sustainability

MDPI

Article

Efficiency and Running Time Robustness in Real Metro Automatic Train Operation Systems: Insights from a Comprehensive Comparative Study

María Domínguez, Adrián Fernández-Rodríguez *, Asunción P. Cucala and Antonio Fernández-Cardador

Institute for Research in Technology (IIT), ICAI School of Engineering, Comillas Pontifical University, 23 Alberto Aguilera St., 28015 Madrid, Spain; maria.dominguez@iit.comillas.edu (M.D.); paloma.cucala@iit.comillas.edu (A.P.C.); antonio.fernandez@iit.comillas.edu (A.F.-C.)

* Correspondence: adrian.fernandez@iit.comillas.edu

Abstract

Automatic Train Operation (ATO) systems are widely deployed in metro networks to improve punctuality, service regularity, and ultimately the sustainability of rail operation. Although eco-driving optimisation has been extensively studied, no previous work has provided a systematic, side-by-side comparison of the two ATO control philosophies most commonly implemented in metro systems worldwide: (i) Type 1, based on speed holding followed by a single terminal coasting at a kilometre point, and (ii) Type 2, which uses speed thresholds to apply either continuous speed holding or iterative coasting–remotoring cycles. These strategies differ fundamentally in their control logic and may lead to distinct operational and energetic behaviours. This paper presents a comprehensive comparison of these two ATO philosophies using a high-fidelity train movement simulator and Pareto-front optimisation via a multi-objective particle swarm algorithm. 40 interstations of a real metro line were evaluated under realistic comfort and operational constraints, and robustness was assessed through sensitivity to three different passenger-load variations (empty train, nominal load and full load). Results show that, once nominal profiles are implemented, Type 1 has up to 5% variability in running times, and Type 2 has up to 20% variability in energy consumption. In conclusion, a new ATO deployment combining both strategies could better balance energy efficiency and timetable robustness in metro operations.

Keywords: automatic train operation; ATO; energy optimisation; metro systems; railway systems; energy efficiency; speed profile design; train simulation; eco-driving; robustness; sustainable transport

Received: 21 November 2025
Revised: 9 December 2025
Accepted: 16 December 2025
Published: 18 December 2025

Citation: Domínguez, M.; Fernández-Rodríguez, A.; Cucala, A.P.; Fernández-Cardador, A. Efficiency and Running Time Robustness in Real Metro Automatic Train Operation Systems: Insights from a Comprehensive Comparative Study. *Sustainability* **2025**, *17*, 11371. https://doi.org/10.3390/su172411371

1. Introduction

Automatic Train Operation (ATO) systems are increasingly adopted in metropolitan and suburban rail networks as a means to enhance operational regularity, service quality, and energy efficiency [1]. These improvements not only contribute directly to the sustainability of rail transport but also strengthen its role as an attractive alternative to less sustainable modes. By increasing the quality of service and reducing energy consumption, ATO-equipped systems can encourage a modal shift toward rail, thereby amplifying their environmental and societal benefits.

The integration of ATO technologies with modern signalling and protection systems contributes to improved safety and reduced human-induced variability, allowing more

predictable running times, higher transport capacity, and reduced energy consumption through systematic eco-driving [2,3]. Although ATO has traditionally been associated with urban metros, equivalent automated driving strategies are now being explored in high-speed [4–7] and even freight operations [8], reflecting a broader industry trend towards leveraging automation to improve operational efficiency across multiple railway domains. Consequently, ATO has evolved into a key driving technology within metro systems, while simultaneously expanding its relevance to other railway applications.

From a functional viewpoint, ATO is an onboard controller that receives real-time information from the Automatic Train Protection (ATP) and Automatic Train Supervision (ATS) systems, such as speed limits, movement authority, gradient profile, train speed and position, and computes the traction and braking commands required to meet a target running time while ensuring safety and comfort [9]. The ATP is responsible for guaranteeing that safety constraints are never violated and must therefore always be present. The ATS, in contrast, may or may not exist, depending on the line's level of automation. When available, it manages the overall traffic flow and is responsible for headway regulation and timetable adherence.

The ATO executes predefined speed profiles that are typically determined offline. In conventional metropolitan implementations, each interstation is associated with a small set of four predefined driving profiles: a flat-out (minimum-time) profile, a nominal one, and two additional slower profiles. The traffic regulation system selects one of these profiles depending on real-time headway and timetable constraints and sends it to the train for execution [10].

Each speed profile is internally defined through ATO commands: a constant-speed (holding) value, a coasting triggering speed or position, a remotoring threshold, and a braking deceleration. Maximum acceleration is always applied when required, whereas coasting cancels propulsion until speed falls to the remotoring threshold. This simple yet effective command structure produces a wide family of possible speed profiles, each characterised by a unique pair of running time and energy consumption. The design of these profiles is therefore a multiobjective optimisation problem aimed at minimising energy while meeting a target running time provided by the timetable.

1.1. Eco-Driving and the State of the Art

The search for optimal energy-efficient train trajectories dates back to the analytical formulation of Ichikawa in 1968 [11], based on the Pontryagin Maximum Principle. Since then, extensive research has addressed the optimal eco-driving problem through analytical [12,13] and numerical methods [14]. Although the theoretical optimal speed profile has traditionally been considered one that combines holding and coasting stages, ATO systems may or may not be technically capable of executing the optimal coasting or speed holding commands. Consequently, most analytical solutions, as well as many numerical ones, cannot be directly applied to real ATO systems due to the discrete nature of their control logic and hardware constraints.

Within practical ATO-oriented optimisation, two main methodological lines can be distinguished. The first relies on exhaustive simulation of all feasible combinations of discrete ATO possible command values, an approach particularly suited to fixed-block signalling systems where the limited bandwidth of balises or track circuits restricts the number of command values that can be transmitted to the train [15]. Under these conditions, the command space remains sufficiently small to allow full enumeration and evaluation of all candidate speed profiles.

The second approach corresponds to multiobjective optimisation methods, required in systems using Communications-Based Train Control (CBTC) [16]. Unlike fixed-block

signalling, CBTC relies on continuous, high-bandwidth, bidirectional radio communication, enabling the use of fine-resolution, almost continuous, ATO command values. In these cases, the solution space grows exponentially, and exhaustive search becomes infeasible. Genetic Algorithms [17–19], Tabu Search [20], Fuzzy Control [15,21], Ant Colony [22,23] and Multi-Objective Particle Swarm Optimisation (MOPSO) [24,25] have demonstrated strong performance in generating Pareto-optimal speed profiles under realistic ATO constraints.

1.2. Robustness: The Overlooked Effect of Train Mass

Another issue frequently overlooked in the literature is the robustness of the implemented set of speed profiles to changes in operating conditions. The robustness of an ATO speed profile is the capacity to maintain its running time and energy consumption goals when the operating conditions change. It can be measured as the variation in the resulting running time and energy consumption under the most unfavourable conditions with respect to the nominal design. For example, passenger load variations substantially alter the train mass throughout the day. Mass changes affect acceleration, coasting decay rates, maximum feasible coasting distance, braking distances, and ultimately both energy consumption and running time. Although mass variability has been identified as one of the major sources of uncertainty in trajectory planning [26], most eco-driving optimisation studies, with a few exceptions [27,28], do not incorporate its effects. Some other studies propose robust optimisation formulations where mass variability is explicitly modelled [26,29], but these frameworks have not yet been applied to the context of real ATO system constraints.

Given that ATO profiles are pre-programmed and then executed under varying load conditions, robustness to mass variations becomes essential to avoid deviations in travel time and excessive energy penalties under different occupancy situations.

1.3. ATO Control Philosophies

In practice, two distinct families of ATO technology exist, which have been named as Type 1 and Type 2 for this study:

- Type 1: systems that employ a single speed-holding phase followed by a coasting period at a specific kilometre point before braking.
- Type 2: systems that rely either on continuous speed holding throughout the interstation or on iterative coasting–re-motoring cycles.

These two logics are mutually exclusive, as their onboard control architectures and parameterisation differ fundamentally. As a result, they generate different operational behaviours, different feasible sets of driving strategies, and different Pareto fronts when evaluated in realistic metro environments.

Liu et al. [30] examine both strategies analytically across generic gradient patterns and derive structural optimality properties for continuous control regimes. However, they used continuous controls, no discrete ATO commands, no comfort filters and no mass variation.

The present work addresses exactly this gap by comparing real ATO implementations rather than theoretical driving regimes, assessing how the two commercial philosophies behave once translated into actual ATO command structures under realistic metro infrastructure.

1.4. Objectives and Contributions

This paper presents a systematic, simulation-based comparison of the two existing ATO strategies in metropolitan systems (Type 1 and Type 2) under realistic conditions, focusing on their operational behaviour, energy performance, and robustness to variations in train mass. To that end, a high-fidelity comparative methodology integrating detailed train dynamics, ATO control logic and infrastructure constraints is used. Then, a Pareto-

front methodology based on MOPSO is applied to generate feasible speed profiles that meet comfort and operational constraints for each ATO type.

The main contributions are:

- A detailed comparative analysis identifying the interstation characteristics (gradients, speed limits, length) under which each ATO strategy performs optimally.
- A robustness study with respect to train mass, evaluating how load variability affects both energy and travel time for each ATO strategy, an aspect largely neglected in the previous literature.

Overall, this work provides a comprehensive comparison of the two ATO control philosophies currently implemented in commercial systems, offering new insights for operators, manufacturers, and signalling providers on how to tune or select ATO strategies to balance energy efficiency, punctuality, and real-world robustness.

The remainder of this paper is organised as follows. Section 2 describes the methodology, including the high-fidelity simulator, the ATO control models, and the multi-objective optimisation process used to generate Pareto-optimal speed profiles. Section 3 presents the characteristics of the metro line selected as a case study. Section 4 reports the results of the comparative analysis, examining both the operational performance of the two ATO strategies and their sensitivity to train mass variability. Section 5 discusses the implications of these findings in the context of sustainable metro operation and ATO design. Finally, Section 6 summarises the main conclusions and outlines future research directions.

2. Methodology

This study adopts a simulation-based approach to evaluate and compare the performance of the two distinct ATO strategies named Type 1 and Type 2 under realistic metro operating conditions. A detailed train movement simulator [25] has been employed, capable of emulating the behaviour of an ATO system by reproducing the full sequence of driving commands. Based on the selected ATO parameters, the simulator generates the resulting speed profile between each pair of consecutive stations.

This section is organised as follows: First, the simulation model is described, including its modular architecture and the implementation of its core components: the two ATO system types, the motor model, the dynamics model, and the energy consumption model. Secondly, the Pareto-based optimisation framework used to evaluate trade-offs between energy efficiency and travel time is introduced.

2.1. Simulation Model

As previously mentioned, the design of speed profiles is governed by two main criteria: travel time and energy consumption. Therefore, accurate simulation of both variables is essential. It must be noted that, for each interstation run, the difference between the four predefined ATO speed profiles is typically only a matter of a few seconds. Consequently, calculation errors must be kept at a significantly lower order of magnitude to ensure reliability.

To achieve the required accuracy and facilitate proper validation and calibration, the simulator adopts a modular structure comprising four interconnected submodules: the ATO module, the motor module, the dynamics module, and the consumption module. Each submodule represents a different subsystem of a real train and executes specific calculations based on predefined input parameters, as schematically illustrated in Figure 1. The simulation model and results were benchmarked against operational data from metro lines employing similar ATO systems. This step ensured the robustness and practical relevance of the simulated outcomes, particularly in terms of energy savings and service

regularity. The different submodules are briefly explained below; for a more detailed description, refer to [25].

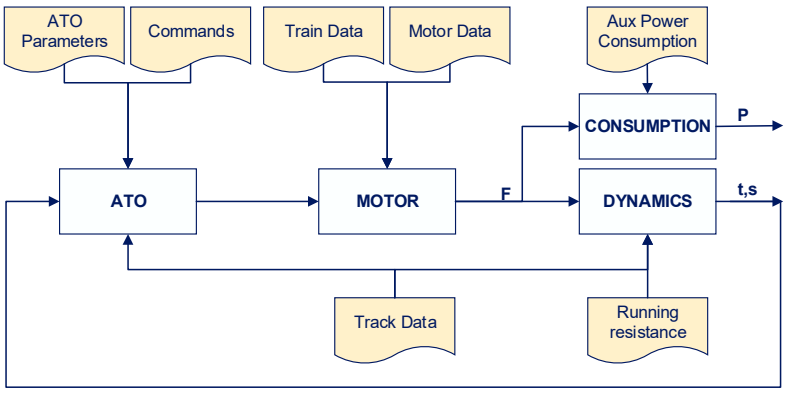

Figure 1. Structure of the simulation framework.

2.1.1. ATO Model

The ATO module emulates the behaviour of real onboard equipment and operates as a proportional controller that computes, at each simulation step, a traction or braking command a_{ATO}. This value represents the ratio between the requested traction/braking effort and the maximum available effort from the train.

As schematically shown in Figure 1, the ATO receives as inputs the current train speed and position, together with infrastructure data such as the maximum speed profile, the gradient profile, and the station stopping point. It also incorporates internal control cycles and hysteresis effects that are characteristic of real ATO equipment. Based on this information, the controller distinguishes between three operational modes: traction mode, braking mode, and final braking mode (approaching station), described below.

1. Traction Mode

In this mode, the train is accelerating or running at a constant speed and does not yet need to decelerate due to a speed restriction or the upcoming station. The ATO computes the traction demand using the proportional control law in Equation (1):

$$a_{ATO} = K_t \cdot (v_t - v) + a_g \quad (1)$$

where

- K_t is a proportional gain for traction that assures that the target speed is achieved in a certain period of time (s^{-1});
- v is the current train speed (m/s);
- v_t is the target speed, defined as the minimum between the maximum permitted line speed and the active holding-speed command;
- a_g is a feedforward acceleration term compensating the gradient (m/s^2).

2. Braking Mode (Speed Restriction)

At each simulation step, the ATO checks whether a speed reduction is required using a braking detection curve [25]. When the detection curve intersects the maximum speed profile, braking mode is triggered.

The target speed becomes the value of the braking curve one step ahead, and the ATO computes the command using Equation (2):

$$a_{ATO} = K(v) \cdot \left(v_{ref} - v\right) - a_{ref} + a_g \quad (2)$$

where

- $K(v)$ is a speed-dependent proportional gain;
- v_{ref} is the reference speed from the braking curve;
- a_{ref} is the target deceleration of the braking curve to reach v_{ref}.

3. Final Braking Mode (Approaching Station)

Before each simulation, a final braking curve is generated from the station stopping point using the designated deceleration command b. When the train intercepts this curve, the final-braking mode becomes active.

The ATO uses the same control structure as in Equation (2) with $a_{ref} = b$ and v_{ref} given by the final braking curve at the next simulation step.

After computing a_{ATO}, the command is saturated within the interval $[-1, 1]$. Additional logic is applied when operating with coasting/remotoring commands: once the train reaches the coasting speed cs, and as long as it remains above the remotoring threshold rm, the maximum allowed command is limited to zero (pure coasting).

2.1.2. Motor Model

The motor model determines the traction and braking effort required to track the setpoint generated by the ATO module. For this purpose, it incorporates the main mechanical characteristics of the rolling stock, including total mass, train length, and rotational inertia. Traction and braking capabilities are represented through the predefined force–speed characteristics of the motors, which define the maximum available effort at each operating speed.

At every simulation step, the ATO output a_{ATO} is interpreted by the motor controller as a fraction of the maximum admissible force. The requested traction or braking effort is then computed as Equation (3):

$$F_{train} = a_{ATO} \cdot F_{max}(v) \tag{3}$$

where $F_{max}(v)$ represents the maximum traction force (for $a_{ATO} \geq 0$) or the maximum braking effort (for $a_{ATO} < 0$) according to the motor and braking system performance, which depends on the speed of the train.

To reproduce realistic behaviour, the model incorporates several dynamic aspects of traction equipment. First, it includes the intrinsic time delay between ATO command generation and motor response. Second, a jerk limitation is applied, since excessive jerk values produce sudden changes in acceleration that can lead to noticeable discomfort for passengers and are therefore unacceptable from a comfort perspective.

Together, these elements enable the motor module to reproduce closely the physical behaviour of real rolling stock under diverse operating conditions, providing accurate traction, braking, and current profiles for subsequent consumption calculations.

2.1.3. Dynamics Model

The dynamics model updates the train's kinematic state (acceleration, speed, and position) at each simulation step by solving the longitudinal equations of motion. The model accounts for all forces acting on the train, including the traction or braking effort provided by the motor module and the resistances arising from the infrastructure and operating conditions.

The total resistance to motion is computed using the classical Davis equation [31], which accounts for rolling resistance, mechanical friction, and aerodynamic drag. It is shown in Equation (4):

$$F_r = A + B \cdot v + C \cdot v^2 \tag{4}$$

where

- F_r is the total running resistance (N);
- v is the train speed (m/s);
- A, B, C are empirically determined coefficients accounting for rolling resistance, mechanical drag, and aerodynamic drag, respectively.

In addition to Davis' formulation, the model includes the effect of track gradients, represented as equivalent additional forces as in Equation (5):

$$F_g = mg\sin(\theta) \tag{5}$$

where

- θ is the local gradient angle;
- m is the total train mass (kg).

Once all forces are computed, the instantaneous acceleration a is obtained from Newton's second law in Equation (6). Speed and position are then updated using uniformly accelerated motion equations (Equations (7) and (8))

$$a = \frac{F_{train} - F_r - F_g}{m} \tag{6}$$

$$a = \frac{dv}{dt} \tag{7}$$

$$v = \frac{dx}{dt} \tag{8}$$

This detailed dynamics model provides an accurate reconstruction of acceleration, coasting, and braking phases across a wide range of track geometries, enabling a reliable evaluation of the speed profiles generated by the two ATO strategies under study.

2.1.4. Consumption Model and Motor Efficiency

The consumption model estimates the electrical energy demand of the train by combining the mechanical power calculated in the motor module with the kinematic behaviour obtained from the dynamics module. At each simulation timestep, the instantaneous electrical power is computed as Equation (9):

$$P(t) = F_{train}(t) \cdot v(t) \tag{9}$$

Auxiliary onboard consumption is also incorporated, ensuring a complete representation of the total energy demand.

A key modelling challenge lies in accurately representing motor efficiency. Full efficiency maps, detailing the variation in efficiency with both speed and torque, are rarely available for metro rolling stock. In most cases, only the current–speed curve associated with maximum traction effort is supplied. A simple approach is to scale this curve linearly according to the ratio between the requested and maximum force, implicitly assuming constant efficiency across the load range. However, such an assumption is unrealistic, since traction motors exhibit significant efficiency degradation when operating at partial load.

To overcome this limitation, the consumption model applies a corrective factor that penalises electrical consumption when the traction effort is below its maximum value. This adjustment captures the efficiency drop characteristic of partial-load operation, leading to a more realistic estimation of energy use, particularly for speed profiles with prolonged low-power cruising.

This behaviour is consistent with both theoretical expectations and empirical evidence and must therefore be explicitly incorporated when comparing ATO control strategies from an energy-efficiency perspective.

2.2. ATO Control Strategies

The ATO module supports the two driving logics analysed in this study. Although both rely on the same proportional control framework described above, each strategy imposes different restrictions on how coasting and traction phases may be combined. In real metropolitan deployments worldwide, these logics are mutually exclusive and are configured at the system level.

- ATO Type 1—Constant-speed holding with final coasting: In this approach, the train accelerates to a target holding speed and maintains it up to a predefined kilometre point. Upon reaching this point, the train transitions into a final coasting phase, during which no traction is applied, allowing the train to decelerate naturally due to resistance forces. Braking is then applied at the appropriate location to stop precisely at the next station. An illustrative example of this speed profile is shown in Figure 2.
- ATO Type 2—Constant-speed holding or iterative coasting–remotoring: This control logic allows the train to either hold a constant speed throughout the interstation when permitted by speed limits, or perform a sequence of coasting–remotoring cycles. When the latter strategy is used, after initial acceleration, the train coasts; when its speed drops below a predefined threshold, a remotoring phase restores speed. This cycle may occur multiple times, subject to comfort constraints and timing requirements, before the final braking phase. A representative example is shown in Figure 2.

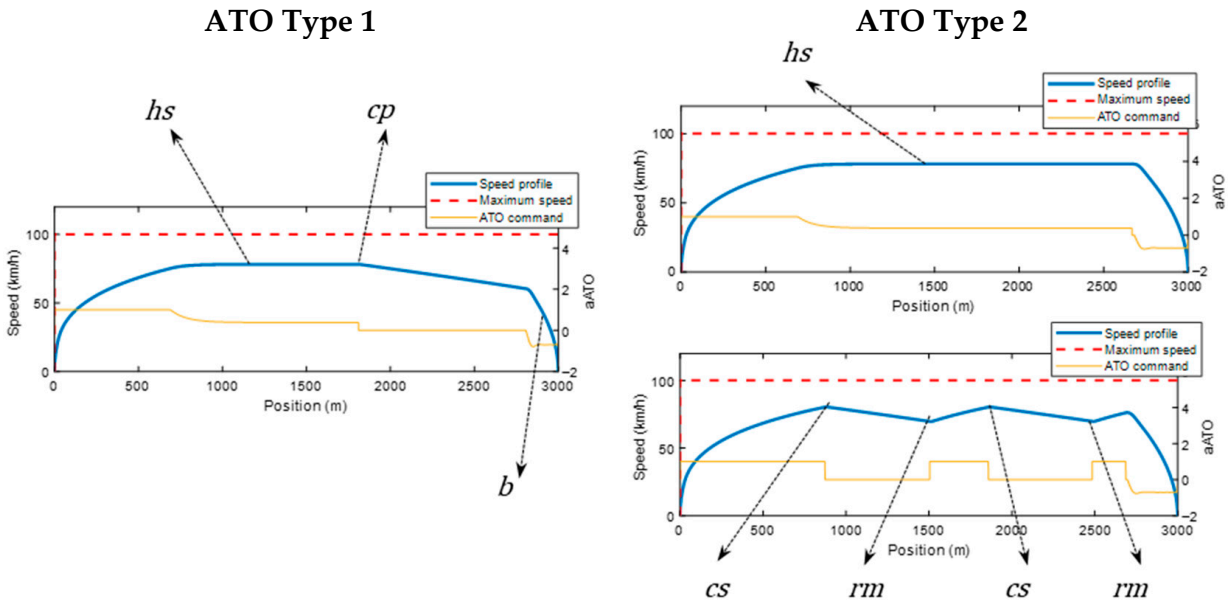

Figure 2. Representative speed profiles generated by the two ATO control strategies.

2.3. Speed Profile Optimisation

The optimisation of ATO speed profiles is a multi-objective problem in which travel time and energy consumption must be jointly minimised while ensuring that the resulting profiles remain feasible under real-world operating conditions. This type of optimisation framework has already been implemented in several metropolitan rail systems, such as those of Madrid, Barcelona, Valencia, and Bilbao in Spain, where ATO-equipped lines rely on precomputed, energy-efficient speed profiles.

To systematically explore the solution space and extract only those profiles that offer optimal performance, a Pareto-based optimisation framework is employed. This framework integrates comfort and operational restrictions directly into the search process, ensuring that only implementable profiles are retained. The following subsections describe the constraints applied and the optimisation procedure used to obtain the Pareto fronts for both ATO strategies.

2.3.1. Comfort Restrictions

To ensure that the optimised speed profiles are not only energy and time-efficient but also operationally viable from a passenger comfort perspective, a set of comfort and operational constraints was incorporated into the optimisation framework. While most studies either overlook comfort considerations or limit them to jerk constraints alone, real-world ATO design requires a broader set of criteria. In practice, an energetically optimal speed profile is of little value if it is not acceptable to passengers or operators. For this reason, comfort-related requirements are typically defined jointly with the operator and systematically verified during the design and validation stages of ATO driving strategies. The following constraints adopted here therefore reflect the practical limitations and user-centred considerations inherent to real metro operations:

- Minimum operating speed: Allowing the train speed to fall excessively during a coasting phase is perceived as uncomfortable for passengers. Therefore, any profile in which speed drops below a prescribed minimum threshold during coasting is rejected.
- Maximum number of remotoring events: Excessive alternation between coasting and remotoring leads to frequent traction transitions, which may reduce comfort and increase wear on traction equipment. A limit on the number of allowable remotoring cycles per interstation is therefore imposed.
- Minimum duration of each driving mode: To avoid abrupt and rapid oscillations between coasting and remotoring, typically associated with discomfort and unnecessary control effort, a minimum time in each driving state is required before transitioning to another mode.

Any candidate speed profile violating one or more of these conditions, regardless of its position in the energy–time objective space, is deemed operationally invalid and removed from further consideration. These constraints are automatically checked within the simulator at each optimisation iteration. Only speed profiles satisfying all operational and comfort requirements are retained and displayed in the final Pareto fronts.

2.3.2. Multi-Objective Particle Swarm Optimisation Algorithm

As previously introduced, a speed profile is defined by a combination of several parameters, each associated with a specific ATO command. The full set of feasible values for these parameters determines the solution space of admissible speed profiles between two stations. The size and nature of this solution space depend largely on the signalling system in use. In fixed-block signalling systems, the available bandwidth for data transmission is extremely limited, and only a small number of discrete values can be sent to the onboard ATO. As a result, the solution space remains relatively small and can be explored exhaustively using simulation.

In contrast, CBTC systems transmit the same configuration parameters through a high-bandwidth wireless communication channel between the train and the wayside controller. Under these conditions, the ATO parameters are no longer restricted to a small discrete set and can be treated as quasi-continuous [10]. Consequently, the number of potential speed profiles grows exponentially, making exhaustive simulation computationally infeasible. In such cases, multi-objective optimisation algorithms become necessary to identify the set

of Pareto-optimal profiles. In previous work [24], a MOPSO algorithm [32] was compared with NSGA-II [33] and the real Pareto front obtained through exhaustive search. The MOPSO algorithm performed successfully to explore the solution space of ATO systems s. The same methodology has been applied in this study. This time to compute the Pareto fronts associated with both the two ATO control strategies analysed (Type 1 and Type 2).

To formulate the optimisation process, each candidate speed profile is encoded as a particle within the MOPSO framework, where the ATO command parameters constitute the decision variables of the search space.

The i-th particle in the population represents a specific combination of ATO parameters and therefore a concrete speed profile. At iteration t, each particle is encoded as a three or four-dimensional vector corresponding to the four decision variables of the ATO control strategy analysed, such that $\vec{x_i}(t) = (b_i, hs_i, cp_i)$ for Type 1 or $\vec{x_i}(t) = (b_i, hs_i, cs_i, rm_i,)$ for Type 2. Each particle also has an associated velocity vector $\vec{v_i}(t) = (vb_i, vhs_i, vcp_i)$ for Type 1 or $\vec{v_i}(t) = (vb_i, vhs_i, vcs_i, vrm_i)$ for Type 2. The best position previously visited by the particle i, i.e., the one yielding the best objective-function value for that particle, is stored as $\vec{p_i}(t) = (pb_i, phs_i, pcs_i, prm_i,)$. Likewise, the best combination of parameters found by the entire swarm up to iteration t is denoted $\vec{p_g}$.

At each iteration t, the velocity of the particle i is updated based on its previous velocity and the distances to both its personal best (pbest) and the global best (gbest), as shown in Equation (10). The new particle position is then obtained by applying the velocity update as specified in Equation (11).

$$\vec{v_i}(t) = w \cdot \vec{v_i}(t-1) + c_1 \cdot r_1 \cdot \left(\vec{p}_i - \vec{x_i}(t-1)\right) + c_2 \cdot r_2 \cdot \left(\vec{p}_g - \vec{x_i}(t-1)\right) \quad (10)$$

$$\vec{x_i}(t) = \vec{x_i}(t-1) + \vec{v_i}(t)\ i = 1, 2, \ldots\ NP\ and\ t = 1, 2, \ldots\ T \quad (11)$$

where

- NP : Number of particles in the swarm (population size).
- T : Total number of optimisation iterations.
- $\vec{x_i}(t)$: Position of particle i , representing a candidate speed-profile parameter set.
- $\vec{v_i}(t)$: Velocity of particle i .
- $\vec{p_i}(t)$: Personal best (pbest) position of particle i
- $\vec{p_g}(t)$: Global best (gbest) or guide particle selected from the archive.
- w: Inertia weight controlling the influence of previous velocity. It ranges from 0.9 to 0.2.
- r_1, r_2 Random numbers in [0, 1], weighting cognitive and social components.
- c_1 : Positive constant that indicates the maximum influence of the pbest individual experience on its new velocity. A value of 1 has been used.
- c_2: Positive constant that indicates the maximum influence of the gbest on the particle's new velocity. A value of 1 has been selected.

All particles' positions are assessed at each iteration in terms of energy consumption and running time using the Simulation Model described in this section as evaluation function.

The base MOPSO algorithm was enhanced by incorporating a crowding distance (CD) mechanism into the selection of the global best (gbest). The CD of a candidate profile is defined as the average distance to its two nearest neighbours after sorting all non-dominated solutions by objective value. Profiles located at the extremes of the front are assigned an infinite CD to guarantee that edge solutions are always retained. In the resulting MOPSO-CD algorithm, the global guide (gbest) for particle updates is selected stochastically: a probability of 98% (top-select probability) is assigned to solutions belonging to the top 6%

of archive members with the highest CD values, while the remaining 2% is assigned to the low-CD region. This strategy directs the swarm towards sparsely populated areas of the Pareto front, thereby eliminating gaps and enhancing the uniformity of the resulting curve. The pseudocode of the algorithm is shown in Figure 3, where

```
A = ∅                                                  Initialise an empty archive
For i = 1 : NP                                         For the entire population size
    {xi, vi}                                           Randomly initialise the position and velocity of each particle
    yi = f(xi) = [EC(xi), RT(xi)]                      Evaluate particles according to the objective function
    pi = xi                                            Initialise each particle's personal best (pbest)
End For

A = { xi | xi ≺ xj  for i,j = 1,…,NP }                 Create the archive with the non-dominated solutions
pg = Random(A)                                         Generate the global guide from a random element of the archive

For t = 1 : T                                          For T iterations
    For i = 1 : NP                                     For each of the NP particles
        For d = 1 : D                                  For each dimension of the D-dimensional space
            vid = w·vid + r1(pid − xid) + r2(pgd − xid)
                                                       Update particle velocities
            xid = xid + vid                            Update particle positions
        End For

        xi := CheckBounds(xi)                          Ensure boundary constraints are satisfied
        yi := (EC(xi), RT(xi))                         Evaluate particles according to the objective function

        If  xi ≺ u   ∀ u ∈ A                           Check if the particle dominates any archive member
            A = { u ∈ A | u ⊀ xi }                     Remove from the archive all solutions dominated by the new particle
            A = A ∪ xi                                 Add the new dominant particle to the archive
        End If
    End For

    CD(yi) = CD(EC(xi)) + CD(RT(xi))  ∀ i = 1,…,NP     Compute the crowding distance (CD) for each objective

    A = Sort(A)                                        Sort archive solutions in decreasing CD order
    pg = selectGuide(A)                                Update the guide (gbest) by selecting a particle from the top P% of the archive

    If  xi ≤ pi  Then  pi = xi  ∀ i = 1,…,NP           Update pbest for all particles

End For
```

Figure 3. MOPSO algorithm. Source: Authors.

- D: Dimensionality of the search space, i.e., number of ATO parameters defining a speed profile.
- $\vec{y_i}\left(\vec{x_i}\right)$: Objective vector of particle i, consisting of energy consumption $EC\left(\vec{x_i}\right)$ and running time $RT\left(\vec{x_i}\right)$.
- A: External archive containing all current non-dominated (Pareto-optimal) solutions.
- w: Inertia weight controlling the influence of previous velocity. It ranges from 0.9 to 0.2.
- CD: Crowding distance used for diversity preservation in the archive.

3. Case Study

The case study selected for this research is a conventional metro line in a Spanish metropolitan system representative of dense urban rail operations. The line comprises 20 interstations and a total length of approximately 23 km per direction. A defining feature of this corridor is its irregular topography, with gradients reaching up to 40 mm/m, which introduces non-trivial challenges for speed-profile optimisation and energy-efficient ATO. The maximum speed limits are also highly variable, ranging from a minimum of 25 km/h to a maximum of 120 km/h. The average interstation distance is close to 1 km, consistent with typical metropolitan rail configurations.

The comfort and operational restrictions imposed in the study reflect the criteria used in real metro services. In addition to jerk limitation, the following conditions were enforced:

- The minimum allowable speed during any coasting phase is 20 km/h;
- No more than three remotoring events are permitted within an interstation;
- Each driving mode (traction, coasting, speed-holding, remotoring) must last at least 10 s before transitioning to another, preventing abrupt oscillations between modes.

For the comparative assessment of the two ATO systems, the full set of command values defined in Tables 1 and 2 has been assumed. These parameters constitute the decision variables of the optimisation process and are defined as follows:

- *b*: braking to stop deceleration, used to generate the final braking curve.
- *hs*: speed-holding command, representing the target holding speed.
- *cp*: coasting position, expressed as a fraction of the interstation length. For example, $cp = 0.5$ indicates that coasting begins at the midpoint between stations, whereas $cp = 1$ implies that coasting is never applied.
- *cs*: coasting-speed threshold, i.e., the speed at which traction is cancelled, and the train enters a coasting phase.
- *rm*: remotoring-speed threshold, defining the permissible speed drop during coasting before traction is reapplied.

Table 1. Feasible ATO Type 1 command values.

	b **(m/s^2)**		***hs*** **(km/h)**		***cp*** **(Percentage of the Length of the Interstation)**	
	Min	**Max**	**Min**	**Max**	**Min**	**Max**
ATO Type 1	0.6	0.8	30	110	0	100

Table 2. Feasible ATO Type 2 command values.

	b **(m/s^2)**		***hs*** **(km/h)**		***cs*** **(km/h)**		***rm*** **(km/h)**	
	Min	**Max**	**Min**	**Max**	**Min**	**Max**	**Min**	**Max**
ATO Type 2	0.6	0.8	30	110	20	110	5	10

These commands are graphically illustrated in Figure 2, which provides a visual interpretation of their meaning within a typical speed profile.

The minimum and maximum admissible values for each command are determined by operational practice, either imposed by technical limitations of the ATO equipment or by constraints established by the operator. To emulate a CBTC-based system capable of transmitting quasi-continuous ATO parameters, the MOPSO algorithm explores the full interval between these upper and lower bounds for each command, thereby enabling the generation of high-resolution profiles consistent with the capabilities of modern CBTC-equipped metro lines. Each interstation of the line was simulated independently following the methodology described in Section 2, allowing the construction of the corresponding Pareto fronts.

4. Results

Pareto fronts were computed for all interstations of the case study using the methodology described in Section 2. A detailed train movement simulator, implementing both ATO control logics (Type 1 and Type 2) and the full set of feasible ATO command values (Tables 1 and 2), has been used to generate the solution space of speed profiles for each interstation. To isolate the profiles that achieve the best trade-offs between energy consumption and running time, the MOPSO algorithm described in Section 2.3.2 has been applied. Beyond identifying mathematically non-dominated solutions, the workflow enforces the

comfort and operational constraints defined for the case study, ensuring that only feasible and implementable profiles remain on each Pareto front.

The following subsections present the results and main observations drawn from the comparison and analysis of both ATO systems.

4.1. Comparative Analysis of ATO Strategies

This subsection presents a comparative analysis of the two ATO strategies (Type 1 and Type 2), based on the Pareto fronts obtained for every interstation of the case study with the nominal mass (50% load). For each segment between stations, both fronts are plotted together to allow a direct visual and quantitative comparison. The aim is to gain an initial understanding of how each strategy performs in terms of energy efficiency and travel time.

Figure 4 shows the aggregated comparison across all interstations. The solution space is limited at +20 s relative to the flat-out solution (minimum-time, hence maximum-energy). This reflects operational practice: the slowest speed profile selected for off-peak service typically does not exceed the flat-out time by more than 20 s. Profiles beyond this window may exist, but they are rarely selected in timetable design or real-time regulation.

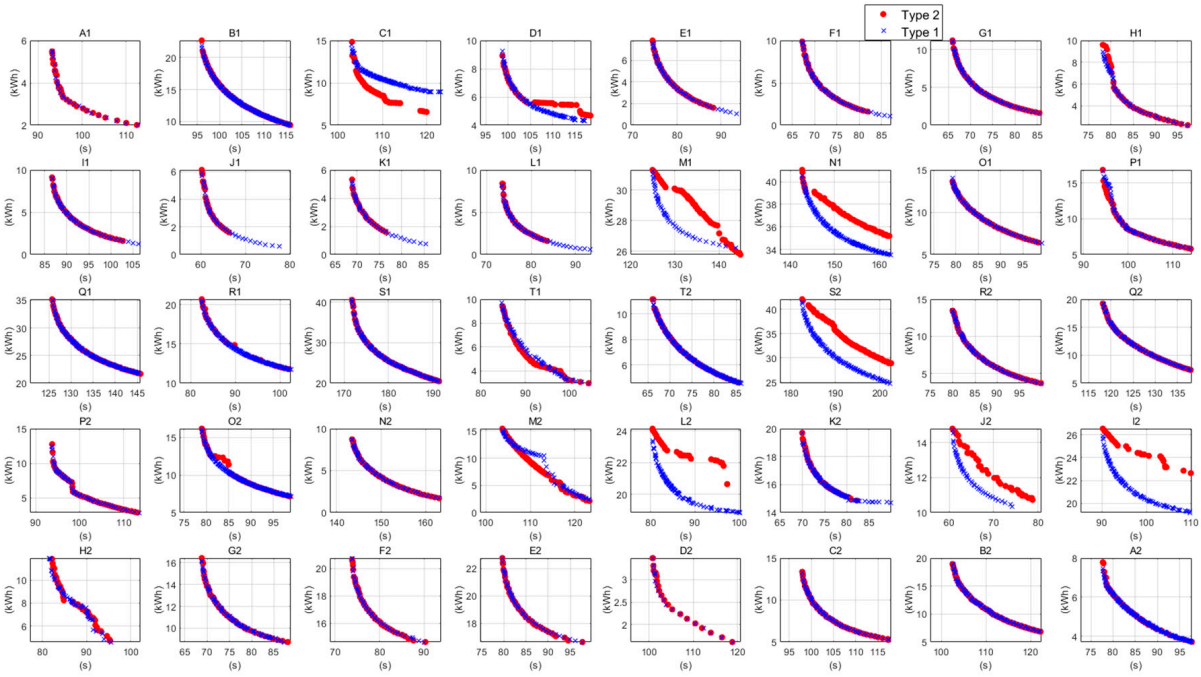

Figure 4. Comparison of Pareto fronts (energy vs. running time) for ATO Type 1 and Type 2 over the 40 interstations (A1–T1 and T2–A2) in the case study.

From this first inspection, it can be seen that Type 1 generally achieves equal or lower energy than Type 2 for the same running time, with the notable exceptions of interstations C1 and M2, where Type 2 offers competitive or superior points.

A particular case is L2, where the Type 2 Pareto front is not only less efficient but also differs in shape. Figure 5 presents an interstation-level analysis. Here, a low speed limit combined with a steep uphill gradient penalises Type 2. Since coasting must start at (and be tied to) a specific coasting command speed that is relatively low due to the speed limit, the train is subsequently forced to remotor. Given the steep ascent, only one remotoring is feasible, resulting in a short total coasting time. By contrast, Type 1 can hold speed until a later, more favourable location, then execute a single terminal coasting that merges smoothly with the final braking curve. For the same running time, this yields lower energy than Type 2 in this interstation.

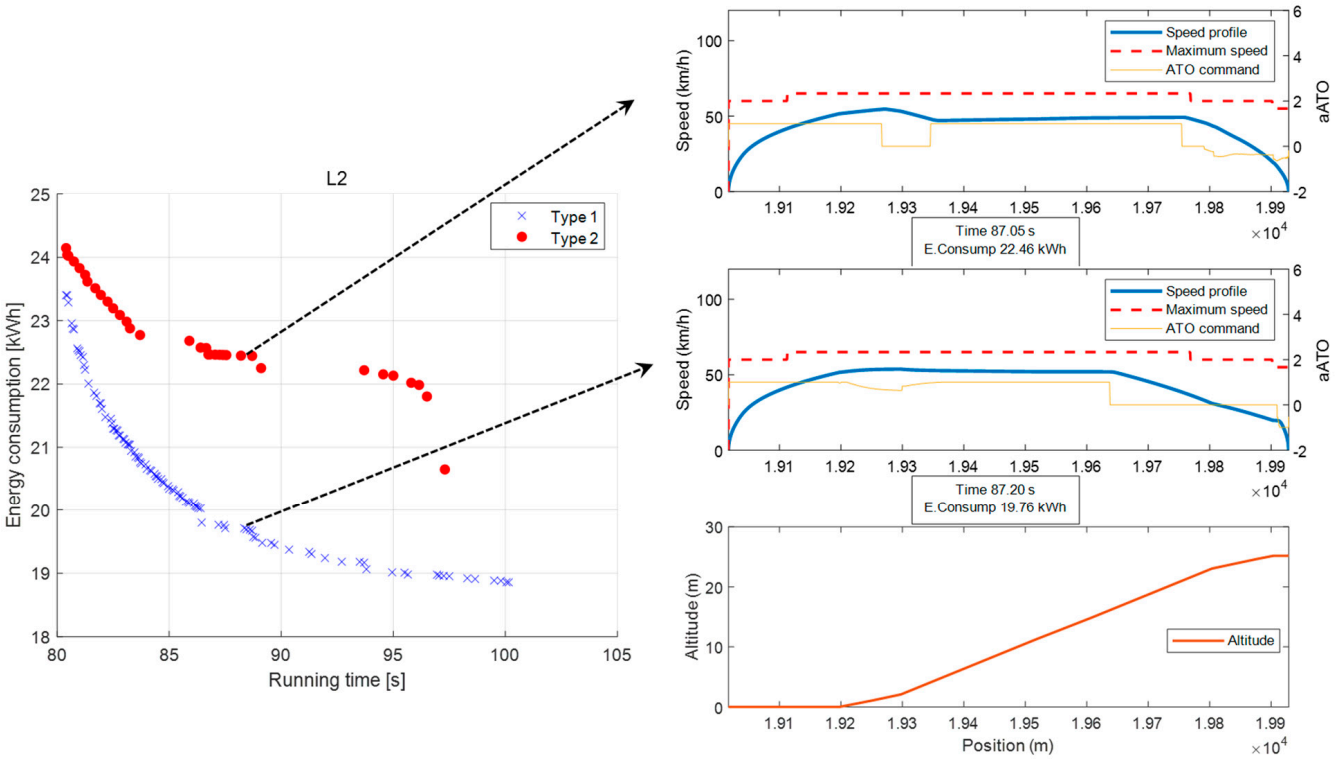

Figure 5. Pareto fronts and individual speed-profile simulations for a fixed running time in interstation L2.

A representative case where the Pareto fronts of both ATO systems are practically superimposed, suggesting identical performance despite their different control logics, has also been examined. It is shown in Figure 6 that isolates interstation B2. From each Pareto front, a speed profile with the same running time (107.5 s) has been selected, and the corresponding commands recovered. For ATO Type 1, those commands are: hold speed at 110 km/h and apply a single coast at kilometre 27.9. For ATO Type 2, the commands are coasting at 67.5 km/h and remotor when speed drops to 37 km/h.

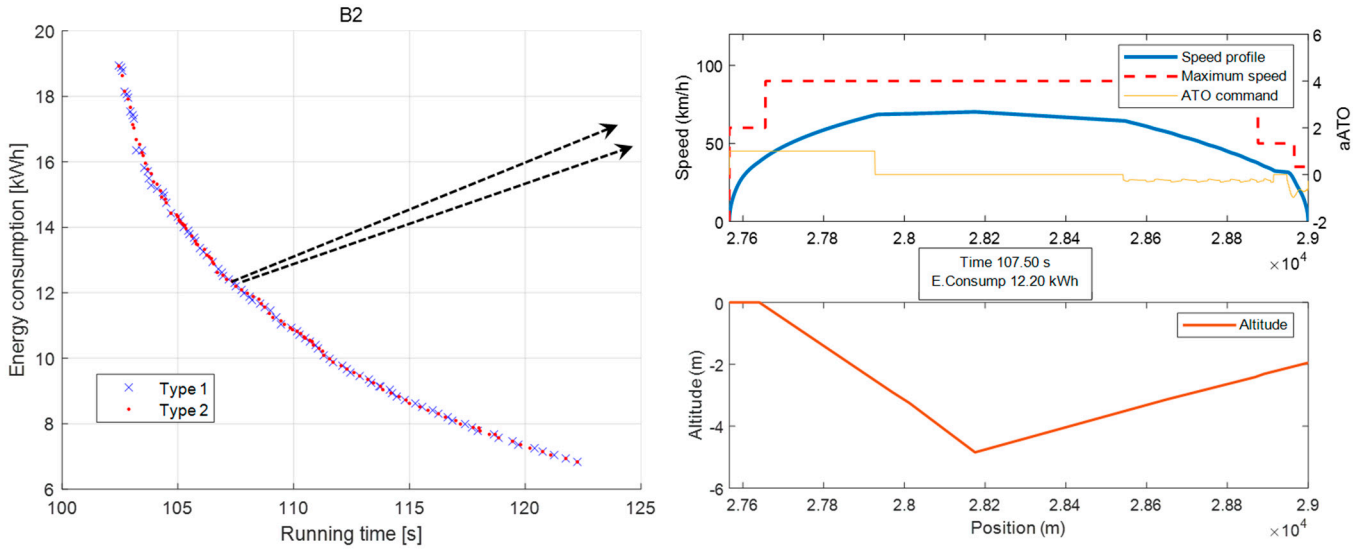

Figure 6. B2 analysis. Overlapped Pareto curves.

Given the abrupt transition from downhill to uphill in B2, both sets of commands degenerate to the same effective driving due to physical reasons. First, the Type 1 holding speed threshold is above the local speed limit, so the controller never actually enters a sustained holding phase; the profile behaves as if accelerating and coasting from the indicated kilometre point. Second, in Type 2, the remotoring threshold is never reached, so the planned remotor does not occur. In practice, both ATOs execute coasting at 67.5 km/h, which occurs at the 27.9 kilometre point, and then brake, yielding identical energy and time outcomes.

In other words, when topography, interstation length, and local speed limits jointly constrain the feasible command space, both ATO strategies can collapse to the same optimal speed profile, despite being parameterised differently.

4.2. Nominal Speed Profile Comparison

To obtain a quantitative measure of the time and energy differences between potential optimal speed profile designs, depending on the ATO type, the nominal speed profiles for both controllers under the 50% load baseline have been compared. This speed profile is defined as the one whose running time is 5 s slower than the flat-out profile. For each interstation, the nominal point is located on the corresponding Pareto front, and the resulting energy and time differences between Type 1 and Type 2 are plotted in Figure 7.

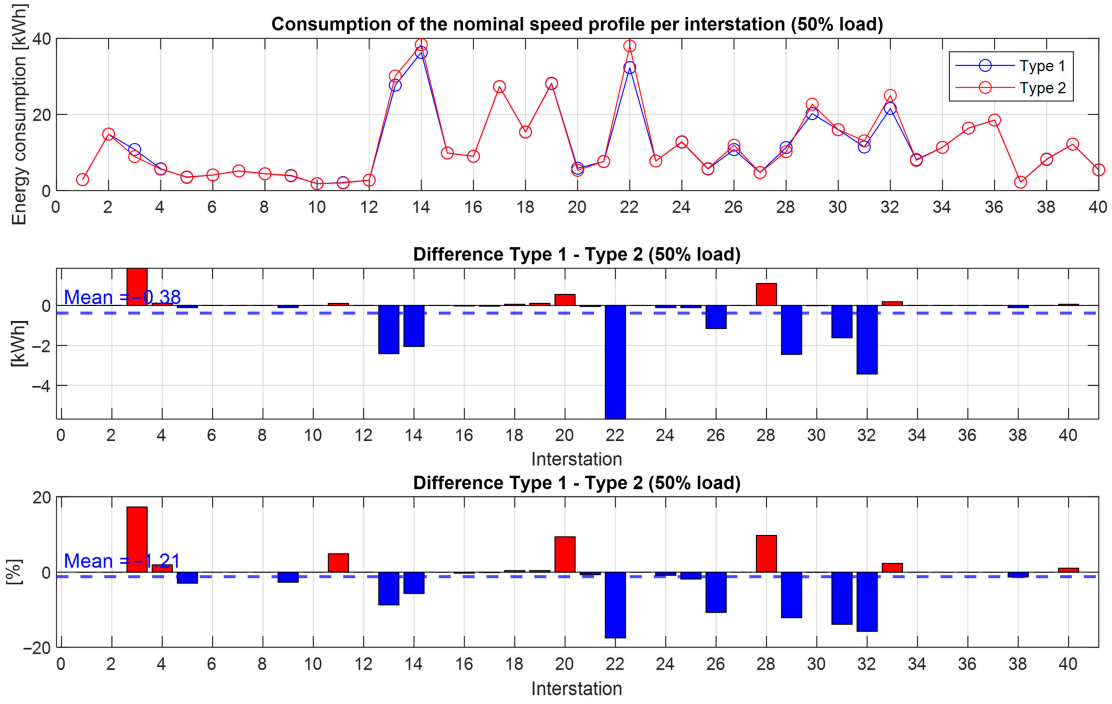

Figure 7. Comparison of nominal speed profiles. (5 s slower than flat-out; 50% load).

Overall, the differences are small for most interstations. On average, Type 1 exhibits slightly lower energy consumption than Type 2 at the nominal (+5 s) target by 1.21% on average, with the aggregate advantage driven by a small subset of interstations with larger gaps. In the remaining segments, the differences are marginal, with the notable exception of interstation 3 (C1), where Type 2 is clearly advantageous.

Interstation 3 (C1) has been analysed in more detail in Figure 8. Here, Type 1 reaches a holding speed and maintains it until a single terminal coasting segment can be initiated, from which the final braking curve is followed. By contrast, Type 2 triggers an early coasting as soon as the downhill gradient becomes beneficial, gaining speed without traction. Upon encountering a binding speed restriction, Type 2 applies braking to comply,

and, once the restriction is cleared, executes a remotoring phase to hold speed up to the final brake. Attempting to reproduce this behaviour with Type 1 (Figure 9) fails for two reasons: (i) without the possibility of a mid-course remotoring, an early coasting combined with the speed limit results in an excessive speed drop, violating comfort requirements (minimum speed), and (ii) even if not rejected on comfort grounds, the resulting profile lies at the inefficient tail of the Pareto front and would not be eligible as nominal. This example illustrates the value of allowing a single remotoring before the terminal brake when early coasting is advantageous to exploit line gradients.

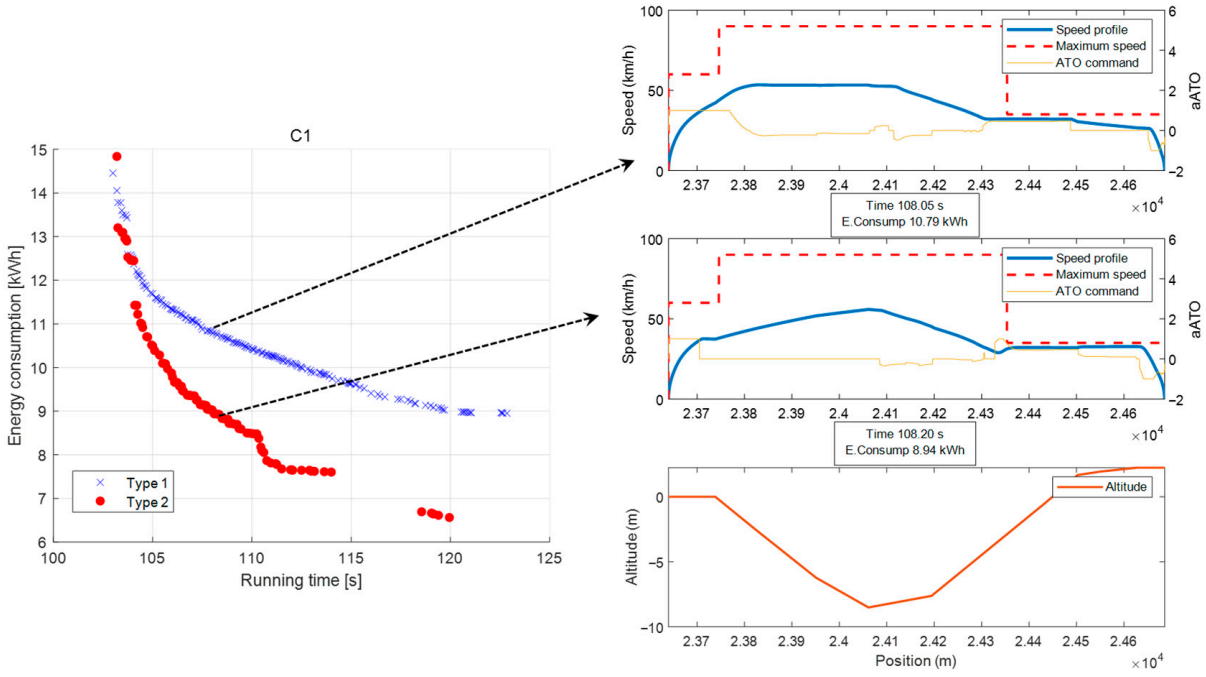

Figure 8. Nominal profiles in Interstation 3 (C1) (Type 1 vs. Type 2).

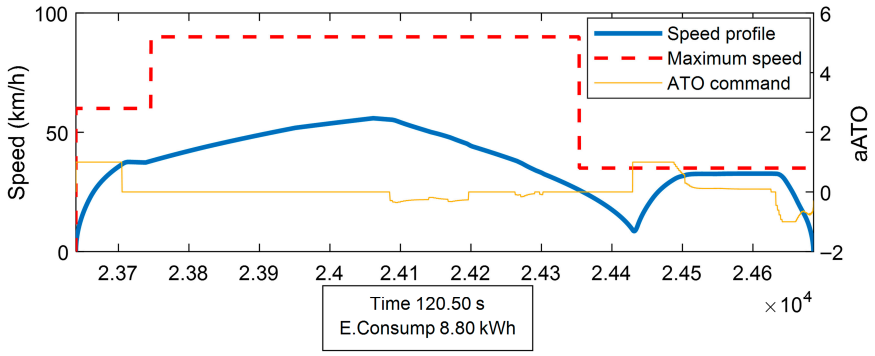

Figure 9. Attempt to reproduce the Type 2 nominal commands with Type 1 (comfort-driven rejection/sub-optimality) at interstation 3 (C1).

In contrast to the previous case, Interstation 22 (S2) (around 3 km long, high permitted speeds) clearly favours Type 1 (Figure 7). As shown in Figure 10, the combination of length (3 times longer than the previous case) and track profile makes it optimal to place a single, well-located coasting point so that natural deceleration aligns with the final braking curve. Type 1 can enforce this behaviour explicitly: it accelerates to the selected speed holding threshold, holds speed, and then triggers one terminal coast at the kilometre point that minimises energy while meeting the nominal time.

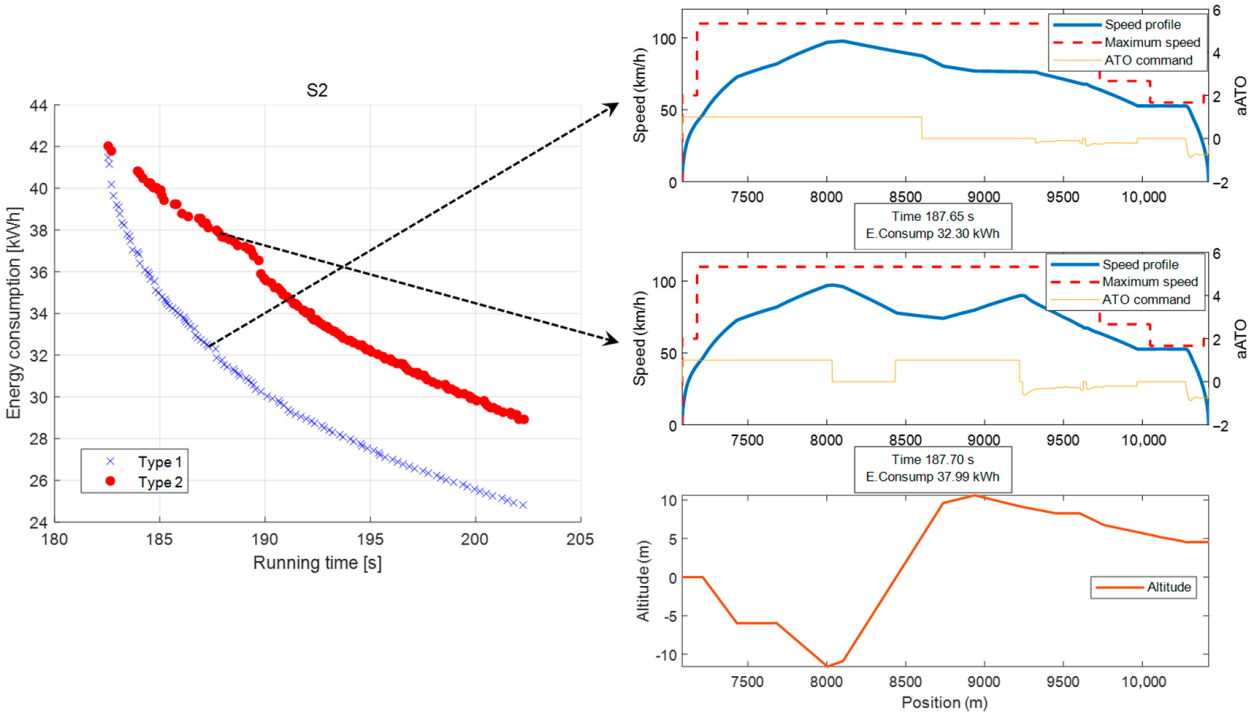

Figure 10. Nominal speed profiles in Interstation 22 (S2) (Type 1 vs. Type 2).

Replicating the same nominal solution with Type 2 is problematic. To force Type 2 to coast at the same optimal point as Type 1, the coasting threshold must be set to the corresponding coasting speed. However, since Type 2's logic is speed threshold-based (not position-based), the train reaches that speed earlier in the run and initiates an unintended early coast. Given the long distance and high limits, the controller then needs to remotor, later coast again, and finally remotor a second time to reach the stopping point (Figure 11). The net effect is longer running time and no compensating energy benefit: extra remotoring segments increase traction energy, while the early coast forfeits part of the gradient's natural braking effect near the stop.

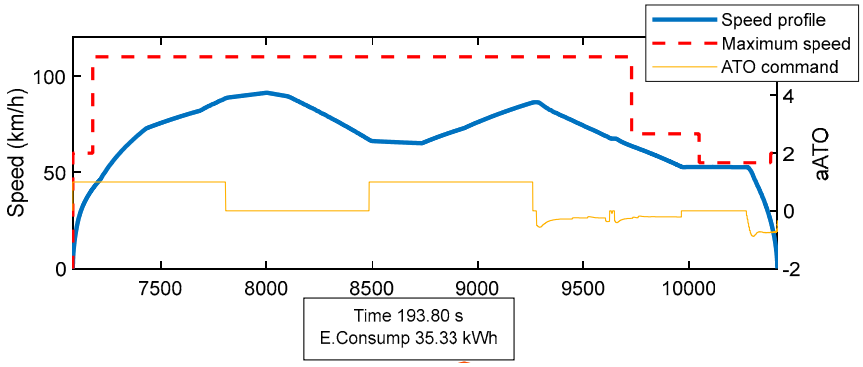

Figure 11. Attempt to reproduce the Type 1 nominal commands with Type 2 at interstation 22 (S2).

In summary, at the nominal (+5 s) target and 50% load, the average advantage slightly favours Type 1, but the ranking is interstation-dependent. Segments whose geometry rewards early coasting with a subsequent remotoring can flip the preference to Type 2, while sections that naturally permit a single, well-placed terminal coast tend to favour Type 1.

4.3. Sensitivity to Train Mass

As previously discussed, only a reduced set of four speed profiles per interstation is ultimately designed under nominal operating conditions. These profiles are tested and then programmed into the ATO equipment, where they remain fixed during daily operation. Consequently, it is not feasible to implement different speed-profile sets for different periods of the day. The selected profiles must therefore be sufficiently robust to real-world variability so that deviations from expected running times, used by traffic regulation systems and timetable planning tools, remain minimal. Therefore, an ATO speed profile robustness can be assessed calculating the variation in the running time and energy consumption with the extreme values of passenger mass.

Although the nominal profiles are designed to satisfy running-time, energy, and comfort criteria, actual operation may depart from these conditions. Among the varying factors, train mass (i.e., passenger load) is one of the most influential [28]. Mass directly affects tractive effort during acceleration, braking behaviour, and the rate of speed decay during coasting. As a result, even a well-designed speed profile may yield different running times or energy consumption when the train is lighter or heavier than assumed. Ensuring robustness to mass variability is therefore essential when parameterising ATO control strategies and selecting the profiles that will be deployed in service [27].

To assess robustness with respect to variations in the design mass, both ATO systems were evaluated under three representative load configurations: light (empty train, no load), nominal (base mass, 50% load, used to construct the Pareto fronts), and heavy (full passenger load). The analysis followed a two-step procedure.

First, for each interstation, and for every Pareto-optimal solution obtained at nominal mass, the same ATO commands set was simulated again (no re-optimised) under light and heavy mass configurations. This approach isolates the effect of mass on running time and energy, given a fixed control strategy.

Second, point-wise deviations with respect to the nominal case (Equations (12) and (13)) were computed for both running time and energy consumption, and the resulting distributions were summarised per-interstation. These quantitative results are analysed in Sections 4.3.2 and 4.3.3 after a preliminary qualitative assessment in Section 4.3.1.

$$\Delta t = t(mass) - t(nominal) \tag{12}$$

$$\Delta E = E(mass) - E(nominal) \tag{13}$$

4.3.1. Qualitative Screening

As in Section 4.1, a qualitative screen was first conducted. Figure 12 plots, for all 40 interstations, six Pareto curves per panel (three masses and two ATO types). This offers an immediate visual assessment of robustness. Ideally, for a given ATO, the three curves would appear as nearly parallel vertical shifts (energy axis) reflecting the expected increase/decrease in energy with mass, without appreciable horizontal displacement (running-time axis). The latter would indicate that timetables are not sensitive to passenger load throughout the day, and traffic control systems would be able to implement speed profiles with predictable travel times.

However, this ideal behaviour is not met in Figure 12. Interstations such as J1 exhibit the desired parallel energy shifts with negligible time drift (a robust case), whereas Q1 shows running-time spreads exceeding 20 s between load configurations for both ATO types, signalling potential susceptibility of schedule adherence to load.

Two additional patterns stand out. In E2, F2 and K2, very few points appear for Type 2 under non-nominal masses, and in M1, the light-mass curve for Type 1 is much more sparsely populated than its other two.

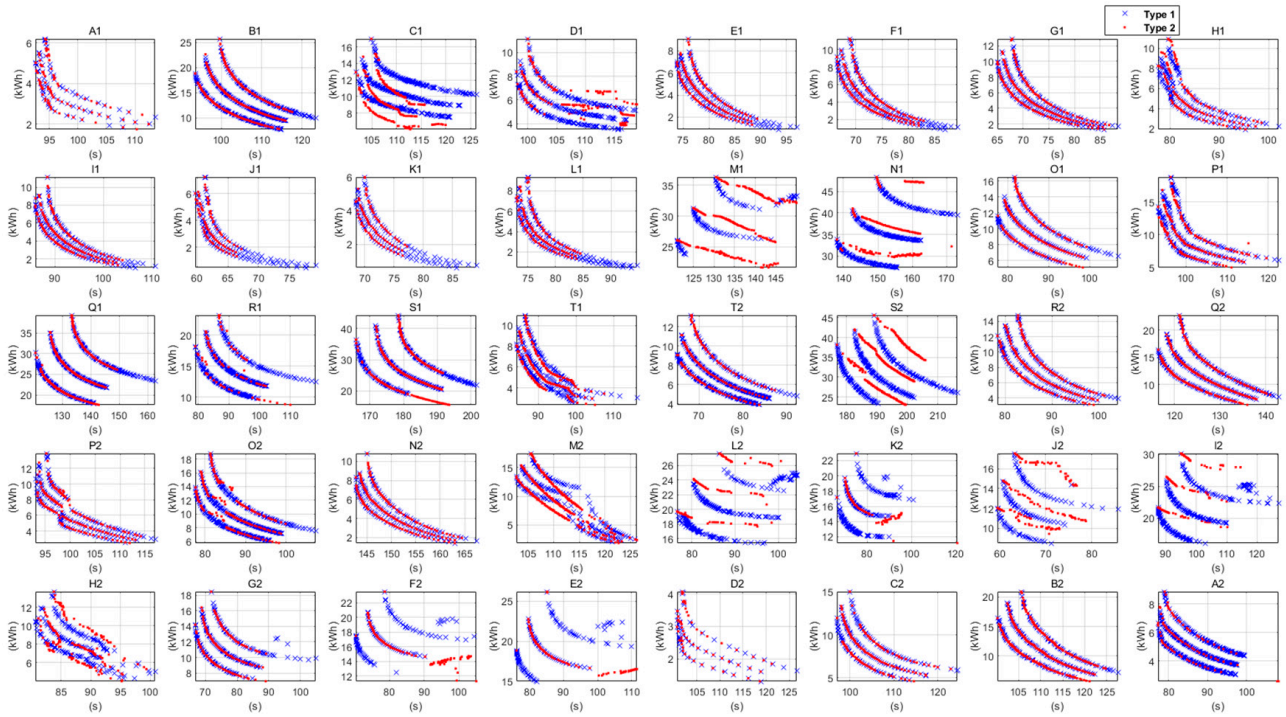

Figure 12. Comparison of Pareto fronts for ATO Type 1 and Type 2 across the 40 interstations (A1–T1 and T2–A2) under three load configurations (light, nominal, heavy).

To analyse those behaviours, Interstation E2 is isolated in Figure 13 to analyse those results. A Pareto-optimal point for Type 2 at the nominal mass has been selected from the left-hand Pareto plot, and its corresponding speed profiles under the three different load configurations are represented, on the right, along with the track gradient profile (fourth graph from the top). The selected solution is defined by the commands $b = 0.79$, $cs = 56.7$, $rm = 50$, interpreted as accelerate to 56.7 km/h, initiate coasting, and remotor when the speed falls to 50 km/h, and finally brake to stop at 0.79 m/s^2.

- Nominal mass (second graph from the top). On this steep uphill segment, the train struggles to reach the coasting threshold because the gradient requires sustained maximum traction for most of the run. As a result, the coasting point is only reached very late, and the final coasting phase merges directly with the braking curve. The remotoring threshold is never activated, so the remotoring command remains inactive. Operationally, the trajectory behaves as a single, late-coasting segment, followed immediately by braking.
- Light configuration (third graph from the top). Because the train is lighter, it reaches the coasting threshold earlier, and the remotor threshold becomes active, resulting in a profile with a pronounced speed fall and reapplication of traction.
- Heavy configuration (first graph from the top). The combination of high mass and steep uphill prevents the train from ever reaching the remotor threshold; most nominal-mass Pareto commands collapse to the flat-out-like behaviour under heavy mass. Consequently, many heavy-mass points cluster near the minimum-time extremity, reducing the visible spread of feasible Type 2 solutions.

For Type 1, this phenomenon is largely absent. The ability to hold speed and trigger a single terminal coasting at a suitable location allows the controller to bridge most of the ramp before coasting, only when it is feasible and beneficial. Hence, Type 1 retains more feasible speed profiles across different loads and shows smaller timing drift.

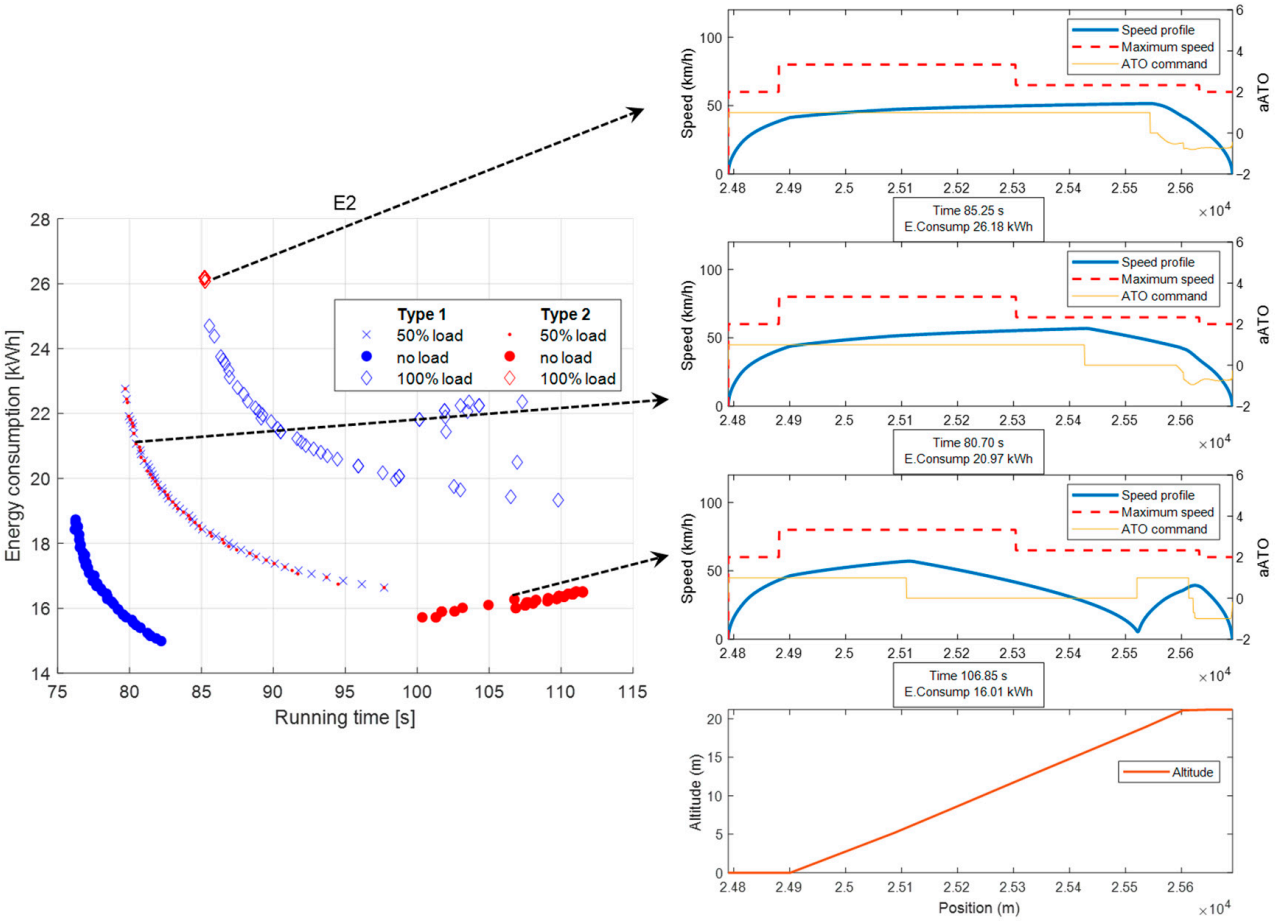

Figure 13. E2 analysis. Representation of a Type 2 speed profile computed with different train loads.

4.3.2. Effects on Running Time

Figure 14 summarises the distributions of Δt using heatmaps. Each panel condenses the 40 interstations into a grid of coloured cells, enabling a rapid visual identification of where running-time variations are largest. For each ATO type, two heatmaps are shown: (i) the mean Δt when re-simulating the nominal-mass Pareto profiles with an empty train, and (ii) the mean Δt when re-simulating them with full load. As a rule of thumb, we expect negative Δt for the no-load case (lighter trains should run faster) and positive Δt for the heavy case. However, two noteworthy exceptions from this rule appear:

- For Type 2 under no load, some cells show large positive Δt. These are the same mechanisms discussed earlier (see Section 4.3): the reduction in mass activates long coasting segments (and sometimes an additional remotoring) that yield very slow and even uncomfortable profiles; they inflate Δt despite the lighter train.
- For Type 2 under full load, interstation D1 exhibits a slightly negative mean Δt. This interstation has been isolated in Figure 15 to verify this result. D1 is a short interstation (around 0.8 km) with very low speed limits, leaving little room for modulation. The slowest Pareto-optimal speed profile has been selected. Its ATO commands are coasting at 32 km/h and re-motor after a 9.5 km/h drop. As seen in the figure, at nominal mass (50%), the train executes two coasting–remotoring cycles. However, at full load, speed decays faster during coasting, and there is just enough distance to trigger a final, short remotoring before the station stop. That late remotoring slightly increases the average speed slightly above the nominal-mass case, yielding a small negative Δt (i.e., marginally shorter time at full load).

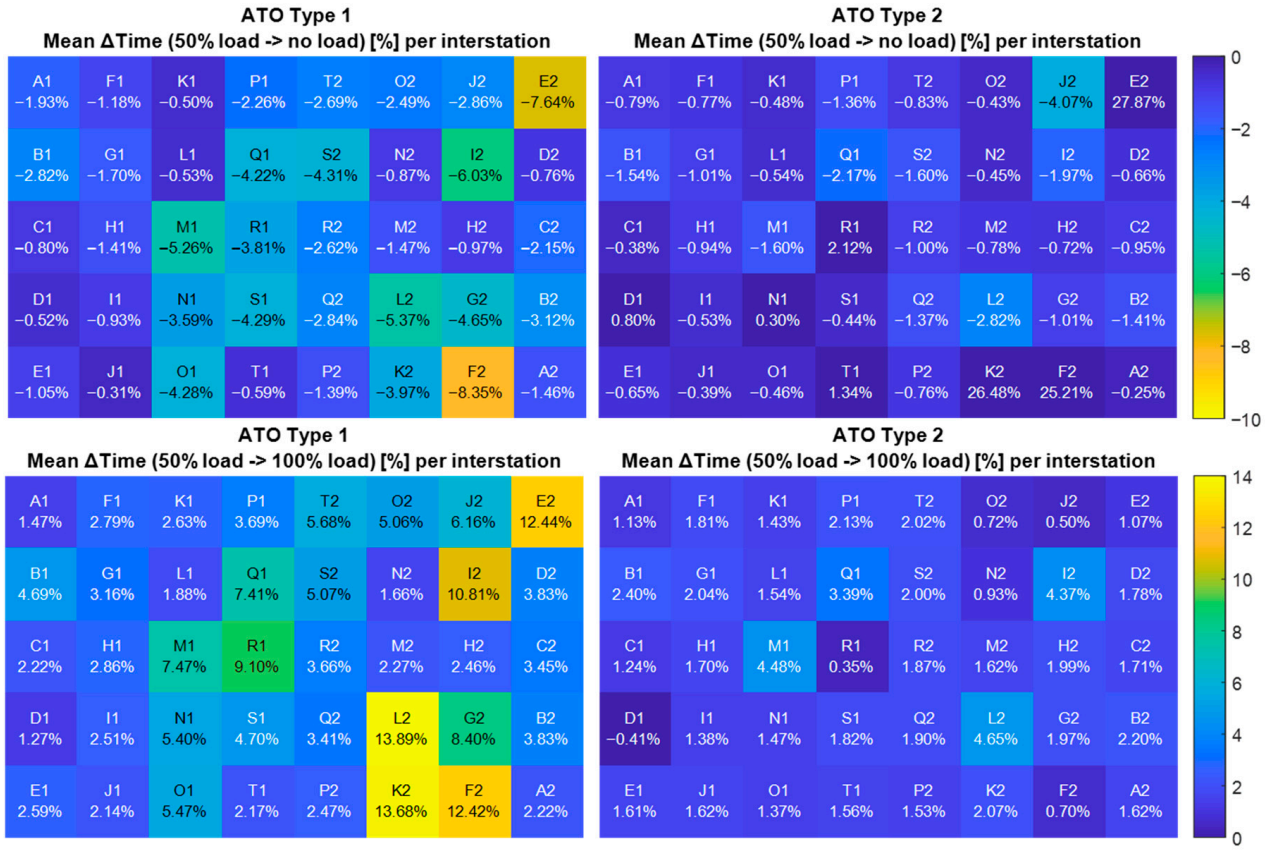

Figure 14. Heatmaps Mean ΔTime.

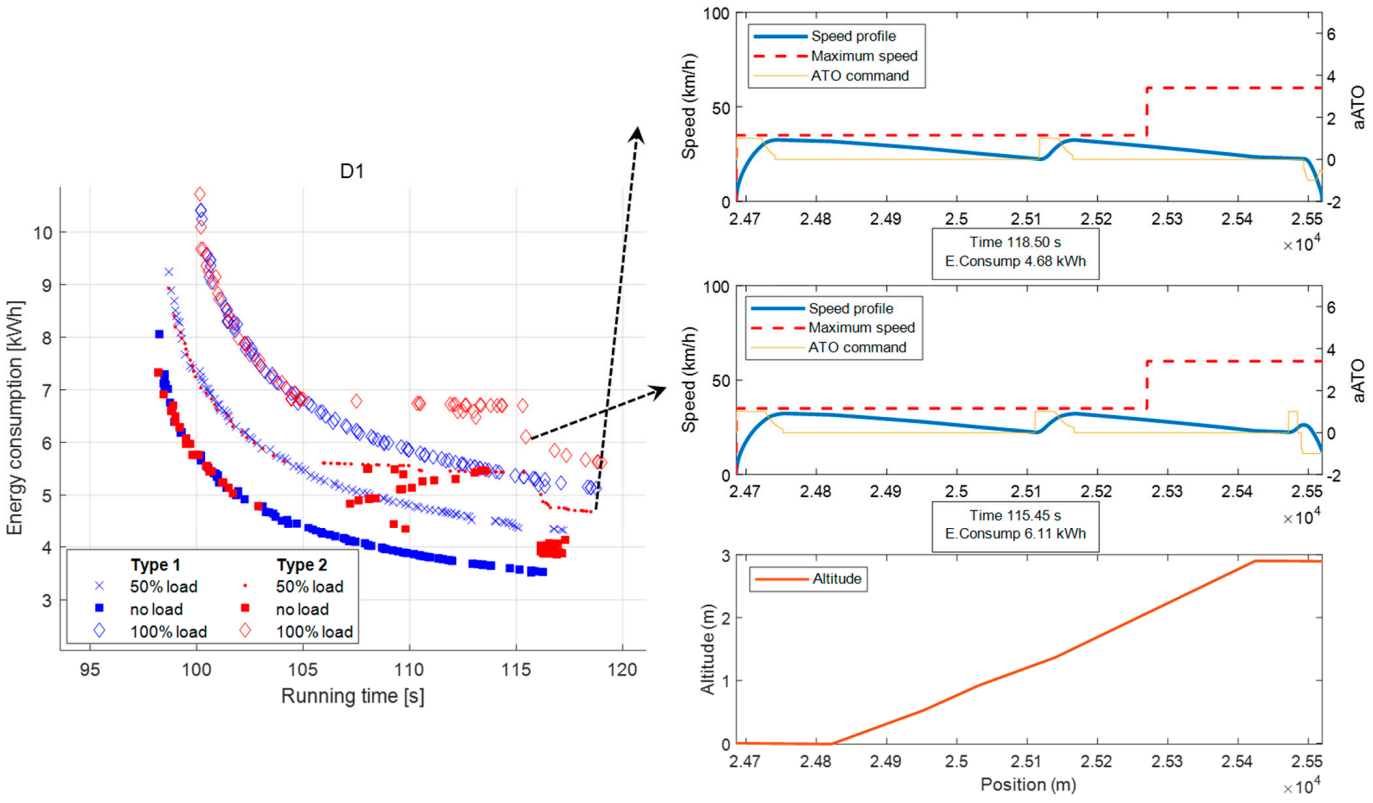

Figure 15. Analysis of interstation D1.

Once these peculiarities are understood, the global picture from Figure 14 is clear: Type 1 exhibits larger timing sensitivity to mass, with more interstations shaded in green/yellow and increases up to 10–14% when moving from nominal to full load in some cases. By contrast, Type 2 generally yields more stable running times once the Pareto-optimal profiles are fixed, with fewer and milder shifts across mass configurations.

In view of these results, from a timetable robustness perspective, when load variability is expected (e.g., peak/off-peak swings), Type 2 profiles tend to preserve running time more consistently. If Type 1 is deployed, slacks in the running-time targets may be considered to prevent mass-induced drifts.

4.3.3. Effects on Energy

Figure 16 reports heatmaps of the mean ΔE (Equation (13)) for light and heavy configurations when the same Pareto-optimal command sets designed at nominal mass are re-simulated with empty and full-load trains. As expected, energy deltas are negative for empty trains (since lighter mass reduces traction work) and positive for full loads (since heavier mass increases kinetic energy and traction demand). Unlike the running-time analysis, no anomalous behaviours are observed here.

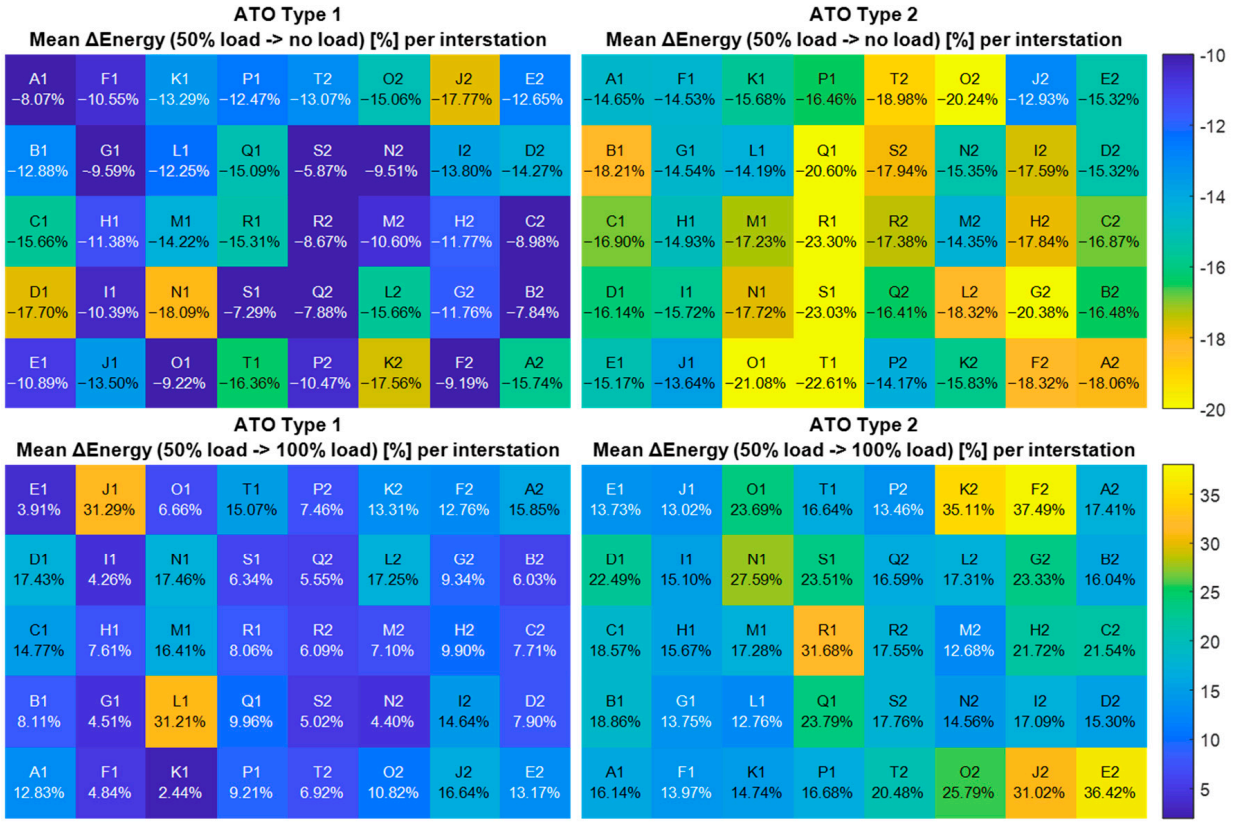

Figure 16. Heatmaps Mean ΔEnergy.

In this case, the patterns emerged from the heatmaps are that ATO Type 1 is markedly more robust in energy since most interstations appear in blue tones, indicating modest shifts with isolated interstations that show larger sensitivities. On the other hand, ATO Type 2 shows higher energy sensitivity to mass, indicated by more green/yellow cells with increases up to 36% under full load.

These differences are consistent with the underlying control logic. In Type 2, heavier trains experience faster speed decay during coasting, which triggers additional remotoring events (or lengthens them) that add traction energy. By contrast, Type 1 maintains a single,

well-placed coasting segment after a stable speed holding, so mass changes primarily scale the traction work during acceleration and the minor adjustments around the final braking merge, yielding smaller ΔE.

Consequently, if energy robustness to load variability is a priority (e.g., lines with large peak/off-peak swings or uncertain occupancy), Type 1 offers a more stable consumption result.

4.3.4. Line-Level Analysis

Nominal speed profiles are commonly selected to minimise the total end to end running time of the line while preserving operational margins [15,25]. To assess how passenger-load variability may affect the resulting timetable, the nominal profiles defined in Section 4.2 have been used to quantify line-level perturbations. For each interstation, the nominal-mass Pareto-selected profile has been re-simulated under light and heavy configurations, keeping the ATO commands fixed. This procedure isolates the impact of mass variations on end-to-end running time, enabling an evaluation of timetable robustness under realistic load fluctuations.

The interstation time deviation has been calculated following Equations (14) and (15), where i is the number of the interstation.

$$\Delta t_i^{light} = t_i(light) - t_i(nominal) \tag{14}$$

$$\Delta t_i^{heavy} = t_i(heavy) - t_i(nominal) \tag{15}$$

These deviations were then aggregated across all 40 interstations and expressed as a percentage of the total nominal running time as shown in Equations (16) and (17).

$$\%\Delta T^{light} = \frac{\sum_{i=1}^{40} \Delta t_i(light)}{\sum_{i=1}^{40} t_i(nominal)} \cdot 100 \tag{16}$$

$$\%\Delta T^{heavy} = \frac{\sum_{i=1}^{40} \Delta t_i(heavy)}{\sum_{i=1}^{40} t_i(nominal)} \cdot 100 \tag{17}$$

Table 1 reports the resulting line-level deviations. The positive $\%\Delta T^{light}$ for Type 2 may initially appear counter-intuitive; however, as explained in Section 4.3.2, lighter trains can experience prolonged coasting phases that activate additional remotoring cycles, producing localised increases in running time. When aggregated along the line, these effects can persist despite partial cancellation with negative deltas, yielding the reported 0.87%

A similar analysis has been conducted on energy consumption. The resulting energy deviations were computed for each interstation following Equations (18) and (19) and aggregated to line level according to Equations (20) and (21).

$$\Delta E_i^{light} = E_i(light) - E_i(nominal) \tag{18}$$

$$\Delta E_i^{heavy} = E_i(heavy) - E_i(nominal) \tag{19}$$

$$\%\Delta E^{light} = \frac{\sum_{i=1}^{40} \Delta E_i(light)}{\sum_{i=1}^{40} E_i(nominal)} \cdot 100 \tag{20}$$

$$\%\Delta E^{heavy} = \frac{\sum_{i=1}^{40} \Delta E_i(heavy)}{\sum_{i=1}^{40} E_i(nominal)} \cdot 100 \tag{21}$$

Results are summarised in Table 3. Negative values for $\%\Delta E^{light}$ reflect, as expected, the reduced energy consumption of empty trains, while positive deviations for $\%\Delta E^{heavy}$

reflect the higher traction demands of fully loaded trains. Type 1 again shows smaller deviations across both load changes, indicating notably greater energy robustness. Type 2, in contrast, exhibits much larger variation due to the sensitivity of its coasting–remotoring logic to mass-dependent speed decay. Heavier trains tend to trigger additional or longer remotoring phases, whereas lighter trains may coast for longer durations, altering energy use in both directions. When accumulated along the line, these effects lead to significantly wider fluctuations in total energy consumption.

Table 3. Impact of mass variations on line-level running time and energy consumption.

	$\%\Delta T^{light}$	$\%\Delta T^{heavy}$	$\%\Delta E^{light}$	$\%\Delta E^{heavy}$
ATO Type 1	−2.85%	5.15%	−12.1%	10.1%
ATO Type 2	0.87%	1.66%	−17.3%	23.1%

Overall, these figures suggest that, once nominal profiles are implemented, Type 2 delivers more timetable-stable behaviour under load variability, whereas Type 1 exhibits larger end-to-end drift (both faster when empty and slower when full). However, Type 1 offers better energy robustness. This highlights the importance of selecting the ATO strategy according to operational priorities (punctuality versus energy sustainability), especially on lines subject to strong variations in passenger load.

4.3.5. Final Aggregate Measures

To conclude the sensitivity study, Figure 17 presents boxplots of the interstation-level mean deltas in running time and energy consumption when transitioning from the nominal mass to the light and heavy configurations, separately for ATO Type 1 and ATO Type 2. Each box aggregates the distribution across the 40 interstations.

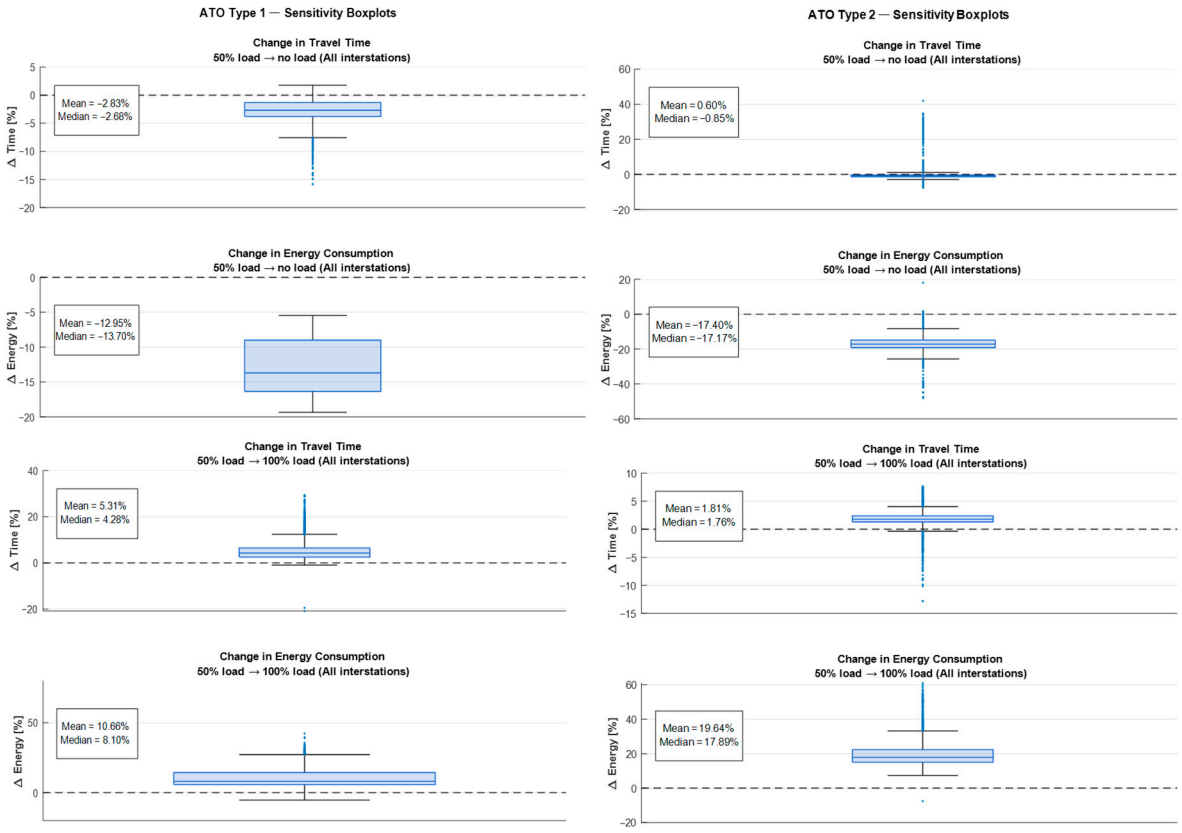

Figure 17. Boxplots of interstation mean ΔTime and ΔEnergy under light/heavy mass, for ATO Types 1 and 2.

Two robust patterns emerge:

- Energy robustness. The dispersion of ΔEnergy is consistently smaller for Type 1, with tighter interquartile range (low central variability) and shorter whiskers (low overall variability) under both mass perturbations. This confirms that Type 1 is more resilient to load variability in terms of energy, i.e., energy deviations remain closer to the nominal design when passenger load changes.
- Time robustness. Conversely, the dispersion of ΔTime is smaller for Type 2, indicating greater stability of running times under load changes. In operational terms, timetables derived from Type 2 nominal profiles are less sensitive to day-to-day fluctuations in occupancy.

These aggregate patterns are consistent with the interstation-level analyses: Type 1 is preferable when energy invariance under variable load is the priority, whereas Type 2 is advantageous when schedule adherence and runtime stability dominate (e.g., tight headways, punctuality KPIs).

5. Discussion

The results presented in the previous section highlight that the behaviour of each ATO strategy is shaped not only by its internal control logic but also by the geometric and operational characteristics of the line. The comparison of Pareto fronts across all interstations shows that neither control philosophy can be regarded as intrinsically superior. Instead, the relative performance of the two ATO types emerges from how effectively each can exploit (or mitigate) the local gradients, speed restrictions, and distances present in each segment.

The analysis of representative interstations illustrates this interaction. Type 1 benefits from its ability to sustain a stable speed holding phase and delay coasting to a precise position point, which is particularly advantageous in long segments with gentle gradients or high speed limits. In such cases, Type 2's threshold-based coasting triggers earlier and often leads to additional remotoring, which increases energy consumption without providing compensating time benefits. Conversely, in segments where an early terminal coast aligns naturally with the line's gradients and braking requirements, Type 2 can outperform Type 1 by exploiting favourable terrain through a controlled coasting–remotoring pattern. These contrasting behaviours indicate that local geometry does not merely influence energy levels but actively determines which ATO strategy can attain solutions close to the true energetic optimum.

A second dimension examined in this study is robustness to passenger load, a critical operational consideration often overlooked in ATO design. As shown before, speed profiles are traditionally computed for a fixed nominal mass (50% load in our case study), with the expectation that running times will remain sufficiently stable during real operation despite fluctuations in train occupancy. This stability is essential, as running times directly affect traffic regulation and timetable adherence.

The sensitivity analysis revealed a systematic trade-off between the two ATO philosophies: Type 2 maintains running times much more consistently when mass varies, while Type 1 maintains energy consumption more consistently under the same conditions. The underlying cause is structural. Type 2's coasting–remotoring logic naturally adapts to changes in speed-decay rates induced by load fluctuations, which stabilises runtime but simultaneously introduces larger variations in energy due to additional or prolonged remotoring episodes. Type 1, by contrast, exhibits predictable scaling of traction energy as mass increases, yet its inability to remotor mid-interstation makes it more susceptible to runtime drift: two sources of delay appear as load increases: first, the train reaches the coasting kilometre point at a lower speed, and second, its coasting deceleration is higher, leading to

longer total running times. These contrasting sensitivities imply that the selection between ATO strategies cannot be decoupled from the level of load variability expected on the line.

Another important insight is that certain interstations force both ATO systems into essentially identical driving behaviours, despite using different command parameters. This convergence occurs when steep gradients, abrupt curvature changes, or low speed limits severely restrict the feasible command space, leaving only a small region of the Pareto set physically manageable. These cases illustrate that the theoretical difference between the two control philosophies may be neutralised in practice by infrastructural constraints.

Taken together, these findings highlight the importance of understanding the interaction between ATO logic and infrastructure when designing speed profiles for metro operations. The optimal strategy varies spatially along the line, and is influenced by gradients, interstation length, speed limits, and expected load variability. Beyond comparing the two systems, the results also suggest potential directions for controller tuning. For example, moderating Type 2's remotoring thresholds could improve its energy robustness, while fine-grained adjustment of Type 1's cruising phases could reduce its runtime sensitivity.

Another possibility to be explored could be the development of a hybrid onboard controller capable of adopting Type 1-like behaviour when energy efficiency is the dominant objective, and switching to Type 2-like strategies when timetable adherence becomes critical. Such an adaptive approach could leverage the strengths of both philosophies and reduce their respective weaknesses in demanding operational contexts.

Ultimately, the discussion underscores that the value of an ATO system depends not only on its nominal efficiency but on its ability to deliver predictable, robust behaviour under the full range of operating conditions. This provides the foundation for the conclusions that follow, which synthesise the practical implications of these findings for metro ATO deployment and optimisation.

6. Conclusions

This work presents the comparative evaluation of two ATO control strategies: Type 1, based on constant-speed holding followed by a single coasting phase defined by a coasting starting position, and Type 2, which incorporates either speed holding or coasting–remotoring cycles defined by speed thresholds, using a high-fidelity simulation framework and Pareto-front analysis under comfort and operational constraints. The results show that, for equal running times, Type 1 generally delivers equal or lower energy consumption than Type 2; however, this trend is not universal. Local geometry, particularly track segments that reward early coasting followed by a limited remotoring, can invert the advantage in favour of Type 2. In some constrained interstations, both strategies ultimately converge to the same effective driving profile due to physical limits imposed by slopes, speed restrictions, and interstation length.

At the nominal operating point, differences between the two controllers remain modest on average, with aggregate advantages driven by a reduced number of interstations where geometric effects are especially pronounced. The sensitivity analysis with varying train mass reveals a clear trade-off between the two ATO strategies. Type 2 exhibits smaller variations in running time and therefore offers greater timetable stability when passenger load fluctuates, while Type 1 is more robust in terms of energy consumption, showing smaller deviations from the nominal energy values under the same mass variability.

These findings highlight that neither ATO system dominates in all contexts. Instead, their comparative performance is shaped by the interplay between control logic and interstation characteristics. Practical improvements are possible: Type 2's sensitivity in energy could be mitigated by constraining remotoring events or incorporating mass-aware thresholds, whereas Type 1's sensitivity in runtime could be alleviated through fine adjust-

ments of speed-holding parameters or the introduction of strategic slack in geometrically sensitive interstations.

In summary, the study shows that a single ATO strategy cannot optimally serve an entire metro line under all operating conditions. A hybrid deployment policy, selecting the controller best suited to each interstation class and incorporating robustness-oriented parameterisation, offers a more effective approach to aligning energy-efficiency objectives with operational reliability in real-world environments.

Author Contributions: Conceptualization, A.F.-R., A.P.C. and A.F.-C.; Methodology, A.F.-R. and A.P.C.; Software, M.D.; Validation, A.F.-R., A.P.C. and A.F.-C.; Formal analysis, M.D. and A.F.-R.; Investigation, M.D. and A.F.-R.; Writing—original draft, M.D.; Writing—review & editing, A.F.-R., A.P.C. and A.F.-C.; Visualization, M.D.; Supervision, A.P.C. and A.F.-C.; Project administration, A.P.C. and A.F.-C. Funding acquisition, A.F.-C. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data presented in this study are available on request from the corresponding author.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Corlu, C.G.; de la Torre, R.; Serrano-Hernandez, A.; Juan, A.A.; Faulin, J. Optimizing Energy Consumption in Transportation: Literature Review, Insights, and Research Opportunities. *Energies* **2020**, *13*, 1115. [CrossRef]
2. Douglas, H.; Roberts, C.; Hillmansen, S.; Schmid, F. An assessment of available measures to reduce traction energy use in railway networks. *Energy Convers. Manag.* **2015**, *106*, 1149–1165. [CrossRef]
3. Emery, D. Towards Automatic Train Operation for long distance services: State-of-the art and challenges. In Proceedings of the 17th Swiss Transport Research Conference, Ascona, Switzerland, 17–19 May 2017. Available online: https://www.semanticscholar.org/paper/Towards-Automatic-Train-Operation-for-long-distance-Emery/bfc0d13d1d45d18d75692a42fdf3c60abe7c0790 (accessed on 18 November 2025).
4. Dai, X.; Zhao, H.; Yu, S.; Cui, D.; Zhang, Q.; Dong, H.; Chai, T. Dynamic Scheduling, Operation Control and Their Integration in High-Speed Railways: A Review of Recent Research. *IEEE Trans. Intell. Transp. Syst.* **2022**, *23*, 13994–14010. [CrossRef]
5. Feng, M.; Wu, C.; Lu, S.; Wang, Y. Notch-based speed trajectory optimisation for high-speed railway automatic train operation. *Proc. Inst. Mech. Eng. Part F J. Rail Rapid Transit* **2022**, *236*, 159–171. [CrossRef]
6. Suzuki, W. Safety Analysis and Design Improvement for Semi-Automatic Train Operation (STO) in High-Speed Rail Using STPA. Master's Thesis, Massachusetts Institute of Technology, Cambridge, MA, USA, 2025. Available online: https://dspace.mit.edu/handle/1721.1/162531 (accessed on 19 November 2025).
7. Yasui, Y. Automatic train operation system for the high speed Shinkansen train. *Comput. Railw. X* **2006**, *1*, 441–446. [CrossRef]
8. Kahl, N.; Weigelt, N.; Mazzone, A. Rail freight automation in Shift2Rail—Development of prototypes. In Proceedings of the 7th Transport Research Arena TRA 2018, Vienna, Austria, 16–19 April 2018; Zenodo: Geneva, Switzerland, 2018; pp. 1–9. [CrossRef]
9. Yin, J.; Tang, T.; Yang, L.; Xun, J.; Huang, Y.; Gao, Z. Research and development of automatic train operation for railway transportation systems: A survey. *Transp. Res. Part C Emerg. Technol.* **2017**, *85*, 548–572. [CrossRef]
10. Cidoncha, Á.; Fernández-Rodríguez, A.; Cucala, A.P.; Fernández-Cardador, A. Predictive Traffic Regulation Model for Railway Mass Transit Lines Equipped with Continuous Communication Systems. *IEEE Access* **2024**, *12*, 96862–96877. [CrossRef]
11. Ichikawa, K. Application of optimization theory for bounded state variable problems to the operation of train. *Bull. Jpn. Soc. Mech. Eng.* **1968**, *11*, 857. [CrossRef]
12. Howlett, P. The Optimal Control of a Train. *Ann. Oper. Res.* **2000**, *98*, 65. [CrossRef]
13. Scheepmaker, G.M.; Goverde, R.M.P.; Kroon, L.G. Review of energy-efficient train control and timetabling. *Eur. J. Oper. Res.* **2017**, *257*, 355–376. [CrossRef]
14. Caramia, P.; Lauro, G.; Pagano, M.; Natale, P. Automatic train operation systems: A survey on algorithm and performance index. In Proceedings of the 2017 AEIT International Annual Conference, Cagliari, Italy, 20–22 September 2017; pp. 1–6. [CrossRef]

15. Sánchez-Contreras, G.; Fernández-Rodríguez, A.; Fernández-Cardador, A.; Cucala, A.P. A Two-Level Fuzzy Multi-Objective Design of ATO Driving Commands for Energy-Efficient Operation of Metropolitan Railway Lines. *Sustainability* **2023**, *15*, 9238. [CrossRef]
16. *IEEE Std 1474.1-2004*; IEEE Standard for Communications-Based Train Control (CBTC) Performance and Functional Requirements. IEEE: New York, NY, USA, 2005. [CrossRef]
17. Chang, C.S.; Sim, S.S. Optimising train movements through coast control using genetic algorithms. *IEE Proc. Electr. Power Appl.* **1997**, *144*, 65. [CrossRef]
18. Pu, Q.; Zhu, X.; Liu, J.; Cai, D.; Fu, G.; Wei, D.; Sun, J.; Zhang, R. Integrated Optimal Design of Speed Profile and Fuzzy PID Controller for Train with Multifactor Consideration. *IEEE Access* **2020**, *8*, 152146–152160. [CrossRef]
19. Seong Ho, H.; Yun Sub, B.; Jong Hyen, B.; Tae Ki, A.; Su Gil, L.; Hyun Jun, P. An optimal automatic train operation (ATO) control using genetic algorithms (GA). In Proceedings of the IEEE. IEEE Region 10 Conference. TENCON 99. 'Multimedia Technology for Asia-Pacific Information Infrastructure' (Cat. No.99CH37030), Cheju, Republic of Korea, 15–17 September 1999; IEEE: New York, NY, USA, 1999; Volume 1. [CrossRef]
20. Liu, S.; Cao, F.; Xun, J.; Wang, Y. Energy-Efficient Operation of Single Train Based on the Control Strategy of ATO. In Proceedings of the 2015 IEEE 18th International Conference on Intelligent Transportation Systems, Washington, DC, USA, 15–18 September 2015; IEEE: Gran Canaria, Spain, 2015; pp. 2580–2586. [CrossRef]
21. Chang, C.S.; Xu, D.Y.; Quek, H.B. Pareto-optimal set based multiobjective tuning of fuzzy automatic train operation for mass transit system. *IEE Proc. Electr. Power Appl.* **1999**, *146*, 577–583. [CrossRef]
22. Fernández, P.M.; Font Torres, J.B.; Sanchís, I.V.; Franco, R.I. Multi objective ant Colony Optimisation to obtain efficient metro speed profiles. *Proc. Inst. Mech. Eng. Part F J. Rail Rapid Transit* **2023**, *237*, 232–242. [CrossRef]
23. Ke, B.-R.; Lin, C.-L.; Yang, C.-C. Optimisation of train energy-efficient operation for mass rapid transit systems. *IET Intell. Transp. Syst.* **2012**, *6*, 58–66. [CrossRef]
24. Domínguez, M.; Fernández-Cardador, A.; Cucala, A.P.; Gonsalves, T.; Fernández, A. Multi objective particle swarm optimization algorithm for the design of efficient ATO speed profiles in metro lines. *Eng. Appl. Artif. Intell.* **2014**, *29*, 43–53. [CrossRef]
25. Blanco-Castillo, M.; Fernández-Rodríguez, A.; Su, S.; Fernández-Cardador, A.; Cucala, A.P. Efficient automatic operation of a metro line: Eco-driving design with optimized use of regenerative energy and rolling stock consideration. *Eng. Optim.* **2025**, *57*, 2044–2076. [CrossRef]
26. Chen, X.; Li, K.; Zhang, L.; Tian, Z. Robust Optimization of Energy-Saving Train Trajectories Under Passenger Load Uncertainty Based on p-NSGA-II. *IEEE Trans. Transp. Electrif.* **2023**, *9*, 1826–1844. [CrossRef]
27. Fernández-Rodríguez, A.; Fernández-Cardador, A.; Cucala, A.P.; Domínguez, M.; Gonsalves, T. Design of Robust and Energy-Efficient ATO Speed Profiles of Metropolitan Lines Considering Train Load Variations and Delays. *IEEE Trans. Intell. Transp. Syst.* **2015**, *16*, 2061–2071. [CrossRef]
28. Mo, P.; Yang, L.; Gao, Z. Energy-Efficient Train Operation Strategy with Speed Profiles Selection for an Urban Metro Line. *Transp. Res. Rec.* **2019**, *2673*, 348–360. [CrossRef]
29. Yang, X.; Chen, A.; Ning, B.; Tang, T. A stochastic model for the integrated optimization on metro timetable and speed profile with uncertain train mass. *Transp. Res. Part B Methodol.* **2016**, *91*, 424–445. [CrossRef]
30. Liu, X.; Tian, Z.; Jiang, L.; Lu, S.; Zeng, P. An improved energy-efficient driving strategy for routes with various gradients and speed limits. *IET Intell. Transp. Syst.* **2024**, *18*, 949–963. [CrossRef]
31. Davis, W.J. The Tractive Resistance of Electric Locomotives and Cars. *Gen. Electr. Rev.* **1926**, *XXIX*, 685–707.
32. Coello, C.A.C.; Pulido, G.T.; Lechuga, M.S. Handling multiple objectives with particle swarm optimization. *IEEE Trans. Evol. Comput.* **2004**, *8*, 256–279. [CrossRef]
33. Deb, K.; Pratap, A.; Agarwal, S.; Meyarivan, T. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Trans. Evol. Comput.* **2002**, *6*, 182–197. [CrossRef]