.1.1 Railway network in consideration

We run our experiments on the PES₂-SFM₂ service of Metro Line 8 in the Shanghai Railway Network, which is the world’s longest metro system by route length and carried approximately 3.88 billion passenger trips in 2019 (International Association of Public Transport (UITP), 2022). Metro Line 8 serves some of Shanghai’s densest residential districts and carried an average daily ridership of about 1.08 million at the time of the study of Das Gupta et al. (2016). The line extends 37.4 km and has 28 operating stations (Shanghai Shentong Metro Group, 2026).




Figure 6: Shanghai Metro Line 8 network used in the numerical study. Squares denote platforms: red on Line 1, green on Line 2; the blue squares are the terminal platforms joined to the opposite line by crossing-overs.

Three services run on Metro Line 8, each with its own CBTC-generated timetable; here we optimize the PES_2 - SFM_2 service shown in Figure 6. The corresponding network consists of two lines, denoted Line 1 and Line 2, and 14 stations, labeled in uppercase in the figure. Each station has one platform on each line; for example, LXM includes platforms LXM_1 and LXM_2 serving Line 1 and Line 2, respectively. The portion of the line considered here is 37 km long. The mean inter-station spacing is 1.4 km, with a minimum of 738 m between YHR and ZJD and a maximum of 2.6 km between PJT and LHR. Track grades lie in $[-2.00453^\circ, 2.00453^\circ]$, the maximum acceleration is 1.04 m/s^2 , and the maximum braking deceleration is 0.8 m/s^2 . The electrical-to-kinetic conversion efficiency is 0.9, the kinetic-to-regenerative-braking conversion efficiency is 0.76, and the regenerative-electricity transmission-loss factor is 0.1.

The affine energy models underlying these instances were fitted offline from KDC output; generating the training data took roughly one hour of simulation per network, and all least-squares fits complete in under a minute. The 29 track-specific traction fits are shared by all trains traversing the corresponding track and by all 11 Shanghai instances. Their coefficients remain fixed; the traction energy predicted for a traversal depends on its timetable-dependent trip time. Five simulated trip times per track, chosen across the admissible window, give coefficients of determination with median 0.98 (fifth percentile 0.83, minimum 0.76) and relative root-mean-square errors of at most 4.2% over the 29 track fits (the crossing-over fits behave similarly, with coefficients of determination between 0.82 and 0.94). The robust trip-time windows are themselves narrow: 0.4–21.2 seconds wide for (TRACK) (median 6.8) and 8.4–31.4 seconds for (CROSS). This makes Assumption 4.1 quantitative for these instances. The 25 transfer fits use 42–108 measurements each, covering effective overlaps from -32.9 to 23.0 seconds; their coefficients of determination have median 0.90 (fifth percentile 0.70, minimum 0.32), all fitted slopes and intercepts are nonnegative, and the aggregate effect of the residual fitting error on the reported reductions is bounded by the independent SPSIM evaluation of Section 6.1.2. The fitted slopes of the effective-phase offsets in (EFF_ACCEL_&_BRK_PHASES) are negligible over windows this narrow, so the instance data supply the offsets as constants, and the experiments in effect use constant effective phases.


6.1.2 Results of the numerical study

We evaluate the model on 11 operational instances of the PES_2 - SFM_2 service on Metro Line 8. The instances correspond to the eleven service configurations studied in Das Gupta et al. (2016), with baseline timetables and operational data from a more recent generation of the CBTC system. The number of trains and headway durations vary across instances. Hitachi Rail Canada supplied the corresponding baseline timetables, and we measure the benefit of the optimized schedules by comparing their energy usage against these existing timetables.


The findings of this computational study are summarized in Table 2. Across the 11 instances,

Table 2: Summary of numerical results for Metro Line 8 of the Shanghai Railway Network.

#Trains	#Variables	#Constraints	Solution time (s)	Initial effective energy consumption (kWh)	Predicted final effective energy consumption (kWh)	Reduction predicted by our model (%)	Reduction predicted by SPSIM (%)
1000	47581	151259	0.448	353721.62	254062.35	28.17	30.27
1032	49104	156102	0.432	364832.93	266932.98	26.83	31.34
1066	50726	161254	0.481	376017.98	274746.55	26.93	33.04
1100	52339	166391	0.476	386087.42	288773.33	25.21	29.10
1132	53865	171239	0.526	389844.65	294141.71	24.55	26.56
1166	55487	176391	0.530	407717.20	307967.10	24.47	26.55
1198	57010	181234	0.579	417767.32	336232.08	19.52	19.27
1232	58626	186376	0.585	447025.23	340027.49	23.94	25.86
1266	60242	191518	0.596	438657.08	337269.06	23.11	26.94
1298	61765	196361	0.651	449379.24	336794.26	25.05	24.83
1332	63387	201513	0.665	475354.59	354114.70	25.51	25.97



(a)



(b)

Figure 7: Shanghai Metro Line 8 results. (a) Model-predicted effective energy consumption of the existing and optimized timetables. (b) Percentage reductions in effective energy consumption predicted by our model and computed with SPSIM. Each reduction compares the optimized timetable with the existing timetable using the same evaluator.

all minimization solves terminated optimally, and the median solution times range from 0.432 to 0.665 seconds. Both the initial and optimized energy values in Table 2 and Figure 7a are computed with (PRED_EFF_ENG_CNSM), evaluated at the corresponding timetable’s event times. The model-predicted percentage reduction compares these two values using the model-evaluated baseline energy as its denominator. On the other hand, the SPSIM percentage reduction is computed independently by evaluating both timetables with the SPSIM simulator and using its baseline energy as the denominator. Here we note that SPSIM and KDC simulators operate independently. SPSIM computes the network-level energy consumption of an entire timetable, whereas KDC (that we used in Section 4) generates the single-train power-profile data underlying our affine fits. After solving (OPT_MODEL), we evaluate the effective energy consumption of each optimized timetable with (PRED_EFF_ENG_CNSM). The optimized timetables improve substantially on their real-world counterparts, with predicted reductions in effective energy consumption ranging from 19.52% in the least favorable instance to 28.17% in the most favorable one. Figure 7b plots these reductions.

To check how accurate the model’s predictions are, we input each optimized timetable into SPSIM and let the simulator compute the effective energy consumption and the resulting reduction. The last two columns of Table 2 and Figure 7b show that the two predictions agree closely, with our model being slightly more conservative on average. The gap between our model and SPSIM has a

simple explanation: SPSIM computes transferred regenerative energy from the exact area of overlap between the power-versus-time curves of an accelerating–braking pair, whereas our model uses the rectangular approximations of Section 4.2 that sometimes capture less energy than the exact overlap does.

Table 3 decomposes the predicted reduction into its two mechanisms: (i) lower traction energy from trip times lengthened within their admissible windows, and (ii) additional regenerative transfer from synchronized braking and acceleration. Both mechanisms contribute separately in every instance: the traction component supplies 51.0%–73.4% of the reduction and the transfer component the remaining 26.6%–49.0%, with transferred regenerative energy growing by a factor of 1.6–3.4 over the baseline timetables.

Table 3: Decomposition of the predicted effective-energy reduction for Metro Line 8: traction energy E^{con} and transferred regenerative energy E^{reg} at the baseline and optimized timetables (kWh, evaluated with the model), and each mechanism’s share of the reduction.

#Trains	Baseline		Optimized		Share of reduction (%)	
	E^{con}	E^{reg}	E^{con}	E^{reg}	traction	transfer
1000	378545.86	24824.24	327207.60	73152.61	51.5	48.5
1032	390658.93	25826.00	337978.83	71045.85	53.8	46.2
1066	403529.69	27511.71	350675.36	75941.72	52.2	47.8
1100	416400.45	30313.03	361444.27	72669.77	56.5	43.5
1132	428513.51	38668.86	373276.48	79135.21	57.7	42.3
1166	441384.27	33667.08	383841.65	75874.71	57.7	42.3
1198	453497.34	35730.02	393653.84	57403.06	73.4	26.6
1232	466368.10	19342.87	405924.57	65833.13	56.5	43.5
1266	479238.86	40581.78	418401.17	81128.18	60.0	40.0
1298	491351.92	41972.68	427633.94	90795.08	56.6	43.4
1332	504222.69	28868.10	442391.52	88264.42	51.0	49.0

6.2 Docklands Light Railway

6.2.1 Railway network in consideration

The second case study uses the Docklands Light Railway (DLR) network shown in Figure 8. The DLR is an automated light-metro network in London that operates under the SelTrac CBTC system (Hitachi Rail, 2026), and it has served as a benchmark for energy-efficient timetabling. Its topology features loops and turn-around movements, so it tests the model on a structurally different network. The benchmark network contains two directed train-lines, denoted Line 1 and Line 2. Stations are labeled by uppercase abbreviations, and each station has a platform in each direction; for example, BAN has one platform on each line. The labels 1067 and ISP1 denote intermediate stopping points on Line 1, whereas BANH and CROH denote turn-around points on Lines 1 and 2, respectively.

6.2.2 Results of the numerical study

We evaluate the model on 17 real-world instances. Each instance corresponds to an operational schedule with a different nominal service headway, ranging from 1 min 22 s to 4 min 17 s; shorter headways allow for more trains, so the instances span 108–336 active trains over a service period of approximately 4.4 hours. The baseline (initial) effective energy consumption of each instance is obtained by evaluating (PRED_EFF_ENG_CNSM) at the existing timetable. As in the Shanghai study, both the baseline and the optimized values are computed with the model, so the comparison uses a single evaluator. The energy coefficients for these instances were fitted network-wide from KDC simulations of representative Docklands Light Railway runs (a single traction fit shared by all




Figure 8: Docklands Light Railway network used in the numerical study.

trips and one transfer fit per event side), and the robust trip-time windows are 8–44 seconds wide (median 18), with crossing-over windows of 10–11 seconds. Table 4 summarizes our findings.

Table 4: Summary of numerical results for Docklands Light Railway instances.

#Trains	#Variables	#Constraints	Solution time (s)	Initial effective energy consumption (kWh)	Predicted final effective energy consumption (kWh)	Reduction predicted by our model (%)
108	2669	10109	0.020	43317.82	34942.62	19.33
121	3691	12497	0.026	48029.98	37115.47	22.72
144	4783	15531	0.036	57145.17	43842.99	23.28
156	4345	15437	0.034	62554.02	49058.68	21.57
180	5547	18703	0.044	71421.42	55989.84	21.61
192	6495	20919	0.047	74739.52	57419.50	23.17
204	6777	22025	0.049	80164.56	61514.73	23.26
216	7755	24291	0.055	84901.98	63246.11	25.51
228	8507	26171	0.060	88791.14	67805.80	23.63
240	10211	29647	0.066	92279.25	69430.24	24.76
252	9449	29013	0.083	98841.16	73490.23	25.65
263	9833	30231	0.069	102097.22	77727.41	23.87
275	8783	29117	0.077	107981.96	83646.10	22.54
288	11155	33755	0.084	111463.79	82726.07	25.78
300	11077	34261	0.105	117099.32	89096.60	23.91
311	10489	33859	0.090	123214.48	95110.66	22.81
336	14529	41913	0.108	131531.13	93620.88	28.82




Figure 9: Docklands Light Railway results. (a) Model-computed effective energy consumption of the existing timetables and of the energy-optimal timetables. (b) Corresponding percentage reduction predicted by our model.

Across the 17 instances, the number of trains ranges from 108 to 336, while the linear programs contain 2,669–14,529 variables and 10,109–41,913 constraints. All minimization solves terminated optimally, and the median solution times range from 0.020 to 0.108 seconds.

The model-computed energy values in Figure 9a generally increase with the number of trains, but the percentage reductions in Figure 9b are not monotone because they depend on how energy-efficient the initial timetables are. The predicted savings range from 19.33% to 28.82%, indicating that the benefit depends on the synchronization opportunities and timing structure of an instance, not only on its train count. Unlike in the Shanghai study, we did not have sufficient physics-related data required to predict energy savings through SPSIM accurately, so we just present the savings as predicted by our model for this network.

Table 5 again decomposes the predicted reductions in the same way as Shanghai study. With sparser traffic the traction component dominates more strongly, supplying 61.7%–89.8% of the reduction, yet synchronization still contributes 10.2%–38.3%: transferred energy grows by a factor of 2.6–6.8 over the baselines, and in the two sparsest instances (nominal headways 2 min 58 s and 4 min 17 s) the baseline timetables does not allow for any transfer at all, so the entire transferred-energy term is created by our optimization model

7 Architectural framework for industrial integration of the proposed model

This section describes the architectural framework for incorporating the proposed optimization model into Hitachi Rail Canada’s industrial timetable compiler. Figure 10 separates the existing timetable-compiler components from the energy-optimal timetable generator and shows the two data flows between them.

Table 5: Decomposition of the predicted effective-energy reduction for the Docklands Light Railway instances, in the format of Table 3.

#Trains	Baseline		Optimized		Share of reduction (%)	
	E^{con}	E^{reg}	E^{con}	E^{reg}	traction	transfer
108	43317.82	0.00	35797.28	854.66	89.8	10.2
121	48499.51	469.54	40072.60	2957.13	77.2	22.8
144	57708.61	563.44	47681.28	3838.29	75.4	24.6
156	62554.02	0.00	51691.04	2632.36	80.5	19.5
180	72099.41	677.99	59565.28	3575.44	81.2	18.8
192	76944.82	2205.30	63575.04	6155.55	77.2	22.8
204	81790.23	1625.66	67584.97	6070.24	76.2	23.8
216	86635.64	1733.66	71596.67	8350.56	69.4	30.6
228	91335.62	2544.48	75459.02	7653.22	75.7	24.3
240	96181.02	3901.77	79468.79	10038.55	73.1	26.9
252	101026.43	2185.27	83478.58	9988.35	69.2	30.8
263	105390.13	3292.91	87077.48	9350.07	75.1	24.9
275	110235.53	2253.58	91087.55	7441.45	78.7	21.3
288	115417.23	3953.44	95362.62	12636.56	69.8	30.2
300	120262.64	3163.32	99380.09	10283.49	74.6	25.4
311	124626.33	1411.85	102971.55	7860.88	77.1	22.9
336	134653.43	3122.30	111256.31	17635.42	61.7	38.3




Figure 10: Architecture for integrating the energy-optimal timetable generator with an industrial timetable compiler. The CSCPP interface sends prepared network and timetable data to the generator and returns the resulting energy-optimal timetable to the compiler.

For the purpose of this integration framework, the existing timetable compiler is partitioned into three discrete modules:

- *Timetable data*: This module houses the necessary data to solve the optimization problem. Additionally, it includes other information specific to the railway network under consideration, such as gradient attributes and speed limitations in various track segments, all presented in raw form.
- *Optimization class*: This module transforms the raw data into a structure that is compatible with the optimization algorithm we employ.
- *CSCPP interface*: This interface is the compiler’s data-exchange layer; it provides bidirectional communication between the industrial timetable compiler and the energy-optimal timetable generator. It sends the prepared network and timetable data to the generator, retrieves the energy-optimal timetable, and converts the result into a format suitable for deployment in a CBTC system.

The energy-optimal timetable generator is similarly sub-divided into three modules:

- *Type definitions*: This module contains the data structures essential for describing a railway network, as described in Sections 3 and 5.
- *Utility functions*: This module constructs the optimization model (OPT_MODEL), predicts effective energy consumption via (PRED_EFF_ENG_CNSM), and transforms the output timetable into a CSCPP-compatible representation.
- *Optimization model and solver*: This module contains the linear program to solve (OPT_MODEL). The resulting timetable is returned to the CSCPP interface through the utility functions.

8 Conclusion

We presented a single-stage, data-driven linear programming model that designs energy-optimal timetables for CBTC-operated metro railway networks by minimizing effective energy consumption: the total traction energy minus the regenerative braking energy transferred from braking trains to accelerating trains. The model embeds the transfer of regenerative energy in linear hypograph constraints and predicts the effective energy consumption of a timetable without physics-based simulation. On 11 full-day operational instances of Shanghai Metro Line 8 with up to 1,332 active trains, it computed energy-optimal timetables in under 0.7 seconds per instance and predicted effective-energy reductions of 19.52%–28.17% relative to the existing operational timetables, in close agreement with the physics-based simulator SPSIM. On 17 Docklands Light Railway real-world instances with up to 336 active trains, it solved in under 0.11 seconds per instance and predicted reductions of 19.33%–28.82%.

Our model has some limitations. The affine models of traction and regenerative energy are justified only over the narrow trip-time windows of Assumption 4.1 that is typical of CBTC systems; however substantially wider windows that is common in inter-city travels may require nonlinear fits. The FWHM rectangles approximate the power-versus-time profiles conservatively, so the model tends to underestimate transferred regenerative energy relative to SPSIM’s exact overlap computation. Also, the Docklands Light Railway predictions could not be validated against a physics-based simulator such as SPSIM due to lack of appropriate data required for the simulator for Docklands Light Railway.

The model is being integrated into Hitachi Rail Canada’s industrial timetable compiler through the architecture described in Section 7. Natural directions for future work include extending the

model to handle stochastic disturbances and delay recovery, and incorporating passenger-demand information into the timetable design.

References

- Montserrat Abril, Federico Barber, Laura Ingolotti, Miguel A. Salido, Pilar Tormos, and Antonio Lova. An assessment of railway capacity. *Transportation Research Part E: Logistics and Transportation Review*, 44(5):774–806, 2008.
- Amie Albrecht, Phil Howlett, Peter Pudney, Xuan Vu, and Peng Zhou. The key principles of optimal train control—Part 1: Formulation of the model, strategies of optimal type, evolutionary lines, location of optimal switching points. *Transportation Research Part B: Methodological*, 94: 482–508, 2016.
- James F Allen. Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26(11):832–843, 1983.
- Andreas Bärmann, Alexander Martin, and Oskar Schneider. Efficient formulations and decomposition approaches for power peak reduction in railway traffic via timetabling. *Transportation Science*, 55(3):747–767, 2021.
- Stephen Boyd and Lieven Vandenbergh. *Convex Optimization*. Cambridge University Press, 2004.
- David Canca and Alejandro Zarzo. Design of energy-efficient timetables in two-way railway rapid transit lines. *Transportation Research Part B: Methodological*, 102:142–161, 2017.
- Jiaxin Cheng and Phil Howlett. A note on the calculation of optimal strategies for the minimization of fuel consumption in the control of trains. *IEEE Transactions on Automatic Control*, 38(11): 1730–1734, 1993.
- Shuvomoy Das Gupta, Lacra Pavel, and James Kevin Tobin. An optimization model to utilize regenerative braking energy in a railway network. In *American Control Conference (ACC)*, pages 5919–5924. IEEE, 2015.
- Shuvomoy Das Gupta, James Kevin Tobin, and Lacra Pavel. A two-step linear programming model for energy-efficient timetables in metro railway networks. *Transportation Research Part B: Methodological*, 93:57–74, 2016.
- George Gamow. *Thirty Years that Shook Physics: The Story of Quantum Theory*. Dover Publications, 1985.
- Patrick Gemander, Andreas Bärmann, and Alexander Martin. A stochastic optimization approach to energy-efficient underground timetabling under uncertain dwell and running times. *Transportation Science*, 57(6):1627–1645, 2023.
- Arturo González-Gil, Roberto Palacin, and Paul Batty. Sustainable urban rail systems: Strategies and technologies for optimal management of regenerative braking energy. *Energy Conversion and Management*, 75:374–388, 2013.
- Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, Version 13.0, 2026. URL <https://docs.gurobi.com/projects/optimizer/en/current/>.
- Hitachi Rail. Hitachi rail acquires Thales’ ground transportation systems for EUR 1,660m. <https://www.hitachi.com/en-us/press/>

- hitachi-rail-acquires-thales-ground-transportation-systems/, May 2024. Press release, May 31, 2024; accessed July 15, 2026.
- Hitachi Rail. SelTrac communications-based train control for urban rail. <https://www.hitachirail.com>, 2026. Formerly a Thales Ground Transportation Systems product; accessed July 15, 2026.
- Phil Howlett. The optimal control of a train. *Annals of Operations Research*, 98:65–87, 2000.
- Phil Howlett, Peter Pudney, and Amie Albrecht. Optimal driving strategies for a fleet of trains on level track with prescribed intermediate signal times and safe separation. *Transportation science*, 57(2):399–423, 2023.
- Phil G Howlett, Peter J Pudney, and Xuan Vu. Local energy minimization in optimal train control. *Automatica*, 45(11):2692–2698, 2009.
- Philip G Howlett and Peter J Pudney. *Energy-efficient train control*. Springer, 1995.
- International Association of Public Transport (UITP). Statistics Brief: World Metro Figures 2021. <https://www.uitp.org/publications/metro-world-figures-2021/>, 2022. Brussels; rider-ship data reported for 2019; accessed July 18, 2026.
- IRSE International Technical Committee. Semi-automatic, driverless and unattended operation of trains. ITC Report 06-02, Institution of Railway Signal Engineers, 2009. <https://webinfo.uk/webdocssl/irse-kbase/ref-viewer.aspx?refno=1559377145&document=itc%20report%2006-02%20semi-automatic,%20driverless%20and%20unattended%20operation%20of%20trains.pdf>, accessed July 15, 2026.
- Eugene Khmelnitsky. On an optimal control problem of train operation. *IEEE Transactions on Automatic Control*, 45(7):1257–1266, 2000.
- Zhao Le, Keping Li, Jingjing Ye, and Xiaoming Xu. Optimizing the train timetable for a subway system. *Proceedings of the Institution of Mechanical Engineers, Part F: Journal of Rail and Rapid Transit*, 229(8):852–862, 2015. doi: 10.1177/0954409714524377.
- Xiang Li and Hong K Lo. An energy-efficient scheduling and speed control approach for metro rail operations. *Transportation Research Part B: Methodological*, 64:73–89, 2014a.
- Xiang Li and Hong K Lo. Energy minimization in dynamic train scheduling and control for metro rail operations. *Transportation Research Part B: Methodological*, 70:269–284, 2014b.
- Pei Liu, Lixing Yang, Ziyong Gao, Yeran Huang, Shukai Li, and Yuan Gao. Energy-efficient train timetable optimization in the subway system with energy storage devices. *IEEE Transactions on Intelligent Transportation Systems*, 19(12):3947–3963, 2018.
- Rongfang Rachel Liu and Iakov M Golovitcher. Energy-efficient operation of rail vehicles. *Transportation Research Part A: Policy and Practice*, 37(10):917–932, 2003.
- Xiaojie Luan, Yihui Wang, Bart De Schutter, Lingyun Meng, Gabriel Lodewijks, and Francesco Corman. Integration of real-time traffic management and train control for rail networks—Part 2: Extensions towards energy-efficient train operations. *Transportation Research Part B: Methodological*, 115:72–94, 2018.
- Miles Lubin, Oscar Dowson, Joaquim Dias Garcia, Joey Huchette, Benoît Legat, and Juan Pablo

- Vielma. JuMP 1.0: Recent improvements to a modeling language for mathematical optimization. *Mathematical Programming Computation*, 15:581–589, 2023. doi: 10.1007/s12532-023-00239-3.
- Sanjoy Mahajan. *Street-Fighting Mathematics: The Art of Educated Guessing and Opportunistic Problem Solving*. The MIT Press, 2010.
- Ian P Milroy. *Aspects of automatic train control*. PhD thesis, Loughborough University, 1980.
- Ahmed Mohamed, Andrew Reid, and Thomas Lamb. White paper on wayside energy storage for regenerative braking energy recuperation in the electric rail system. White paper, Con Edison, New York, NY, USA, 2018.
- Robert D Pascoe and Thomas N Eichorn. What is communication-based train control? *IEEE Vehicular Technology Magazine*, 4(4):16–21, 2009.
- Maite Peña-Alcaraz, Mort Webster, and Andrés Ramos. An approximate dynamic programming approach for designing train timetables. Working paper, Massachusetts Institute of Technology, Engineering Systems Division, 2011. <https://dspace.mit.edu/handle/1721.1/102831>.
- Maite Peña-Alcaraz, Antonio Fernández, Asuncion Paloma Cucala, Andres Ramos, and Ramon R Pecharromán. Optimal underground timetable design based on power flow for maximizing the use of regenerative-braking energy. *Proceedings of the Institution of Mechanical Engineers, Part F: Journal of Rail and Rapid Transit*, 226(4):397–408, 2012.
- Gerben M. Scheepmaker, Rob M. P. Goverde, and Leo G. Kroon. Review of energy-efficient train control and timetabling. *European Journal of Operational Research*, 257(2):355–376, 2017.
- Shanghai Shentong Metro Group. Line 8 route information (in Chinese). <http://service.shmetro.com/axlcz08/index.htm>, 2026. Official Shanghai Metro service portal; accessed July 18, 2026.
- Lingyuan Shi, Xin Yang, Wei Zhang, Huijun Sun, and Gilbert Laporte. Integrated optimization of train rescheduling and speed management during partial blockages: A decomposition approach. *Transportation Science*, 60(3):484–507, 2026.
- Shuai Su, Xiang Li, Tao Tang, and Ziyu Gao. A subway train timetable optimization approach based on energy-efficient operation strategy. *IEEE Transactions on Intelligent Transportation Systems*, 14(2):883–893, 2013.
- Shuai Su, Xuekai Wang, Yuan Cao, and Jiateng Yin. An energy-efficient train operation approach by integrating the metro timetabling and eco-driving. *IEEE Transactions on Intelligent Transportation Systems*, 21(10):4252–4268, 2019.
- United Nations. World urbanization prospects 2025: Summary of results. Technical Report UN DESA/POP/2025/TR/NO. 12, United Nations, Department of Economic and Social Affairs, Population Division, 2025. <https://desapublications.un.org/publications/world-urbanization-prospects-2025-summary-results>, accessed September 14, 2026.
- Pengling Wang and Rob MP Goverde. Multi-train trajectory optimization for energy efficiency and delay recovery on single-track railway lines. *Transportation Research Part B: Methodological*, 105:340–361, 2017.
- Pengling Wang and Rob MP Goverde. Multi-train trajectory optimization for energy-efficient timetabling. *European Journal of Operational Research*, 272(2):621–635, 2019.

- Pengling Wang, Nikola Bešinović, Rob MP Goverde, and Francesco Corman. Improving the utilization of regenerative energy and shaving power peaks by railway timetable adjustment. *IEEE Transactions on Intelligent Transportation Systems*, 23(9):15742–15754, 2022.
- Songpo Yang, Feixiong Liao, Jianjun Wu, Harry JP Timmermans, Huijun Sun, and Ziyu Gao. A bi-objective timetable optimization model incorporating energy allocation and passenger assignment in an energy-regenerative metro system. *Transportation Research Part B: Methodological*, 133: 85–113, 2020.
- Songpo Yang, Feixiong Liao, Jianjun Wu, and Yanyan Chen. An efficient train timetable scheduling approach with regenerative-energy supplementation strategy responding to potential power interruptions. *IEEE Transactions on Intelligent Transportation Systems*, 23(9):14267–14282, 2021.
- Xin Yang, Xiang Li, Ziyu Gao, Hongwei Wang, and Tao Tang. A cooperative scheduling model for timetable optimization in subway systems. *IEEE Transactions on Intelligent Transportation Systems*, 14(1):438–447, 2013.
- Xin Yang, Xiang Li, Bin Ning, and Tao Tang. A survey on energy-efficient train operation for urban rail transit. *IEEE Transactions on Intelligent Transportation Systems*, 17(1):2–13, 2016.
- Jiateng Yin, Lixing Yang, Tao Tang, Ziyu Gao, and Bin Ran. Dynamic passenger demand oriented metro train scheduling with energy-efficiency and waiting time minimization: Mixed-integer linear programming approaches. *Transportation Research Part B: Methodological*, 97:182–213, 2017.
- Huiru Zhang, Limin Jia, Li Wang, Xinyue Xu, and Fei Dou. Energy-efficient timetable optimization empowered by time-energy pareto solution under actual line conditions. *IEEE Transactions on Intelligent Transportation Systems*, 25(7):7106–7124, 2024.