Staged cost-capacity optimisation of trackside train detection layouts under ETCS Hybrid Train Detection

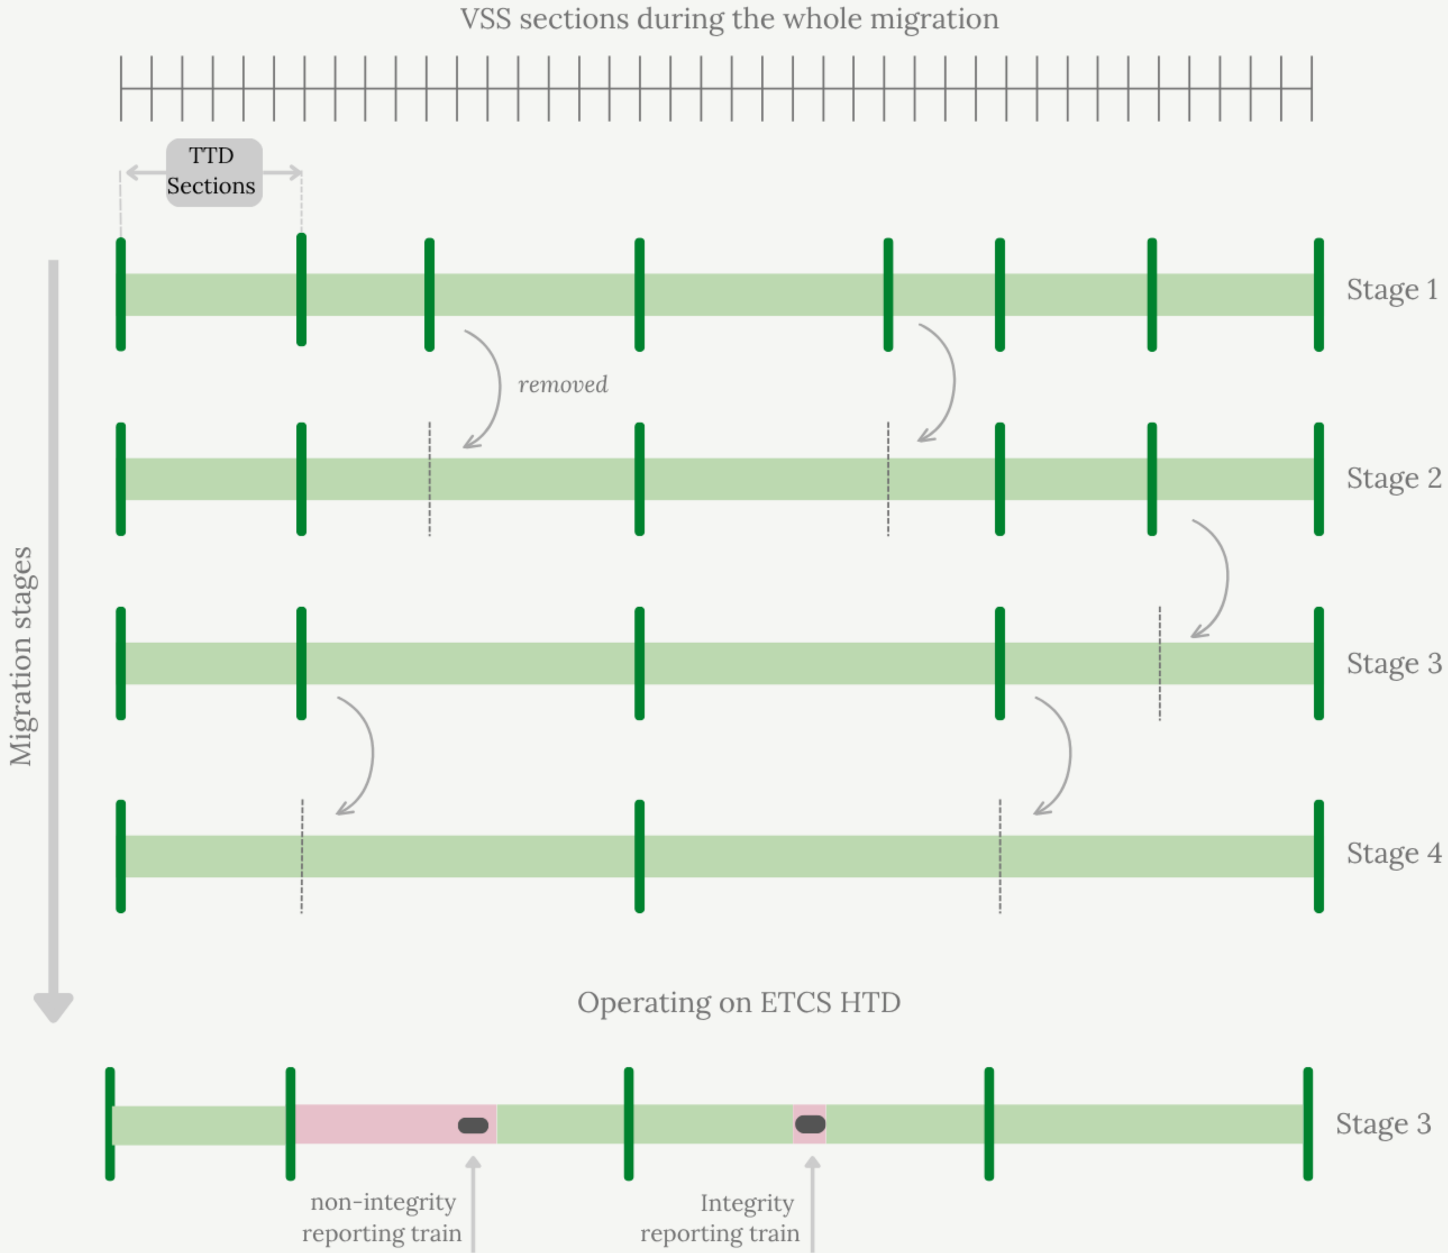

TU Delft

ProRail

Staged cost-capacity optimisation of trackside train detection layouts under ETCS Hybrid Train Detection

---

Delft University of Technology

MSc Transport, Infrastructure and Logistics

18/08/2026

Thesis Candidate:	J.I.J. (Julia) Lakeman
Student number:	5171156
Chair TU Delft:	E. (Egidio) Quaglietta
Supervisors TU Delft:	M.Y. (Yousef) Maknoon, J. (Joelle) Aoun
Supervisors ProRail:	D.J. (David) Koopman, T.B. (Thomas) van den Berg

TU Delft

Preface

This thesis investigates the optimisation of physical trackside train detection layouts during a migration towards ETCS Hybrid Train Detection. It focuses on how remaining physical detection boundaries can be selected over successive migration stages, while balancing railway capacity effects with lifecycle and deployment costs. The case study applies the framework to a single-direction section of the SAAL corridor, in collaboration with ProRail.

I carried out this research as part of the Master's programme in Transport, Infrastructure and Logistics at Delft University of Technology, in collaboration with ProRail. The thesis combines railway signalling, traffic management, infrastructure planning and optimisation. As a result, the project became both a technical and personal learning process, requiring me to connect operational railway knowledge with mathematical modelling and simulation-based evaluation.

My interest in this topic developed during my studies at TU Delft. The course Decision Making for Logistics and Transport Systems, taught by Yousef Maknoon, introduced me to the role of optimisation and decision support in complex transport systems. Subsequently, the course Railway Traffic Management, taught by Egidio Quaglietta, showed me how railway capacity, signalling and train interactions can be analysed in a structured way. Together, these courses inspired me to work on a topic at the intersection of optimisation and railway networks. At the start of this thesis, I was still relatively unfamiliar with ERTMS, and ETCS Level 2 Hybrid Train Detection felt like a black box. Understanding the concept, its terminology and its operational implications took considerable effort, but this also made the research process challenging and rewarding.

This thesis addresses researchers and practitioners working on railway capacity, signalling design, ETCS deployment and infrastructure migration strategies. More specifically, it may be relevant for infrastructure managers and railway stakeholders interested in introducing Hybrid Train Detection in a cost-effective way while maintaining sufficient operational performance.

I would like to thank my graduation committee for their guidance throughout this research. I am grateful to my chair, E. (Egidio) Quaglietta, and my TU Delft supervisors, M.Y. (Yousef) Maknoon and J. (Joelle) Aoun, for their academic guidance, critical feedback and encouragement. I would also like to thank my ProRail supervisors, D.J. (David) Koopman and T.B. (Thomas) van den Berg, for their practical insights, technical feedback and support in connecting the research to railway practice.

I especially appreciate how my supervisors challenged me to think more critically, improve the structure of my work and push the research further than I initially thought possible. Their feedback often encouraged me outside my comfort zone, but it also helped me get the most out of myself and this thesis. I am thankful for the time, patience and expertise they contributed throughout the process.

Finally, I would like to thank TU Delft and ProRail for the opportunity to work on this research topic. The thesis process was not always easy, but it taught me a great deal about railway systems, modelling, academic writing and perseverance.

J.I.J. (Julia) Lakeman
Delft, September 2026

Executive summary

European railway networks face growing transport demand while existing rail infrastructure is increasingly congested and costly to expand. Across Europe, European Rail Traffic Management System (ERTMS) is being progressively deployed to modernise and harmonise railway signalling, with European Train Control System (ETCS) providing its train control function. Conventional ETCS Level 2 (L2) fixed block operation uses continuous cab signalling but still relies on physical Trackside Train Detection (TTD) to determine whether track sections are clear. As a potential next step, ETCS Hybrid Train Detection (HTD) is being developed and evaluated as a concept that subdivides physical TTD sections into fixed Virtual Sub-Sections (VSSs). When train position is reported and integrity is confirmed, infrastructure behind the train can be released at the finer VSS level. Without confirmed integrity, release remains dependent on the underlying physical TTD section.

This creates a migration problem. During the gradual deployment of Train Integrity Monitoring (TIM), which provides the required train integrity confirmation, physical train detection remains important for part of the traffic. Retaining the complete inherited detection layout, however, may become unnecessarily expensive as more movements can use finer VSS level release. The design question is therefore not simply how much physical detection can be removed, but which TTD boundaries should remain, where they should be located, and during which migration stages they continue to justify their cost. This thesis therefore addresses the following research question:

> *How can TTD boundary layouts under ETCS HTD be designed to balance capacity, infrastructure cost, and migration feasibility under mixed-traffic conditions?*

A review of existing signalling layout methods and ETCS HTD research showed that methods for designing conventional fixed block layouts are well established, while previous HTD studies have mainly examined predefined physical detection layouts and different VSS configurations. However, no systematic method was identified that determines which physical TTD boundaries should remain during the different stages of TIM deployment.

This thesis therefore treats the remaining physical train detection as a staged location selection problem and formulates it as a bi-objective Mixed Integer Programming (MIP), solved through an objective-equivalent Mixed Integer Linear Programming (MILP) reformulation. Before the optimisation model is formulated, the physical and virtual layout structure under ETCS HTD and the engineering requirements that constrain its design are established. Physical TTD sections and VSSs determine how finely infrastructure can be released, whereas ETCS L2 fixed blocks and ETCS HTD virtual blocks determine the locations up to which a train may be authorised to proceed. The block layout design requirements are translated into mandatory, forbidden, and optional TTD boundary positions so that required physical detection functions are preserved around constrained infrastructure. A fixed VSS grid of approximately 50 m is then constructed, defining the design domain within which the optimisation selects the physical TTD boundaries for each migration stage.

Within this design domain, the model searches for staged TTD layouts that balance capacity performance against the costs directly associated with the HTD migration. Capacity is represented by the expected line headway between consecutive trains, while the cost objective includes the placement, removal, and maintenance of physical TTD boundaries and the deployment of TIM. The optimisation generates a range of cost-capacity trade-off alternatives that together form a Pareto frontier, from which a multi-criteria assessment selects one for detailed evaluation. After computational verification, the selected layouts are evaluated in RailSys, a microscopic railway simulation model. The selected T1 and T2 layouts are compared with the analytical model for analytical-microscopic headway consistency, while nominal UIC 406 capacity consumption is evaluated for all three case-study types. Operational stability is assessed separately for the primary T1 case through deterministic disturbance recovery.

The model is applied to a 14.3 km section of the Schiphol-Amsterdam-Almere-Lelystad (SAAL) corridor

between Weesp and Almere Centrum. T1 represents the primary brownfield case, starting from 44 inherited physical TTD boundaries and using heterogeneous stopping and non-stopping services. T2 uses the same traffic but assumes that the physical detection layout can be designed from scratch, while T3 combines this greenfield starting condition with homogeneous intercity traffic. T1-T2 therefore examines the value of inherited infrastructure, while T2-T3 examines the influence of traffic composition.

The results first show that the location of physical detection matters independently of the number of boundaries. At the reference condition, the existing 44-boundary layout produces an expected line headway of 155.15 s. Keeping the same number of boundaries but optimising their locations reduces this to 151.98 s. An optimised layout with only 26 selected boundaries achieves comparable analytical performance within the predefined equivalence tolerance, with an expected line headway of 156.42 s. Boundary count alone is therefore an incomplete indicator of capacity performance: a boundary has nominal capacity value only when it is located where it can reduce the line headway.

Whether ETCS HTD is introduced by adapting an existing brownfield detection layout or by designing a new greenfield layout mainly affects the cost of achieving comparable capacity performance. The selected T1 and T2 alternatives achieve almost identical maximum expected line headways of 162.18 s and 162.58 s, while their modelled migration costs are approximately € 2.03 M and € 2.82 M, respectively. The sensitivity analysis also shows that physical detection and TIM deployment are coupled migration decisions, since changing the assumed cost of TIM changes the selected physical infrastructure trajectory.

Traffic composition changes the spatial value of physical detection as well. With heterogeneous stopping and non-stopping services, consecutive trains progress differently along the corridor, causing useful residual boundaries to concentrate around specific capacity-critical locations, such as around intermediate stations where stopping patterns create larger differences between train movements. In the homogeneous T3 case, the identical intercity services have no intermediate stops within the case study section, so these differences do not arise and the selected boundaries are distributed more evenly.

For the primary T1 case, the selected migration trajectory contains 21, 20, 16, and 12 active physical TTD boundaries over the four migration stages, eventually converging to the 12 mandatory-only positions. Retaining all 44 inherited boundaries throughout the migration would reduce the line headway from 162.18 s to 154.51 s, but would increase the modelled migration cost from approximately € 2.03 M to € 4.61 M. The selected trajectory therefore reduces modelled discounted lifecycle and deployment cost by approximately € 2.58 M, or 56%, while increasing the maximum expected line headway by 7.67 s, or approximately 5%.

Microscopic RailSys evaluation shows that, for T1 and T2, capacity consumption is up to approximately 2.5 percentage points above the heterogeneous ETCS L2 fixed block baseline during the early migration stages, while T3 shows a larger difference relative to the homogeneous baseline before all three cases converge towards their traffic-matched baseline by the final stage.

The disturbance analysis shows, however, that nominal capacity and operational stability are distinct design criteria. Because the optimisation balances lifecycle cost against nominal capacity, the early T1 layout recovers more slowly from the evaluated deterministic 15 min standstill than the heterogeneous fixed block baseline. An ex-post RailSys diagnostic shows that retaining one or two additional physical boundaries within the long Almere Poort (AMPO)-Almere Muziekwijk (ALMM) TTD section restores the evaluated recovery approximately to the baseline level without changing nominal UIC 406 capacity consumption, showing that a boundary can have little nominal capacity value while still supporting disturbance recovery.

The thesis concludes that TTD boundary layouts under ETCS HTD should be designed as a staged location selection problem rather than reduced according to a fixed boundary count. Capacity and lifecycle cost can be balanced by retaining physical detection where it remains valuable under the prevailing traffic mix and TIM deployment level, while engineering requirements and inherited infrastructure constrain which boundaries can be removed and when. Because a cost-capacity optimal layout does not necessarily provide sufficient operational stability, the selected layout should also be screened for disturbance recovery, while further research should quantify how TIM deployment and the residual TTD layout jointly affect this performance.

Contents

Preface **i**

Executive summary **ii**

1 Introduction **1**
1.1 Background Context . . . 1
1.2 Scientific Contribution . . . 2
1.3 Report Structure . . . 2

2 Background **3**
2.1 Conventional lineside railway signalling and trackside train detection . . . 3
2.1.1 Block principle . . . 3
2.1.2 Track circuits and axle counters . . . 4
2.1.3 Interlocking principle . . . 4
2.1.4 Automatic train protection . . . 4
2.2 Blocking time theory . . . 4
2.2.1 Blocking time of a block section . . . 4
2.2.2 Blocking time stairway and minimum line headway . . . 5
2.3 European Rail Traffic Management System and European Train Control System (ERTMS/ETCS) . . . 6
2.3.1 Blocking time in ETCS L2 fixed block operation . . . 8
2.4 ETCS HTD . . . 10
2.4.1 Section and block hierarchy . . . 10
2.4.2 Blocking time under ETCS HTD . . . 12
2.4.3 Implications for physical TTD layout design . . . 12

3 Literature Review **13**
3.1 Block layout design methods . . . 13
3.2 ETCS Hybrid train detection . . . 14
3.2.1 Capacity and operational evidence . . . 14
3.2.2 Infrastructure placement constraints . . . 14
3.2.3 Design methods for ETCS HTD . . . 15
3.3 Synthesis and research gap . . . 15
3.4 Key performance indicators in the literature . . . 18

4 Layout design requirements under ETCS HTD **20**
4.1 TTD layout requirements under ETCS L2 HTD . . . 21
4.1.1 Safety critical aspects . . . 22
4.1.2 Operationally critical aspects . . . 24
4.1.3 General recommendations . . . 25
4.1.4 System behaviour constraints . . . 26
4.2 Block layout requirements under ETCS L2 HTD . . . 27
4.3 Implications for TTD layout optimisation . . . 28

5 Optimisation framework for TTD layout design under ETCS HTD **29**
5.1 Analytical headway formulation under ETCS L2 . . . 30
5.1.1 Blocking time and minimum line headway . . . 30
5.1.2 Minimum headway components . . . 31
5.2 Analytical headway formulation under ETCS HTD . . . 32
5.2.1 Release regimes and section level approximation . . . 32
5.3 Lifecycle and deployment cost model . . . 34
5.4 Design procedure for VSS grid construction and TTD boundary placement . . . 35

5.5 Optimisation model for TTD boundary placement . . . 36
5.5.1 Modelling scope and assumptions . . . 37
5.5.2 Precomputed parameters . . . 38
5.5.3 Notation . . . 38
5.5.4 MIP formulation . . . 40
5.5.5 Pareto frontier generation and layout selection . . . 43
5.5.6 Computational verification and spatial layout check . . . 44
5.6 Evaluation approach . . . 45
5.6.1 Case study comparisons . . . 45
5.6.2 Analytical-microscopic consistency . . . 45
5.6.3 KPI evaluation . . . 46

6 Case study: SAAL corridor 48
6.1 Case study corridor and current operation . . . 48
6.2 Case study design . . . 49
6.3 Shared case study parameters . . . 49
6.4 TIM deployment and migration horizon . . . 50
6.5 Lifecycle and deployment cost assumptions . . . 51
6.6 Consistency-check configurations . . . 52
6.7 RailSys case study setup . . . 52
6.7.1 Timetable and deployment trajectory . . . 53
6.7.2 Disturbance scenario for operational stability (KPI 3) . . . 53

7 Results 54
7.1 Analytical-microscopic consistency and spatial response . . . 54
7.2 Infrastructure inheritance and staged asset reuse . . . 55
7.3 Traffic-dependent spatial value of TTD boundaries . . . 56
7.4 Capacity, cost, and operational stability . . . 57
7.4.1 Capacity consumption . . . 58
7.4.2 Lifecycle cost relative to full retention . . . 59
7.4.3 Operational stability under deterministic disturbance . . . 60

8 Discussion 63
8.1 Interpretation of the main findings . . . 63
8.2 Nominal capacity and operational stability . . . 63
8.3 Methodological limitations and layout selection . . . 64
8.4 Implications and future research . . . 64

9 Conclusion 65

References 66

A Additional optimisation data 69
A.1 Computational setup and solution settings . . . 69
A.2 Pareto frontiers by case study type . . . 69
A.3 Initial physical TTD reference layouts . . . 70
A.4 Selected physical TTD layouts by case study type . . . 70
A.4.1 T1 (brownfield, heterogeneous traffic) . . . 71
A.4.2 T2 (greenfield, heterogeneous traffic) . . . 72
A.4.3 T3 (greenfield, homogeneous IC-only traffic) . . . 73
A.5 T1 TIM retrofit cost sensitivity . . . 74

B MILP reformulation 76
B.1 MILP formulation . . . 76
B.2 Big-M calibration . . . 78
B.3 Certified lower bound . . . 78

C Optimisation verification and preprocessing diagnostics 79
C.1 Computation of the headway components . . . 79
C.1.1 Approach time . . . 79

C.1.2 Clearing time 80
C.1.3 Constant components 80
C.2 Sensitivity of the analytical headway components 80
C.3 Station-stop motion correction 82
C.4 Big-M and LP diagnostics 83
C.4.1 LP strengthening cut effectiveness 83
C.4.2 Scenario-specific big-M constants 84
C.4.3 Big-M implementation check 84
C.4.4 LP relaxation diagnostics 84
C.4.5 Handcase validation output 85
C.5 Validation suite 85

D RailSys simulation data 87
D.1 Analytical versus microscopic line headway comparison 87
D.2 Deterministic disturbance time-distance diagrams 87

E AI Statement 90

List of Figures

2.1 Three-aspect signalling principle (The Railway Technical Website, 2023). 3
2.2 Blocking time components in fixed block signalling, based on Pachl (2021). 5
2.3 Blocking time stairway of two consecutive trains: the signal headway of the critical block section determines the minimum line headway, based on Pachl (2021). 6
2.4 Blocking time components in ETCS L2 fixed block operation, based on Pachl (2021). . . 9
2.5 Section and block hierarchy under ETCS HTD, showing the relationship between the ETCS L2 fixed block (layer 1), physical TTD sections (layer 2), VSS (layer 3), and virtual blocks (layer 4). 10
2.6 Release mechanism in ETCS L2 fixed block operation and ETCS HTD. 12

4.1 Section and block hierarchy under ETCS HTD (reproduced from Figure 2.5). 20
4.2 Overview of the TTD boundary placement requirements. 24
4.3 Overview of the block boundary placement restrictions. 28

5.1 Overview of the design framework for optimal TTD boundary placement. 29
5.2 Blocking time of a block section under ETCS L2 fixed block operation, based on Pachl (2021). 30
5.3 Blocking time stairways of a leading and a following train, showing the critical block section that determines the minimum line headway, based on Pachl (2021). 31
5.4 Design procedure for TTD and VSS layout construction. 35
5.5 Illustrative example of the two coupled decisions the optimisation model determines jointly: the yearly TIM deployment trajectory (left) and the staged physical TTD boundary layout on the fixed VSS grid (right). 37
5.6 Schematic illustration of the two-step decision procedure in which the Pareto frontier represents the trade-off between headway and cost (a), after which the D37.1 Multi-Criteria Analysis (MCA) weights are used to select one alternative (b). 44
5.7 UIC 406 capacity consumption concept, distinguishing occupation time, additional time, and unused capacity within the defined evaluation period (International Union of Railways, 2013). 46

6.1 Speed profiles and mandatory TTD positions for the case study corridor section. 48

7.1 Comparison of the current, optimised 44-boundary, and least-cost TTD layouts on the SAAL corridor at $\rho = 0$. 54
7.2 Pareto frontiers for the brownfield T1 and greenfield T2 migration approaches. 55
7.3 Selected staged TTD layouts for the brownfield T1 and greenfield T2 migration approaches. 56
7.4 MCA-selected staged TTD layouts across the four migration stages for T2 (greenfield, heterogeneous) and T3 (greenfield, homogeneous). 57
7.5 RailSys UIC 406 compressed timetable of the T1 stage 2 layout, with red arrows marking the binding critical block sections that set the line headway, using the same compression procedure applied to every case study type and migration stage. 59
7.6 Disturbance recovery indicators after the T1 15 min standstill at AMPO, across the four migration stages under best-case and worst-case TIM assignments. 61
7.7 RailSys time-distance diagrams for the T1 stage 2 disturbance under the best-case and worst-case TIM assignments, with red occupation bands indicating trains delayed by more than 60 s. 62

A.1 Pareto frontiers for the three case study types, with the MCA-selected alternative marked in red in each panel. 70

A.2 Selected per-stage physical TTD layout for T1: brownfield initial layout with heterogeneous Intercity (IC)/Stop Train (ST) traffic. 72
A.3 Selected per-stage physical TTD layout for T2: greenfield initial layout with heterogeneous IC/ST traffic. 73
A.4 Selected per-stage physical TTD layout for T3: greenfield initial layout with homogeneous IC-only traffic. 74
A.5 Selected per-stage physical TTD layouts for T1 under the €0 and €42 500 TIM retrofit cost settings; the €84 420 baseline layout is shown in Figure A.2. 75

C.1 Approach time computation by braking curve intersection. 80
C.2 Tornado diagram showing the sensitivity of the slow track-releasing pair headway $h_{\mathrm{TTD},pq}$ to ±20 % variation of each headway component. 82

D.1 Case study time-distance diagrams for the deterministic T1 disturbance at AMPO, stages 1 and 2, under the best-case and worst-case TIM assignments. 88
D.2 Case study time-distance diagrams for the deterministic T1 disturbance at AMPO, stages 3 and 4, under the best-case and worst-case TIM assignments, with the two assignments coinciding in stage 4 where all eight trainsets are TIM-equipped. 89

List of Tables

3.1 Overview of key studies on block section layout design and their relation to the present research. . . . 16

4.1 Overview of the safety-critical TTD placement requirements. . . . 22
4.2 Overview of the operationally critical TTD placement requirements. . . . 22
4.3 Overview of the general recommendations for TTD placement. . . . 22

5.1 Parameters precomputed for each representative leader-follower scenario before optimisation. . . . 38
5.2 Notation for the TTD boundary optimisation model. . . . 39

6.1 Case study types considered in this study. . . . 49
6.2 Main input parameters for the SAAL case study. . . . 50
6.3 Discrete TIM deployment level sets per year for the case study. . . . 51
6.4 RailSys configurations used for the microscopic evaluation. . . . 52
6.5 Best-case and worst-case TIM assignments for the T1 KPI 3 delay-recovery simulations, with effective deployment ratios constrained to multiples of 1/8 and the T1 ratio column reporting the selected end of stage deployment ratio represented by each assignment. . . . 53

7.1 End of stage TIM ratios of the T1, T2, and T3 MCA-selected trajectories, used in the RailSys capacity evaluation. . . . 58
7.2 RailSys UIC 406 capacity consumption for the selected layouts and their traffic-matched ETCS L2 fixed block baselines B1 and B2 across the four migration stages. . . . 58

A.1 Dimensions of the solved MILP for the T1 case study instance. . . . 69
A.2 Brownfield initial TTD layout used for T1 (44 boundaries). . . . 70
A.3 Mandatory TTD boundary positions enforced from stage 1 onward, applied to T1, T2, and T3. . . . 70
A.4 T1 MCA-selected layout, brownfield initial layout with heterogeneous traffic (discounted cost €2,031,448, $\bar{H} = 162.18\,\mathrm{s}$). . . . 71
A.5 T2 MCA-selected layout, greenfield initial layout with heterogeneous traffic (discounted cost €2,820,816, $\bar{H} = 162.58\,\mathrm{s}$). . . . 72
A.6 T3 MCA-selected layout, greenfield initial layout with homogeneous IC-only traffic (discounted cost €2,660,346, $\bar{H} = 94.41\,\mathrm{s}$). . . . 73
A.7 T1 MCA-selected objective results under three TIM retrofit cost settings. . . . 74

B.1 Notation additional to Chapter 5 used in the MILP reformulation. . . . 76

C.1 Effect of the two built-in LP-strengthening components on the root-node LP bound. . . . 83
C.2 Scenario-specific fast-release headway floor F_k and conservative slow-release bound $M_{G,k}$ used in the reported arc-flow optimisation runs. . . . 84
C.3 Big-M implementation check on the Type 1 instance at four cost levels with a 3 % MIP gap, comparing the calibrated tight formulation against a deliberately loose slow headway bound ($M_G = 5000\,\mathrm{s}$), with the numerical tolerance $\epsilon_M = 3 \cdot \mathrm{gap} \cdot \max(\bar{H}_{\mathrm{tight}}, \bar{H}_{\mathrm{loose}})$ below which observed differences are attributable to the reported MIP gap rather than to a bound implementation error. . . . 84
C.4 LP relaxation diagnostics for selected cost levels (Type 1, $c_{\mathrm{train}} = €84{,}420$, both built-in LP-strengthening components active). . . . 85
C.5 Handcase validation results for Type 1 restricted to a 5-position free subset (IC-IC single scenario), where MIP and brute-force headways agree exactly in all cases. . . . 85

C.6 Validation suite results on the Type 1 case study instance. 86

D.1 Reference condition comparison of the analytical and microscopic RailSys expected line headway for the current brownfield 44-boundary layout at $\rho = 0$, with all relevant movements dependent on TTD level release. 87

D.2 Stage-specific comparison of the analytical and microscopic RailSys expected line headway for the selected T1 and T2 layouts under partial TIM deployment, with the difference given as RailSys minus analytical headway. 87

Glossary

ALMM Almere Muziekwijk. iii, 48, 50, 55–58, 62, 63, 65, 71–73

AMPO Almere Poort. iii, vii, viii, 48, 50, 53, 55–58, 61–63, 65, 71–73, 87–89

ATO Automatic Train Operation. 17

ATP Automatic Train Protection. 3, 4

CCS TSI Control-Command and Signalling Technical Specification for Interoperability. 6, 8

EEIG European Economic Interest Group. 7

EoA End of Authority. 8, 9, 11, 12, 27

ERA European Union Agency for Railways. 7

ERTMS European Rail Traffic Management System. ii, 1, 3, 6

ETCS European Train Control System. ii–iv, vii, ix, 1–3, 6–18, 20, 21, 27, 29, 30, 32–34, 48, 49, 52–55, 58–60, 63, 65, 79, 80, 87

EUG ERTMS Users Group. 7, 10, 11, 27

HTD Hybrid Train Detection. ii–iv, vii, 1–3, 7, 10–21, 25–27, 29, 32–34, 52–54, 58, 59, 63, 65, 79, 87

IC Intercity. viii, ix, 48–50, 53, 69, 70, 72–74

L1 Level 1. 7, 20

L2 Level 2. ii–iv, vii, ix, 1, 6–13, 16, 18, 20–28, 30–32, 48, 49, 52–55, 58, 60, 65, 87

L3 Level 3. 13

LP Linear Programming. 79

MA Movement Authority. 8, 10–12, 20, 28

MB Moving Block. 8, 27

MCA Multi-Criteria Analysis. vii, ix, 44, 49, 53, 57, 58, 64, 69–74

MILP Mixed Integer Linear Programming. ii, ix, 2, 15, 18, 43, 45, 64, 69, 76, 77

MIP Mixed Integer Programming. ii, ix, 2, 18, 29, 36–38, 44, 76, 77, 83–85

NTC National Train Control. 7, 8

RBC Radio Block Centre. 8, 11, 12, 26, 30, 34, 81

SAAL Schiphol-Amsterdam-Almere-Lelystad. ii, vii, ix, 2, 14, 48, 50, 54

SMB Stop Marker Board. 10, 11, 20, 21, 27

ST Stop Train. viii, 48–50, 53, 72, 73, 83, 87

TIM Train Integrity Monitoring. ii, iii, vii–x, 1, 2, 10, 11, 29, 32–37, 40, 41, 43–47, 50–65, 69, 71, 74–76, 78, 87–89

TLD Train Length Dependent. 23, 27

TTD Trackside Train Detection. ii–iv, vii–x, 1, 2, 4, 7–12, 14–18, 20–29, 31–39, 41–45, 48–60, 62–65, 69–75, 78–81, 85, 87

VSS Virtual Sub-Section. ii, vii, 1, 7, 10–12, 14, 15, 17, 20, 21, 23, 25–28, 30, 32–38, 41, 45, 50, 52, 53, 56, 58, 59, 63, 64

1 Introduction

This chapter introduces the infrastructure design problem arising during migration to European Train Control System (ETCS) Hybrid Train Detection (HTD), formulates the research questions and scientific contribution, and outlines the structure of the thesis.

1.1. Background Context

The Netherlands operates the most intensively used railway network in the European Union (ProRail, 2025c) and one of the most intensively used worldwide, while demand for additional rail capacity continues to grow. Physical expansion is constrained by cost, limited space, and lengthy planning procedures, while maintenance and renewal needs already place substantial pressure on infrastructure budgets (ProRail, 2025d). Future capacity growth must therefore increasingly be achieved by using the existing railway infrastructure more efficiently.

Within this context, the Netherlands and other European countries are rolling out European Rail Traffic Management System (ERTMS) to support safer and more efficient railway operations (European Commission, Directorate-General for Mobility and Transport, 2026; ProRail, 2025a). Its train control component, ETCS, supervises train movements (European Commission, 2025). Conventional ETCS Level 2 (L2) uses continuous cab signalling and radio-based supervision but remains dependent on fixed Trackside Train Detection (TTD) for train detection and infrastructure release (European Union Agency for Railways et al., 2023b).

ETCS HTD, previously referred to as ETCS Hybrid Level 3, extends this concept by subdividing the underlying physical TTD sections into fixed Virtual Sub-Sections (VSSs) (EEIG, 2026). When confirmed train integrity is available, infrastructure can be released at VSS level, whereas movements behind trains without confirmed integrity remain dependent on the underlying physical TTD section. Train position is reported through the ETCS onboard system, while Train Integrity Monitoring (TIM) provides the required integrity confirmation, and physical train detection also remains relevant under degraded operating conditions (EEIG, 2026; European Union Agency for Railways et al., 2023b).

This hybrid release mechanism creates a staged infrastructure design problem. During partial TIM deployment, TTD boundaries can still support movements that depend on trackside release, but retaining them also incurs lifecycle cost, so their value can change as TIM deployment increases. Determining which TTD boundaries should remain, where they should be located, and for how long forms the central problem addressed in this research. Here, migration feasibility refers to realising a staged TTD trajectory from the initial infrastructure state while respecting boundary-placement requirements and progressive TIM deployment. This leads to the following main research question:

How can TTD boundary layouts under ETCS HTD be designed to balance capacity, infrastructure cost, and migration feasibility under mixed-traffic conditions?

To answer this main question, the following sub-questions are addressed:

1. What is the state of the art on methods for block section layout design for conventional and ETCS signalling?
2. Which infrastructure and operational requirements constrain the placement of TTD boundaries and block boundaries under ETCS HTD?
3. How do TTD boundary layouts and VSS-based release structures influence capacity and lifecycle cost under different traffic and TIM deployment conditions?
4. Which method can enable a cost-effective design of staged TTD boundary placement on a real-world corridor under mixed-traffic conditions?
5. Which guidelines and recommendations can be derived for the cost-effective and operationally stable design of TTD layouts under ETCS HTD?

1.2. Scientific Contribution

Previous research has established the capacity potential of ETCS HTD, but generally evaluates predefined infrastructure configurations rather than determining how the residual TTD layout should evolve during migration. No systematic method was identified that jointly considers spatial TTD design, partial TIM deployment, traffic heterogeneity, infrastructure inheritance, and migration cost. Addressing this gap requires signalling engineering to define the feasible design domain, blocking time theory to quantify capacity effects, operations research to determine the staged migration trajectory, and lifecycle economics to value infrastructure decisions over time. The main contributions are:

- **Formalisation of the ETCS HTD layout design hierarchy.** The study separates the TTD/VSS release hierarchy from the fixed and virtual operational block hierarchy and translates the associated engineering requirements into mandatory, forbidden, and optional TTD boundary positions.
- **Staged spatial optimisation of residual train detection.** A bi-objective Mixed Integer Programming (MIP), solved through an objective-equivalent Mixed Integer Linear Programming (MILP) reformulation, jointly determines which optional TTD boundaries remain or are modified as TIM deployment progresses.
- **Traffic-dependent valuation of TTD boundary locations.** Heterogeneous leader-follower interactions make the nominal capacity value of a TTD boundary dependent on its corridor location and traffic pattern rather than on boundary count alone.
- **Infrastructure inheritance as a migration-design factor.** Brownfield and greenfield comparisons show how inherited TTD assets affect the lifecycle cost of attaining comparable capacity performance.
- **Coupled fleet and infrastructure migration planning.** Endogenous TIM deployment shows that rolling-stock integrity deployment and residual TTD adaptation should be treated as coupled migration decisions.

The contributions are demonstrated on the Schiphol-Amsterdam-Almere-Lelystad (SAAL) corridor and complemented by microscopic RailSys evaluation, which shows that TTD boundaries with limited nominal capacity value may still contribute to disturbance recovery.

1.3. Report Structure

Chapter 2 introduces the technical and operational principles of ETCS HTD, followed by the literature review and research gap in Chapter 3. Chapter 4 derives the TTD and block layout requirements that define the design domain. Chapter 5 develops the analytical headway and cost models and formulates the staged optimisation framework. Chapter 6 defines the SAAL case study and RailSys evaluation, and Chapter 7 presents the resulting layouts and their cost, capacity, and operational stability performance. Chapter 8 interprets the findings and their limitations and implications, after which Chapter 9 answers the research questions and concludes the thesis.

2 Background

This chapter provides the theoretical and technical background required to understand the remainder of this thesis. It covers four topics: conventional fixed block signalling and trackside train detection (Section 2.1), blocking time theory and its relation to line capacity (Section 2.2), the ETCS application levels within ERTMS (Section 2.3), and the operational principles of ETCS HTD (Section 2.4). Readers already familiar with these concepts may proceed directly to the literature review in Chapter 3.

2.1. Conventional lineside railway signalling and trackside train detection

Conventional railway signalling separates consecutive trains by dividing the railway into fixed block sections. Train detection systems determine whether these sections are occupied or clear, while the interlocking uses this information to authorise only conflict-free train movements. The length and placement of block sections therefore influence both railway capacity and the amount of signalling and detection equipment required.

In this conventional arrangement, the different signalling functions operate as one chain. Physical train detection determines which track sections are occupied, the interlocking uses this information together with the required route and point positions to determine whether a movement can be authorised, and lineside signals communicate this authority to the driver. The driver operates the train accordingly, while onboard Automatic Train Protection (ATP) supervises compliance with the permitted speed and movement authority. The block layout provides the spatial structure within which these functions interact.

2.1.1. Block principle

A fundamental principle of fixed block signalling is that a following train may only enter a block section after the preceding train has completely cleared it. In conventional fixed block signalling, lineside signals communicate whether entry into the next block section is permitted, thereby maintaining a safe separation between consecutive trains (Pachl, 2021). Figure 2.1 illustrates this principle for three-aspect signalling.

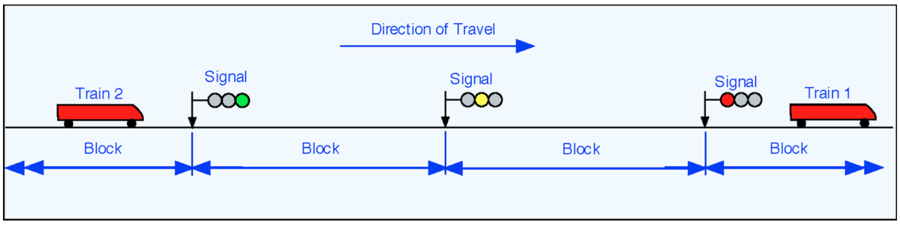

Figure 2.1: Three-aspect signalling principle (The Railway Technical Website, 2023).

Block section length depends on braking distance, permissible speed, train characteristics, safety margins, and operational requirements (Pachl, 2021). Shorter block sections can reduce train separation and increase capacity, but require additional signalling and train detection equipment. Conventional fixed block layout design therefore involves a trade-off between capacity and infrastructure cost (I. A. Hansen & Pachl, 2014). Realising this principle in operation requires a reliable means of detecting whether each section is occupied, which is the function of the physical train detection systems described next.

2.1.2. Track circuits and axle counters

Physical train detection provides this occupancy information, and conventional systems mainly use track circuits or axle counters, which measure occupation in fundamentally different ways. A track circuit uses the two rails of a section as an electrical circuit: when a train enters, its wheelsets short-circuit the rails and divert the current, which the system detects as an occupied section. It therefore senses train presence continuously along the whole section. An axle counter instead places detection points at both ends of a section and counts the axles passing each point, treating the section as occupied while more axles have entered than left and releasing it only once the two counts balance (Pachl, 2021).

Track circuits are constrained by the electrical characteristics of the rails and therefore have practical limits on section length. Axle counters are not subject to the same electrical limitation and can support longer detection sections, although their placement remains constrained by signalling and infrastructure requirements. Both technologies operate according to fail-safe principles: uncertainty or failure results in the affected section being treated as occupied (Pachl, 2021).

This thesis represents the physical boundaries created by these systems as TTD boundaries, whose number and location influence both the granularity at which infrastructure can be released and the amount of physical detection equipment that must be installed and maintained.

2.1.3. Interlocking principle

The interlocking then uses this train detection information to prevent conflicting train movements. Before setting a route, it verifies that the required track sections are clear and that points and other infrastructure elements are safely configured and locked (Pachl, 2021). These functions impose additional constraints on the placement of detection and block boundaries, particularly around stations, junctions, and switches.

2.1.4. Automatic train protection

Whereas the interlocking authorises movements from the trackside, ATP adds a separate onboard safety layer that supervises train speed against the permitted speed profile and enforces braking when required (Flammini, 2013). The trackside signalling and interlocking system therefore determines the permitted movement, while onboard ATP supervises compliance with that authority. Together, these detection, interlocking, and protection functions determine how trains occupy and release the infrastructure over time, which the next section formalises through blocking time theory.

2.2. Blocking time theory

The block layout introduced above ultimately determines how closely trains can follow one another. The minimum headway between successive trains partly governs railway capacity. Blocking time theory provides a microscopic time-distance framework to quantify this relationship by modelling how trains occupy and release infrastructure over time and space (I. A. Hansen & Pachl, 2014; Pachl, 2021). In conventional fixed block signalling, the signalling system releases the track behind the train block by block: it releases a section only after the TTD confirms that the entire train has cleared it. The block structure therefore dictates the minimum train separation and thus the line capacity.

2.2.1. Blocking time of a block section

For each block section, the theory models the blocking time of a block section as the sum of six components, illustrated in Figure 2.2 (I. A. Hansen & Pachl, 2014; Pachl, 2021):

1. **Signal control time**: the time for the interlocking to set and lock the route, including moving the switches, and to clear the signal ahead.
2. **Signal sight and reaction time**: the time for the driver to observe and react to the signal aspect.
3. **Approach time**: the time for the train to cover the distance from the sighting point to the entry of the block section, corresponding to the braking distance at the entry speed.
4. **Running time**: the time between block signals.
5. **Clearing time**: the time for the rear of the train to pass the clearing point at the block section

exit, so that the train detection system registers the block section as vacated.

6. **Release time**: the processing time for the signalling system to confirm that the section is clear and to release it for the following train.

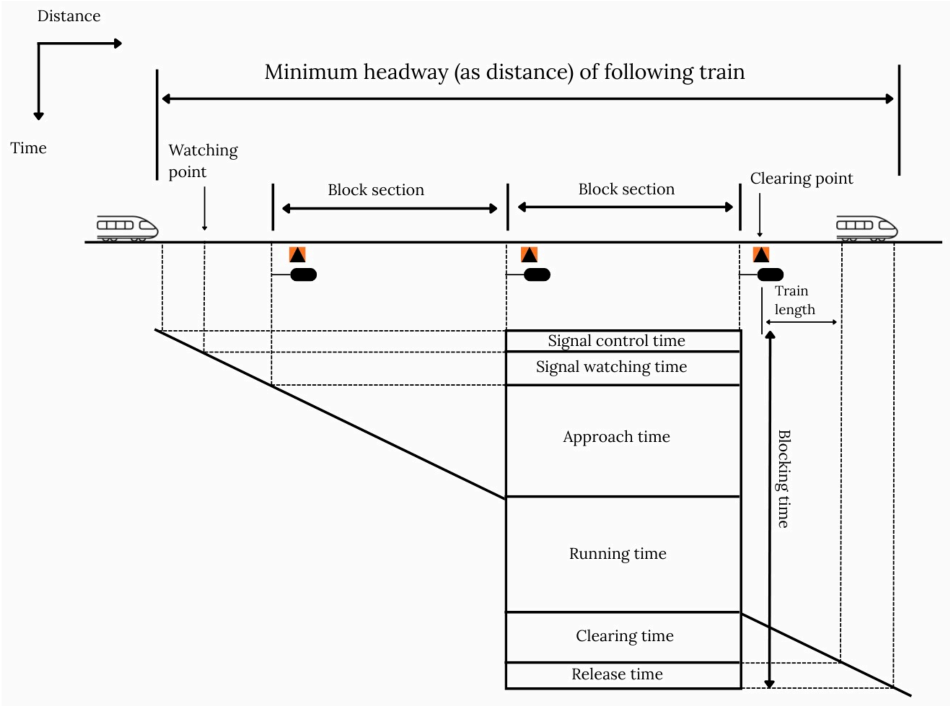

Figure 2.2: Blocking time components in fixed block signalling, based on Pachl (2021).

2.2.2. Blocking time stairway and minimum line headway

Whereas the previous subsection considered a single block section, a corridor consists of multiple consecutive block sections, each with its own blocking time. Placing the blocking times of all sections a train traverses into a time-distance diagram produces the blocking time stairway. Compressing the blocking time stairways of the leading and following trains and finding the tightest feasible separation then yields the minimum headway between two consecutive trains (Figure 2.3).

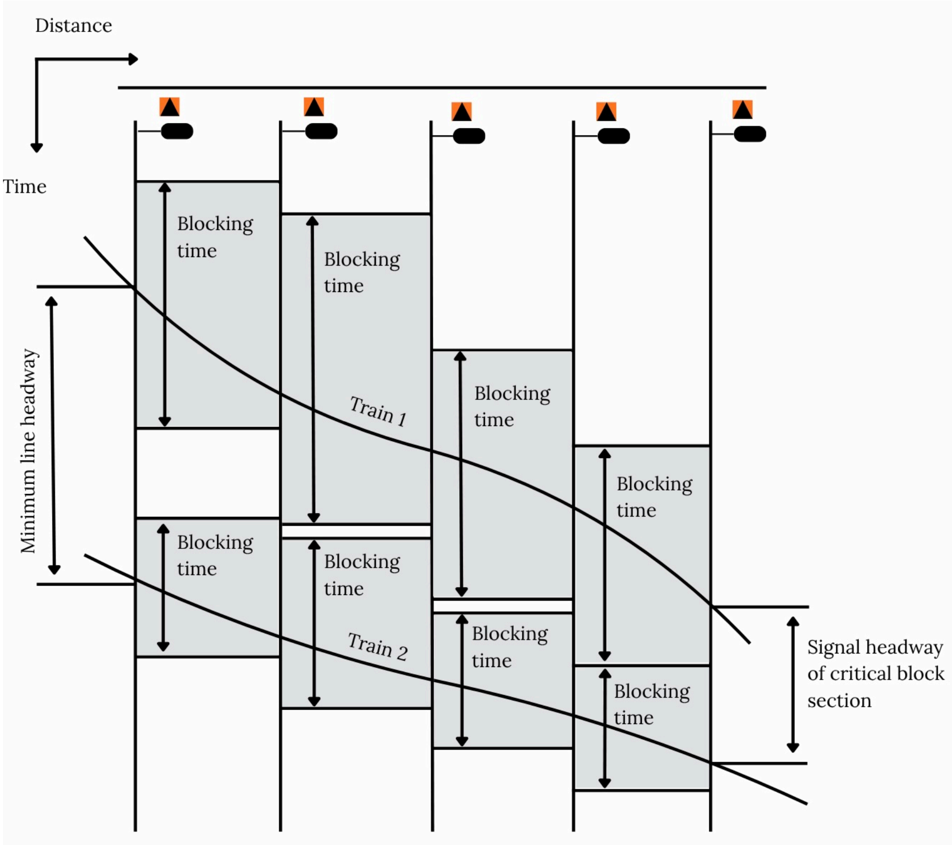

Figure 2.3: Blocking time stairway of two consecutive trains: the signal headway of the critical block section determines the minimum line headway, based on Pachl (2021).

In practice, each block section contributes to the blocking time stairway of a train. Compressing the blocking time stairway of two successive trains determines the minimum feasible headway. The block section that becomes binding in this compression is the critical block section, which determines the minimum line headway for the analysed train sequence. The minimum line headway is a fundamental capacity measure: the theoretical maximum number of trains per hour on a corridor is inversely proportional to it, so reducing the minimum line headway is equivalent to increasing line capacity.

Having established how fixed block layouts determine train separation under conventional lineside signalling, the next section introduces ERTMS/ETCS and explains how train control and supervision change while the underlying fixed block separation principle can remain applicable.

2.3. European Rail Traffic Management System and European Train Control System (ERTMS/ETCS)

As introduced in Chapter 1, ERTMS is the European railway traffic management system whose core train control component, ETCS, supervises train movements and enforces movement authorities (European Commission, 2025). Whereas conventional lineside signalling communicates movement authorities through physical signals along the track, ETCS transfers signalling information and train supervision to the onboard system. Under ETCS L2, movement authorities are transmitted by radio and presented to the driver in the cab, while the onboard equipment continuously supervises the train against the permitted speed profile and movement authority. The signalling interface therefore changes substantially, even though train separation may still rely on the same fixed block principle when infrastructure release remains dependent on physical train detection.

The legal interoperability framework for ERTMS/ETCS is established by the Control-Command and Signalling Technical Specification for Interoperability (CCS TSI), which applies to onboard and trackside control-command and signalling subsystems and refers to the mandatory technical specifications for ETCS, including the ETCS System Requirements Specification SUBSET-026 (European Commission, 2023; European Union Agency for Railways et al., 2023b).

Within this framework, European Union Agency for Railways (ERA) supports the harmonisation and interoperability of the European railway system, while the ERTMS Users Group (EUG) represents the railway user side in the specification process. The EUG is organised as a European Economic Interest Group (EEIG). In this context, EEIG denotes the legal form of the EUG, rather than a separate technical stakeholder.

The ETCS application levels used in this thesis follow the current ERTMS/ETCS specifications. An application level describes the operating relationship between track and train and depends on the trackside equipment, the means by which information is transmitted between track and train, and the allocation of functions between trackside and onboard equipment (European Union Agency for Railways et al., 2023a, 2023b). The currently defined application levels are Level 0, Level National Train Control (NTC), Level 1, and Level 2. Level 3 is no longer defined as a separate ETCS application level (European Commission, 2023; European Union Agency for Railways et al., 2023a). This thesis therefore distinguishes between the official ETCS application levels and the train detection and train separation configurations and concepts relevant to ETCS Level 2 operation:

- **ETCS Level 0** applies when an ETCS-equipped train operates in an area where the trackside is neither equipped with operational ETCS nor with an operational national train control system. The ETCS onboard equipment therefore provides only the functionality available for Level 0 operation, while train detection and train separation are not provided by ETCS (European Union Agency for Railways et al., 2023a, 2023b).
- **ETCS Level NTC** applies when an ETCS-equipped train operates under the supervision of an existing national train control system. The ETCS onboard equipment interfaces with the national system, for example through a Specific Transmission Module (STM), where applicable (European Union Agency for Railways et al., 2023a, 2023b). This thesis uses Level NTC as the current specification term. Older sources may refer to related implementations as Level STM or Level STM/NTC.
- **ETCS Level 1 (L1)** is overlaid on conventional lineside signalling and uses Eurobalises, and optionally Euroloop or radio infill, to transmit movement authorities and other signalling information to the train. Conventional trackside means remain responsible for determining train position and integrity (European Union Agency for Railways et al., 2023a, 2023b).
- **ETCS L2** uses radio communication to transmit movement authorities to the train. In contrast to Level 1, continuous lineside signalling is not required for normal operation. Importantly, the current Level 2 definition allows train-reported position and train integrity information to be used by the trackside to determine the parts of the track occupied by the train (European Union Agency for Railways et al., 2023a, 2023b). Within this Level 2 framework, this thesis distinguishes three train detection and train separation configurations or concepts relevant to the study:
  - **ETCS L2 fixed block** refers in this thesis to conventional Level 2 operation in which train detection and train separation rely on physical trackside train detection. Infrastructure behind a train is released only after the relevant TTD section is confirmed clear. Consequently, the granularity of train separation is determined by the physical TTD layout (European Union Agency for Railways et al., 2023b; Vignali et al., 2020).
  - **ETCS HTD** combines fixed, pre-configured virtual blocks with a limited installation of physical trackside train detection. The EUG defines a VSS as a subdivision of a physical TTD section whose occupation status is determined using information from both train position reports and trackside train detection. Trains reporting confirmed integrity can support release at VSS level, whereas trains without confirmed integrity remain dependent on the underlying physical TTD for release (EEIG, 2026). The use of fixed virtual blocks is a deliberate feature of the current HTD concept. Sufficiently short virtual sections can provide performance comparable to Moving Block while limiting the impact on existing trackside systems (EEIG, 2026). Section 2.4.1 explains this hierarchy in more detail.

– **Moving Block (MB)** is treated in this thesis as a train separation concept that can be implemented within the current ETCS Level 2 framework, rather than as a separate ETCS application level. SUBSET-023 defines a moving block as a block whose length is determined by the position of the train occupying the track ahead (European Union Agency for Railways et al., 2023a). Unlike fixed block and HTD operation, Moving Block therefore does not require train separation to be tied to predefined fixed block boundaries. Historically, this functionality was commonly associated with ETCS Level 3. Under the current CCS TSI, however, an ETCS Level 2 trackside application with no or reduced trackside train detection is explicitly identified as what was formerly ETCS Level 3 (European Commission, 2023). Level 3 therefore no longer constitutes a separate ETCS application level, although the term remains present in older literature and ongoing research programmes (UNIFE, 2024).

The resulting distinction is important for the remainder of this thesis. ETCS Level 0, Level NTC, Level 1, and Level 2 are official ETCS application levels, whereas ETCS L2 fixed block, ETCS HTD, and Moving Block describe different train detection and train separation configurations or concepts relevant within the Level 2 framework. The transition from ETCS Level 1 to conventional ETCS Level 2 removes the dependency on lineside signals for normal operation but does not, by itself, remove the fixed block separation principle. In the ETCS L2 fixed block configuration considered in this thesis, track release remains dependent on physical TTD sections. ETCS HTD modifies this mechanism by introducing fixed virtual subdivisions and the use of reported train integrity, as described in Section 2.4.

2.3.1. Blocking time in ETCS L2 fixed block operation

Section 2.2 introduced the general blocking time model for conventional fixed block signalling. Under ETCS L2, the same fixed block separation principle can still apply, but lineside signals no longer convey movement authority. Instead, the interlocking determines which route it can safely set, the Radio Block Centre (RBC) translates this information into a Movement Authority (MA), and the onboard ETCS equipment supervises the train against the permitted speed profile and the braking curve towards the End of Authority (EoA).

In ETCS L2 fixed block operation, the system still uses trackside train detection for infrastructure release. It releases a block section only after the complete train has cleared the section and the TTD system has confirmed that the section is free. Therefore, the underlying fixed block separation principle remains unchanged. The main difference compared with fixed block lineside signalling lies in the interpretation of the approach time component.

Figure 2.2 showed the blocking time components in conventional fixed block signalling, including signal control time, signal sight and reaction time, approach time, running time, clearing time, and release time. Figure 2.4 shows the corresponding blocking time interpretation for a cab signalling line under ETCS L2 fixed block operation.

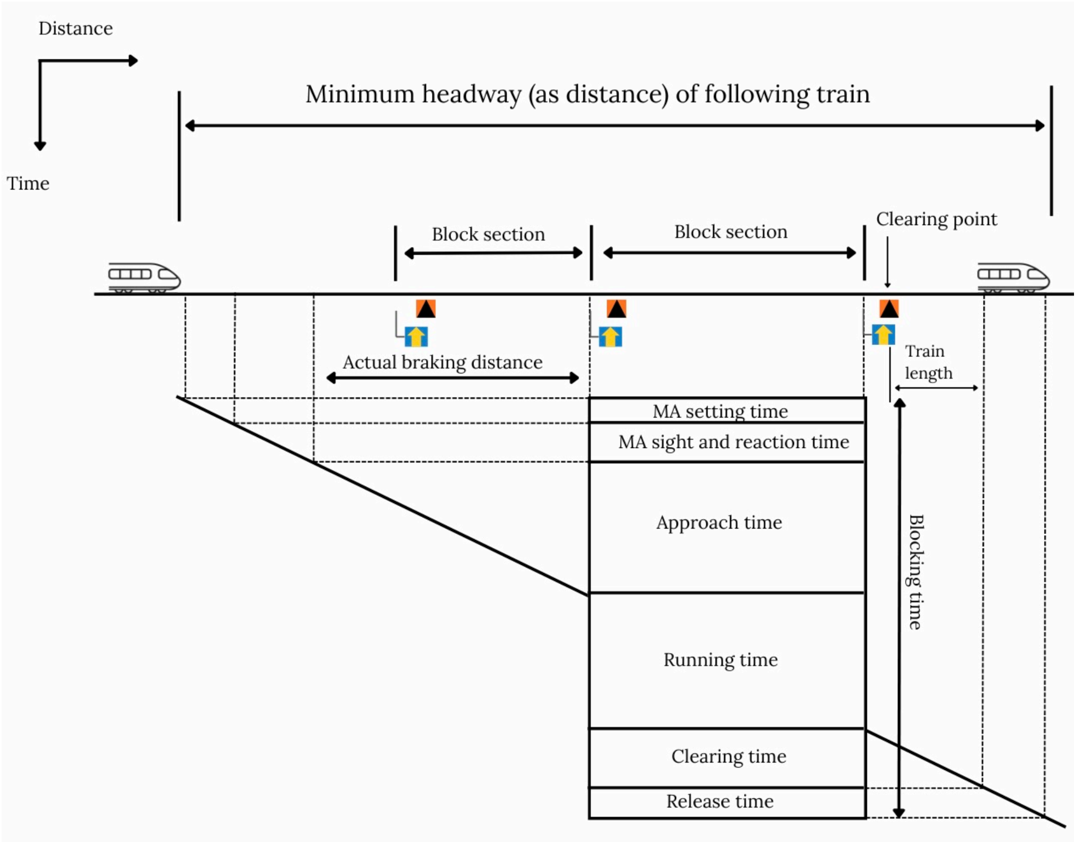

Figure 2.4: Blocking time components in ETCS L2 fixed block operation, based on Pachl (2021).

Following Pachl (2021), the key difference calculating the blocking time between fixed block lineside signalling and ETCS L2 fixed block operation is that the approach time no longer depends on the distance between the signal at the entrance of the block section and the signal in rear that provides the approach indication. Instead, the running time over the actual braking distance determines the approach time, based on the braking curves of the cab signalling system. Since lineside signals are absent, an MA setting time, which includes the time to move, set, and lock the switches along the route, replaces the signal control time, and an MA sight and reaction time replaces the signal sight and reaction time: the time the driver needs to react to the movement authority once it has been notified and acknowledged. The other blocking time components remain conceptually equivalent to those of a fixed block system with lineside signals.

The analytical model developed in this thesis expresses the blocking time of a block section under ETCS L2 fixed block operation as:

$$t_{\text{blocking}} = t_{\text{setting}} + t_{\text{sr}} + t_{\text{approach}} + t_{\text{run}} + t_{\text{clear}} + t_{\text{release}} \tag{2.1}$$

This blocking time is usually considerably longer than the time the train physically occupies the block section: the occupation is only the running time t_{run}, whereas the blocking time also includes the route setting, reacting, approach, clearing, and release times around it.

In ETCS L2 fixed block operation, TTD boundaries support the operational block boundaries, because releasing a block requires trackside confirmation that the corresponding detection section is clear. However, not every TTD boundary necessarily forms an operational block boundary. The TTD layout therefore directly influences the block structure and, through the blocking time components, the achievable minimum headway. Unlike in conventional fixed block signalling, the actual ETCS braking distance towards the EoA governs t_{approach}, not the length of the previous block section. This is the key blocking time difference between lineside signalling and ETCS L2 cab signalling.

2.4. ETCS HTD

The previous subsection established ETCS L2 fixed block operation as the reference case. ETCS HTD, based on the EUG specification for ETCS L2 (EEIG, 2026), modifies this reference case by introducing VSSs that enable earlier infrastructure release and therefore closer train separation. However, whether a train can benefit from VSSs depends on the availability of confirmed train integrity. Train position is reported through the ETCS onboard system, while TIM provides the integrity confirmation required for VSS level release. Trains without this capability still rely on conventional TTD release. This coexistence is what makes the system hybrid, and it is the reason why the system must retain physical TTD sections alongside VSSs. The following subsections introduce the section and block hierarchy under ETCS HTD and explain how this hierarchy affects the release mechanism and blocking time.

2.4.1. Section and block hierarchy

Figure 2.5 illustrates the four layer hierarchy of sections and blocks under ETCS HTD. The following subsections explain each layer and how they interact.

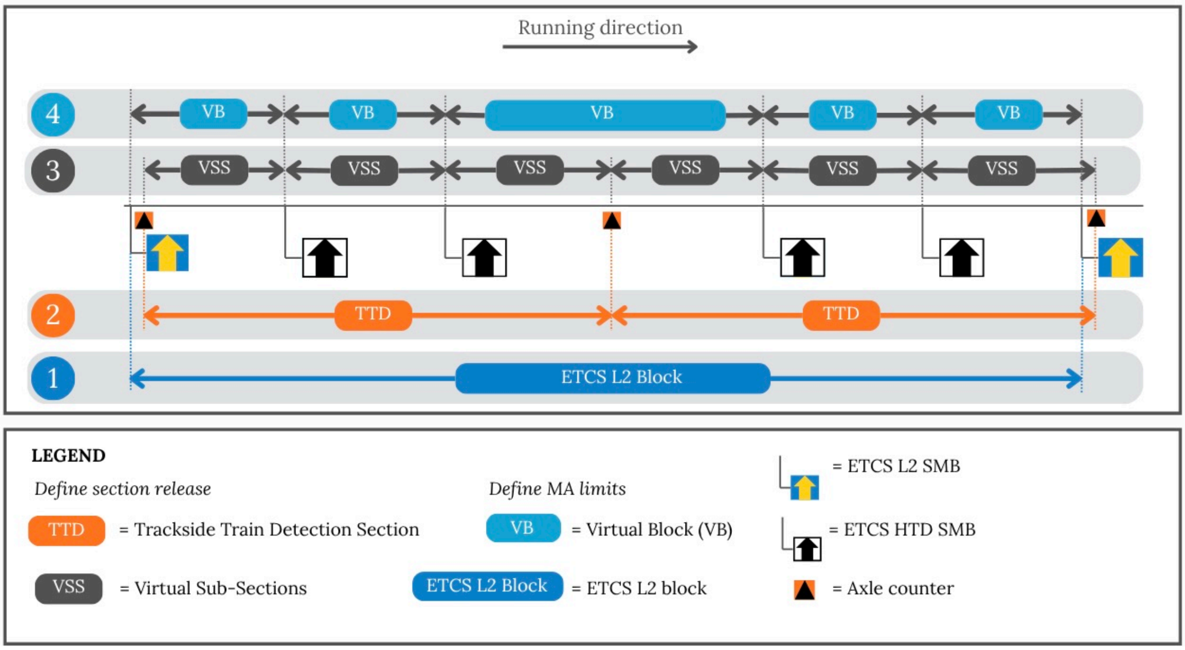

Figure 2.5: Section and block hierarchy under ETCS HTD, showing the relationship between the ETCS L2 fixed block (layer 1), physical TTD sections (layer 2), VSS (layer 3), and virtual blocks (layer 4).

ETCS L2 fixed block release

Under ETCS L2 fixed block operation, train separation relies entirely on physical trackside detection. An ETCS L2 fixed block section (layer 1 in Figure 2.5) runs from one ETCS L2 Stop Marker Board (SMB) to the next ETCS L2 SMB in the direction of travel (European Union Agency for Railways, 2022). A MA authorises a train to proceed up to a defined end point on the track. Under ETCS L2 fixed block operation this end point corresponds to the SMB at the end of the block. The following train may not proceed beyond this SMB until the detection system has confirmed the block free.

Placing an ETCS L2 SMB requires a corresponding trackside detection boundary to support the block release logic. In axle counter-based layouts, this corresponds to an axle counter at that location, although the physical detector may sit with a small design offset relative to the SMB. However, the reverse does not hold: an axle counter may sit without a corresponding SMB, for example to provide detection coverage around infrastructure elements such as switches or level crossings, where the interlocking needs occupancy information but trains should not stop. Section 4.1 discusses these situations further.

TTD sections and VSSs
Before introducing the block release mechanism under ETCS HTD, this section must establish the section definitions. The EUG defines a TTD section (layer 2 in Figure 2.5) as a section delimited by conventional trackside train detection systems such as axle counters (EEIG, 2022). A VSS (layer 3) is a fixed section that subdivides a TTD section, but unlike a TTD section, no axle counter or other trackside detection system physically marks it. Instead, the RBC software stores the VSSs, and determines their occupation status by combining onboard position and train integrity reports with the physical trackside detection information (EEIG, 2022). Every TTD boundary therefore coincides with a VSS boundary, but not every VSS boundary is a TTD boundary. The EUG uses the terms *TTD border* and *VSS border* to classify VSS boundaries: a TTD border is a VSS boundary that coincides with a TTD section boundary, and a VSS border is one that does not (EEIG, 2022). This thesis uses the general term *boundary* throughout (e.g. TTD boundary, VSS boundary) to denote the location where two adjacent sections meet, without implying the specific EUG classification.

It is important to note that TTD sections and VSSs are not by definition operational block sections. They define the physical and virtual detection granularity of the infrastructure, but whether a particular section boundary also serves as a block boundary (i.e. a location where a MA may end) depends on whether a SMB sits at that location. A VSS is therefore not automatically a virtual block. Instead, the SMB configuration defines a virtual block, which may consist of one or more VSSs.

ETCS HTD virtual block release
With the section and block definitions established, the release mechanism can be described. Under ETCS HTD, the ETCS onboard system reports train position, while TIM provides train integrity confirmation. This thesis refers to trains carrying TIM as TIM-equipped trains. When a train is actively reporting confirmed integrity, it operates as an *integrity reporting* train. All other trains, including trains without TIM and TIM-equipped trains whose TIM is temporarily unavailable (for example due to a fault or communication loss), operate as *non integrity reporting* trains. Block release under ETCS HTD therefore depends on the integrity reporting status of the leading train:

- **Non integrity reporting trains** cannot confirm their integrity to the RBC. Track release therefore falls back to the ETCS L2 fixed block mechanism: the detection system must confirm the entire fixed block (from SMB to SMB) free before the following train may enter. The release granularity therefore corresponds to the physical TTD section, while the MA remains governed by the applicable fixed block boundaries, resulting in longer train separation.
- **Integrity reporting trains** enable virtual block release. The RBC can confirm individual VSSs as free based on the train's reported position and confirmed integrity, so the system releases infrastructure at VSS level rather than at TTD level. The RBC can then extend the MA of the following train to the next available ETCS HTD block boundary, which is supported by the finer VSS release information and results in shorter train separation.

A virtual block (layer 4 in Figure 2.5) runs from one SMB to the next SMB in the direction of travel. These virtual blocks may correspond either to the fixed block structure supported by physical TTD boundaries or to the finer ETCS HTD block structure supported by VSS level release. The fixed block boundaries define the MA limits for non integrity reporting trains, while the ETCS HTD boundaries define the finer virtual block boundaries for integrity reporting trains. Both release mechanisms coexist on the same infrastructure, and the applicable level depends on the integrity reporting status of the leading train. The safety justification for releasing VSSs without direct trackside confirmation rests on the TIM system: without confirmed integrity, detached vehicles could remain undetected if the system based release only on train reported position (EEIG, 2026).

The key distinction is therefore between the section hierarchy and the block hierarchy. The section hierarchy (TTD sections and VSSs) defines the release granularity: the smallest unit of infrastructure that the system can confirm free and release behind a train. The block hierarchy (fixed blocks and virtual blocks) defines the MA boundaries: the locations where a train's MA may end, i.e. the EoA locations associated with SMBs.

2.4.2. Blocking time under ETCS HTD

Figure 2.4 introduced the blocking time components for ETCS L2 fixed block operation. The same component structure applies to ETCS HTD: the t_{approach} remains based on the actual ETCS braking distance towards the EoA, and the RBC still generates and supervises the MA. The difference is not in t_{approach}, but in the effective running time component t_{run}, which depends on the release mechanism applicable to a given pair of consecutive trains.

Figure 2.6 illustrates this. In ETCS L2 fixed block operation, shown at the top of the figure, the occupied and released infrastructure ties to the TTD section, so the time to traverse the full TTD section determines t_{run}. Under ETCS HTD, shown in the lower part, the VSSs subdivide the physical TTD section. For a pair in which the leading train reports confirmed integrity, the RBC can release a VSS once the rear of that train has cleared it, resulting in a shorter t_{run}. For a pair in which the leading train does not report confirmed integrity, release remains tied to the TTD boundary and t_{run} corresponds to the full TTD section.

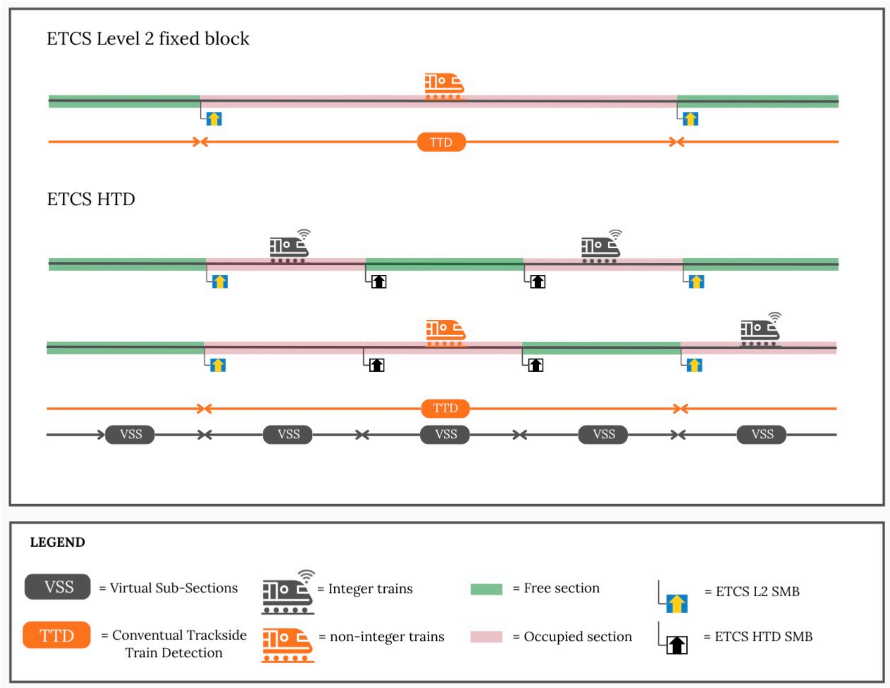

Figure 2.6: Release mechanism in ETCS L2 fixed block operation and ETCS HTD.

The effective t_{run} therefore varies between pairs of consecutive trains. Chapter 5 formalises this pair-dependent release mechanism, representing it analytically by the VSS length when the leading train reports confirmed integrity and by the full TTD length when it does not, and uses it to formulate the TTD boundary placement problem.

2.4.3. Implications for physical TTD layout design

The pair-dependent release mechanism creates a trade-off in the physical TTD layout. Additional TTD boundaries provide finer release for movements that remain dependent on physical train detection, but also increase infrastructure cost. As reliance on VSS level release increases, the nominal capacity value of this additional physical subdivision can therefore change. This relationship motivates the TTD layout design problem examined in the following chapters.

3 Literature Review

This chapter reviews the literature relevant to the design of block section layouts under ETCS HTD. It has three objectives. First, it reviews existing principles and methods for block section layout design. Second, it examines how these methods can be applied in the context of ETCS HTD. Third, it identifies the research gap addressed by this thesis.

The literature search focused primarily on peer-reviewed publications retrieved through Scopus and ScienceDirect, with Google Scholar and the TU Delft institutional repository used to identify additional relevant studies. The search terms covered signalling technology, block and train detection structure, and optimisation, and included "ETCS Level 2", "hybrid train detection (HTD)", "Hybrid Level 3 (L3) (HL3)", "train integrity monitoring (TIM)", "trackside train detection (TTD)", "virtual subsection (VSS)", "virtual block", "moving block", "block section layout", "signalling layout optimisation", and "capacity optimisation". Relevant studies were further identified through backward and forward snowballing by reviewing both the references cited by key publications and later publications that cited them.

In addition to peer-reviewed literature, the review includes the official ETCS HTD specification documents issued by the ERTMS Users Group EEIG (2022, 2026) and research deliverables from the Europe's Rail R2DATO programme (Aoun, 2024, 2025, 2026a). These sources are particularly relevant because they define the current operational principles, engineering requirements, and deployment considerations for ETCS HTD, which are not yet covered in comparable detail by the peer-reviewed literature.

3.1. Block layout design methods

Before reviewing the ETCS HTD-specific literature, this section examines fixed block signalling studies that provide methodological foundations for block section layout design. These studies show how block boundaries can be positioned or spaced to meet capacity requirements while limiting the amount of signalling infrastructure. Quaglietta (2014) formulates signalling block layout design as a cost-capacity problem. In the proposed equi-block framework, block sections in capacity-critical zones determine the minimum line headway, whereas block sections in non-critical, or "stretchable", zones can be lengthened without increasing that headway. The method uses simulation-based optimisation to minimise lifecycle-related costs while satisfying a required capacity level and evaluates the resulting layouts under stochastic operational disturbances. Although the case study concerns a fixed block metro application rather than ETCS HTD, the distinction between capacity-critical and stretchable locations provides a conceptual basis for examining where physical detection boundaries contribute to capacity and where additional boundaries provide limited benefit. Vignali et al. (2020) develop an ETCS L2-specific backward design procedure that determines Marker Board and axle counter positions from prescribed contractual headway requirements. The method combines an analytical blocking time model with OpenTrack-based headway assessment and accounts for ETCS braking curves, route setting and release times, train characteristics, and infrastructure and operational constraints. It evaluates deterministic two-train headway scenarios and iteratively adjusts the layout until the required headways are satisfied, while limiting the amount of signalling equipment. The method therefore designs a physical fixed block layout for prescribed performance requirements, but does not consider ETCS HTD. Wagter (2019) provides a station-focused complement using RailSys microscopic simulations for Schiphol and Almere. In these two case studies, the minimum headway under undisturbed (unhampered) conditions is governed by the critical block section, while modifying other, non-critical block sections does not reduce the governing headway. This illustrates the importance of identifying location-specific bottlenecks rather than shortening block sections uniformly. However, the study is limited to station areas and does not address corridor-wide physical detection layout design, infrastructure cost, or ETCS HTD. Together, these studies show that block layout design should distinguish between locations where block boundaries affect the governing headway and locations where additional subdivision provides little capacity benefit.

They also provide two complementary methodological perspectives: simulation-based optimisation of cost and capacity, and deterministic backward design from prescribed headway requirements. These principles are relevant to the present research, but they do not resolve the TTD layout problem under ETCS HTD. The next section therefore reviews the ETCS HTD-specific literature and examines which parts of this problem have already been addressed.

3.2. ETCS Hybrid train detection

This section reviews the literature on ETCS HTD and closely related next-generation signalling concepts. The studies are discussed according to their contribution to capacity assessment, infrastructure design, and operational modelling, with particular attention to the distinction between VSS layout and TTD layout.

3.2.1. Capacity and operational evidence

Simulation studies show that the capacity benefit of ETCS HTD depends strongly on where virtual subdivision is introduced and on the surrounding traffic and infrastructure conditions. Knutsen et al. (2024) show that shorter VSSs can reduce technical headways, while the corresponding network level improvement can remain limited by station and switch bottlenecks and timetable margins. Their subsequent comparison with moving block confirms this spatial dependence: selectively shortening virtual blocks in lower-speed sections can approach moving block capacity performance more closely than applying a uniform virtual block length across the corridor (Knutsen et al., 2025). Both studies nevertheless treat the physical TTD layout as predefined.

For the Dutch context, Jansen (2019) directly demonstrates the interaction between virtual subdivision and residual physical detection. Shorter VSSs reduce infrastructure occupation, whereas reducing the physical TTD layout partly offsets this benefit for movements behind trains without confirmed integrity. Vergroesen (2020) reports a similar location-dependent effect on the SAAL corridor, where the benefit of finer subdivision varies with service pattern and infrastructure configuration. Together, these studies indicate that the operational value of physical and virtual subdivision cannot be inferred from section length alone, but depends on traffic composition and corridor location.

A complementary research stream considers implementation and operational modelling. D. Hansen et al. (2020) demonstrate how virtual occupation states can be derived from train position, train integrity information, and physical detection states, thereby establishing the practical feasibility of the hybrid concept. Versluis et al. (2025), meanwhile, show that finer virtual discretisation can improve conflict-resolution performance but also increases computational complexity. These studies reinforce the operational relevance of section granularity, although neither determines the residual physical TTD layout.

Taken together, the capacity literature establishes three mechanisms that are central to the present research: finer virtual release can reduce headway, the benefit is spatially and traffic dependent, and reducing physical detection can become restrictive while movements without confirmed integrity remain. What remains unresolved is how the number and location of residual physical TTD boundaries should be selected systematically while accounting for infrastructure cost.

3.2.2. Infrastructure placement constraints

The EUG Principles and Engineering Guideline define the operational and engineering conditions that limit design freedom under ETCS HTD. The HTD concept uses fixed pre-configured virtual blocks for trains that report confirmed integrity, while retaining a limited amount of trackside train detection for trains that cannot provide integrity confirmation and for degraded situations (EEIG, 2026). TTD therefore remains part of the HTD design even when trains with confirmed integrity are separated using VSSs. The Engineering Guideline recommends TTD borders near infrastructure and operational locations where physical train detection remains important. These include points, level crossings, movable bridges, tunnels with safety systems, borders to non-HTD trackside and shunting areas, and locations associated with Start or End of Mission, splitting and joining, system restart, loss of train integrity or communication, and uncertainty in the reported train position (EEIG, 2022). At such locations, TTD information can support reliable infrastructure release, detect trains that are not communicating with the trackside, and limit the propagation of uncertain virtual occupation states.

Deliverable D37.2 then translates these recommendations into a practical section design framework. In this framework, TTD sections containing points, level crossings, or movable bridges are designed as a single VSS, while similar treatment is applied at borders to non-HTD trackside and shunting areas and at locations where trains regularly change direction (Aoun, 2025). Other infrastructure conditions affect the usefulness of additional virtual subdivision without necessarily imposing strict placement rules. For example, high gradients and catenary overlap zones can restrict where particular trains are allowed to stop. Deliverable D37.2 therefore notes that adding VSS boundaries at such locations may provide little operational benefit, while also stating that these conditions do not justify general restrictions on VSS placement because they can depend on train characteristics (Aoun, 2025). The EUG Principles also allow long TTD sections between constrained infrastructure areas to be divided into as many virtual sections as required for the intended performance (EEIG, 2026). Deliverable D37.2 follows the same principle by first placing required TTD borders and any required TTD = VSS sections before using the remaining space for the preferred virtual section design (Aoun, 2025). The document does not define a generally preferred number or length of intermediate TTD sections and identifies their relation to the share of TIM-equipped trains as an open design issue. The main remaining design freedom therefore concerns which additional TTD boundaries should be retained between constrained locations. This is the design freedom addressed in the present research, where TTD layout choices are evaluated in terms of cost, capacity, and operational stability.

3.2.3. Design methods for ETCS HTD

Engels et al. (2023) provide one of the most directly relevant optimisation-based methods for VSS layout design under ETCS HTD. Their MILP determines the minimum number of VSS boundaries required to make a given timetable feasible for an existing TTD layout. Building on this work, Engels et al. (2025) formalise several design tasks for ETCS HTD and analyse their computational complexity. They show that timetable verification is NP-complete and that the considered VSS generation, schedule optimisation, and capacity optimisation tasks are NP-hard. These results provide a theoretical basis for treating HTD layout design as a combinatorial optimisation problem. To reduce this computational burden, Engels and Wille (2024) propose a sequential optimisation pipeline. Train routing is first determined under a moving block representation, after which the resulting routes are fixed and the VSS layout is generated in a second optimisation step. In the initial benchmark experiments, this separation substantially reduces runtime compared with jointly optimising routing and VSS placement, while producing the same or close to optimal solutions in almost all tested instances. The authors note, however, that the sequential approach has no theoretical guarantee of preserving the global optimum. These studies provide formal and optimisation-based methods for virtual layout design, but the TTD layout remains fixed. A complementary design stream is provided by the Europe's Rail R2DATO programme. Deliverable D37.2 develops an engineering-oriented section design methodology that first imposes required physical detection locations and subsequently designs the virtual subdivision (Aoun, 2025). It distinguishes several VSS design regimes but does not prescribe a generally preferred number or spacing of intermediate TTD boundaries, explicitly leaving this dependent on the infrastructure context and the availability of confirmed train integrity.

Deliverable D37.3 complements this design framework with a simulation protocol in which TTD length, VSS length, and train integrity availability are varied and assessed through nominal capacity and deterministic disturbance recovery (Aoun, 2024). The resulting scenarios are evaluated across several European case studies in D37.4, which confirms that virtual subdivision is most effective around capacity-critical locations and that reduced physical detection can become restrictive when movements remain dependent on TTD level release (Aoun, 2026b). At deployment level, D37.1 further identifies capacity and cost as the two highest-weighted criteria for HTD implementation (Aoun, 2026a).

Together, these deliverables provide engineering rules, evaluation procedures, and empirical evidence, but they stop short of a mathematical method that endogenously selects the staged residual TTD layout.

3.3. Synthesis and research gap

The studies reviewed above differ not only in signalling context, but also in methodological purpose: some provide capacity assessment methods, some evaluate specific corridor scenarios, some formalise or validate ETCS HTD principles, and others define engineering constraints or optimisation approaches.

Table 3.1 summarises these works by objective, method, and their relation to the present research.

Table 3.1: Overview of key studies on block section layout design and their relation to the present research.

Study	**Objective**	**Method**	**Relation to present research**
Block layout design for other signalling systems (not ETCS HTD)			
Quaglietta (2014)	Minimise total costs under a headway constraint for fixed block railway signalling.	Simulation-based optimisation for fixed block signalling.	Provides methodological background for cost-constrained block layout design in a non-ETCS HTD signalling context; the cost-capacity framing inspires the present approach.
Vignali et al. (2020)	Backward size ETCS L2 block lengths from a prescribed headway target.	Analytical and iterative block design method for ETCS L2.	Provides methodological background for ETCS L2 block sizing, but concerns ETCS L2 fixed block rather than ETCS HTD.
Wagter (2019)	Assess the impact of block layout on station headways under ETCS.	Microscopic simulation in RailSys for ETCS L2 station areas.	Relevant for identifying capacity-critical station and switch areas, but concerns ETCS L2 station envelopes rather than ETCS HTD layout design.
ETCS HTD-specific studies			
Engels et al. (2023)	MILP for VSS placement to achieve timetable feasibility under ETCS HTD.	MILP with continuous position variables.	Optimises VSS for schedule feasibility, not TTD length for cost-capacity trade-off; ratio of TIM-equipped trains not varied.
Engels and Wille (2024)	Decomposed pipeline for ETCS HTD design.	Sequential optimisation pipeline.	Supports decomposition of ETCS HTD design tasks, but does not address the number and placement of intermediate TTD boundaries or the ratio of TIM-equipped trains.
Engels et al. (2025)	Analyse the computational complexity of ETCS HTD design tasks.	Complexity analysis.	Shows the computational difficulty of several HTD design tasks and motivates efficient formulations and decomposition approaches.
D. Hansen et al. (2020)	Validate ETCS HTD principles and demonstrate Virtual Block Function behaviour.	Formal B model; real-life demonstration.	Safety and system-behaviour focus; no capacity evaluation or layout optimisation.
Versluis et al. (2025)	Impact of track discretisation granularity on conflict detection and resolution.	MILP-based conflict detection and resolution; sensitivity analysis.	Traffic management and rescheduling focus; no TTD layout optimisation method.

Continued on next page

Table 3.1: Overview of key studies on block section layout design and their relation to the present research continued

Study	**Objective**	**Method**	**Relation to present research**
Jansen (2019)	Assess physical TTD reduction and its capacity consequences on the Utrecht-Den Bosch corridor.	RailSys simulation; scenario analysis.	Directly investigates physical TTD reduction and its capacity consequences; a key precursor, but the reduced-TTD layout is predefined rather than optimised.
Vergroesen (2020)	Assess ETCS HTD with Automatic Train Operation (ATO) on the SAAL corridor.	RailSys simulation; scenario analysis.	Adds the HTD+ATO and corridor-specific capacity perspective for the SAAL corridor, but evaluates a predefined TTD layout.
Knutsen et al. (2024)	Assess the effects of VSS length and train integrity availability on headway, capacity utilisation, and punctuality.	Conceptual two-train analysis and RailSys network simulation.	Evaluates predefined HTD configurations, but does not optimise either VSS or TTD boundary placement.
Knutsen et al. (2025)	Compare the capacity performance of ETCS HTD and moving block.	Conceptual headway analysis and RailSys network simulation.	Shows where shorter VSSs provide additional benefit, but does not optimise the physical TTD layout or include infrastructure cost in the design problem.
Aoun (2026a)	Position ETCS HTD within deployment and standardisation strategy.	Stakeholder-oriented deployment framework; SWOT, multi-criteria analysis, and roadmapping.	Clarifies implementation criteria such as capacity, cost, maturity, and standardisation readiness; does not provide a technical optimisation method for TTD layout.
Aoun (2025)	Define a methodology for optimised placement of virtual and physical blocks under ETCS HTD.	Engineering framework based on mandatory TTD locations and VSS design regimes.	Provides an operational ETCS HTD design methodology, but leaves the number and placement of intermediate TTD boundaries as a context-dependent design choice.
Aoun (2024)	Specify requirements for capacity and robustness simulations of ETCS HTD scenarios.	Scenario-based simulation protocol with variables such as TTD length, VSS length, and ratio of TIM-equipped trains.	Defines how ETCS HTD configurations should be evaluated, but does not generate an optimal layout.
Aoun (2026b)	Evaluate the impact of ETCS HTD configurations using simulation.	Simulation-based impact assessment.	Provides evidence on the effects of ETCS HTD design variables, but does not formulate a cost-capacity optimisation model for TTD placement.

Continued on next page

Table 3.1: Overview of key studies on block section layout design and their relation to the present research continued

Study	Objective	Method	Relation to present research
EEIG (2026); EEIG (2022)	Define ETCS HTD operating principles and engineering rules.	System specification and engineering guideline.	Define ETCS HTD states, engineering constraints, and TTD-related considerations, but provide no method for selecting the TTD layout under mixed-integrity operation.

The reviewed literature reveals a consistent research gap. Conventional and ETCS L2 fixed block studies provide established principles for locating capacity-relevant block boundaries, while ETCS HTD research has established the operational principles, capacity effects of virtual subdivision, and engineering constraints on physical detection. However, no method was identified that endogenously selects the number and location of intermediate physical TTD boundaries under partial train integrity availability while explicitly balancing railway capacity against infrastructure cost.

The remaining design freedom is concentrated mainly between infrastructure locations at which physical detection is required for safety or operational reasons. Within these less constrained areas, the capacity value of an additional TTD boundary depends on both its location and the traffic using the corridor, while its economic value also changes over the migration horizon. Residual TTD design is therefore not only a section-sizing problem, but a staged location selection problem.

This thesis addresses this gap by formulating staged TTD boundary placement as a bi-objective MIP, solved through an objective-equivalent MILP reformulation. An analytical model based on blocking time represents the expected headway under mixed fast and slow track-releasing movements, while discounted lifecycle and deployment cost represents the competing economic objective. The resulting Pareto frontier provides alternative migration paths, after which the selected layouts are evaluated microscopically in RailSys for nominal capacity and operational stability.

Theoretically, this thesis adopts blocking time theory as its analytical lens but uses its spatial bottleneck principle in a different design setting. Classical blocking time theory identifies the critical block section after an operational block layout has been defined (**Pachl2021**; I. A. Hansen & Pachl, 2014). Under ETCS HTD, however, the operational block architecture remains fixed while residual TTD boundaries can be redesigned and the applicable release granularity depends on the integrity status of the leader. The present framework therefore evaluates blocking time over the relevant infrastructure release section and treats optional TTD boundaries as endogenous spatial decisions.

This changes boundary criticality from a descriptive property of a fixed signalling layout into a traffic-, deployment-, and location-dependent design property. Three implications follow. First, boundary count alone is insufficient to describe the nominal capacity value of residual TTD because layouts with the same number of boundaries can produce different governing headways. Second, the value of an optional TTD boundary decreases as fewer leader interactions depend on TTD-level release. Third, traffic heterogeneity changes the spatial distribution of this value through cumulative leader-follower running time differences. This perspective enables TTD layout, traffic composition, and TIM deployment to be analysed jointly and distinguishes the present study from HTD approaches that evaluate predefined TTD layouts or optimise only the virtual subdivision.

3.4. Key performance indicators in the literature

The literature identifies cost, capacity, and operational stability as important performance dimensions for assessing railway signalling layouts, particularly in the context of ETCS HTD. This choice is supported by Deliverable D37.1, in which railway experts and stakeholders assessed seven criteria relevant to the deployment of HTD. Capacity and cost received the highest importance weights, at 25% and 21%, respectively, while robustness ranked fourth at 15% (Aoun, 2026a).

Cost refers to the investment and operating costs associated with a signalling layout, while capacity describes the amount of train traffic that the infrastructure can accommodate over time (Aoun, 2026a). These two criteria are particularly relevant for HTD because reducing trackside train detection infrastructure can lower costs, while the resulting physical layout can also affect railway capacity.

Capacity under regular operation does not, however, capture operational stability under disturbed conditions. Deliverable D32.5 considers operational stability as a separate performance dimension and relates it to the capacity available to absorb operational variation and to the recovery of railway operations following a deterministic disturbance (Koopman - van Welij, 2026).

Based on the literature, cost, capacity, and operational stability are selected as KPI 1, KPI 2, and KPI 3, respectively. Section 5.6 explains how these KPIs are applied in this research.

4 Layout design requirements under ETCS HTD

Before the optimisation framework of Chapter 5 can determine where physical TTD boundaries should be placed, this chapter must first identify and classify the design requirements that constrain the ETCS HTD layout. This chapter addresses Research Question 2 by classifying these requirements as safety-critical (SC), operationally critical (OC), or general recommendations (GR). Together, these requirements define the feasible design domain for the layout optimisation developed in Chapter 5. Although the classification is established for ETCS HTD, requirements that are not specific to VSS behaviour or virtual block design remain applicable to conventional fixed block signalling, ETCS L1, and ETCS L2, because infrastructure and interlocking constraints drive them, rather than the signalling system alone.

Section 2.4.1 introduced the section and block hierarchy under ETCS HTD. Figure 4.1 reproduces this hierarchy because the placement requirements in this chapter apply to different layers of the infrastructure.

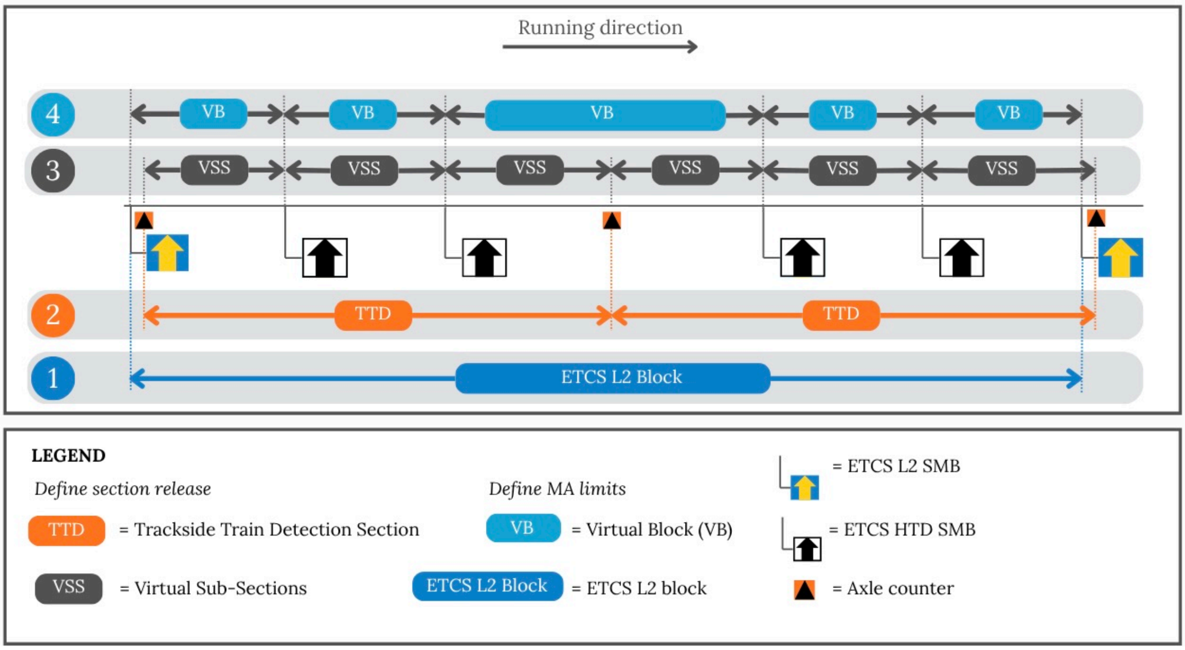

Figure 4.1: Section and block hierarchy under ETCS HTD (reproduced from Figure 2.5).

The key distinction is between sections and blocks. TTD sections and VSSs determine the granularity at which infrastructure can be confirmed free and released behind a train. Operational blocks instead determine where a train's MA may end, typically at a SMB. A section boundary is therefore not necessarily an operational block boundary.

The applicable release level depends on whether the leading train can confirm its integrity. A train without confirmed integrity remains dependent on physical TTD sections for infrastructure release. A train reporting confirmed integrity can release infrastructure at the finer VSS level. In both cases, the movement authority of the following train remains limited by the applicable operational block boundaries. Chapter 5 formalises these two cases as slow and fast track-releasing pairs and derives their effect on headway.

This distinction is important for the requirements developed in this chapter because constraints may apply either to release section boundaries or to operational block boundaries. Every TTD boundary coincides with a VSS boundary, but not every VSS boundary is a TTD boundary, and neither type of section boundary necessarily requires a SMB. This study optimises the physical TTD layout on a predefined VSS grid. It does not optimise the placement of operational block boundaries.

The remainder of this chapter therefore distinguishes between requirements on the physical TTD and VSS layout and requirements on operational block boundaries. Section 4.1 identifies where physical detection must or should be retained, while Section 4.2 identifies where operational block boundaries may or may not be placed.

4.1. TTD layout requirements under ETCS L2 HTD

Under ETCS HTD, trains that report confirmed integrity can release infrastructure at the finer VSS level, whereas trains without confirmed integrity remain dependent on physical TTD sections for release. As the share of trains able to report confirmed integrity increases, a larger proportion of the traffic can therefore use VSS level release and fewer trains remain dependent on the physical TTD layout. This creates the opportunity to reduce the number of TTD boundaries while retaining sufficient physical detection for trains without confirmed integrity, degraded operation, and infrastructure locations where trackside detection remains required or beneficial (EEIG, 2022, 2026). The TTD layout cannot be designed independently of the other layers of the ETCS HTD layout. Infrastructure and operational requirements can constrain the placement of TTD boundaries, the subdivision of TTD sections into VSSs, and the locations at which section boundaries may also function as operational block boundaries. These layers are interdependent, because a requirement imposed at one layer can restrict the feasible design choices at another. Reducing or relocating TTD boundaries therefore requires first identifying the infrastructure and operational requirements that constrain the overall ETCS HTD layout. The EUG Engineering Guideline identifies a range of infrastructure and operational considerations for the design of VSSs and TTD sections (EEIG, 2022). Deliverable D37.2 builds on this guidance and translates it into an engineering framework that includes required TTD borders, required TTD = VSS sections, guidance for the subdivision into VSSs, and additional locations where physical detection may be beneficial (Aoun, 2025). Together with conventional fixed block requirements, these considerations also constrain which section boundaries can be used as block boundaries. Although the existing guidance distinguishes between required and optional design considerations and gives priority to safety-critical requirements, it does not provide the three-level classification used in this study. This chapter therefore classifies the identified layout requirements as safety-critical (SC), operationally critical (OC), or general recommendations (GR). The classification is based on the EUG Engineering Guideline, the ETCS L2 HTD Principles, the R2DATO D37.2 methodology, and conventional Dutch signalling requirements where applicable. Tables 4.1-4.3 summarise the resulting classification and specify the implementation and design purpose of each aspect. The SC and OC aspects are subsequently discussed in detail because they are treated as fixed constraints in the layout design, whereas the GR aspects are advisory and do not impose mandatory placement in the optimisation.

Table 4.1: Overview of the safety-critical TTD placement requirements.

Aspect	Implementation	Design benefit
Points (switches)	TTD boundary on both sides of the switch zone, see Figure 4.2	Safe route setting, locking and release
Level crossings	TTD boundary on both sides of the protected zone, see Figure 4.2; TTD section coincides with single VSS	Safe activation and release of crossing protection
Movable bridges	TTD boundary on both sides of the protected zone, see Figure 4.2; TTD section coincides with single VSS	Safe activation and release of bridge protection
Overhead catenary overlap zones	TTD boundary in the middle of the catenary overlap zone, see Figure 4.2	Prevents block occupation leakage during maintenance
Borders to non-ETCS L2 HTD trackside	TTD boundary at the transition point, see Figure 4.2	Reliable handover between ETCS L2 HTD and conventional detection

Table 4.2: Overview of the operationally critical TTD placement requirements.

Aspect	Implementation	Design benefit
Borders between traffic control areas	TTD boundary at the control boundary, as illustrated in Figure 4.2	Prevents unknown VSS states from propagating into adjacent control areas; reduces dispatcher coordination and traffic control system complexity

Table 4.3: Overview of the general recommendations for TTD placement.

Aspect	Implementation	Design benefit
Tunnels with safety systems	TTD before and after tunnel; TTD section coincides with single VSS	Reliable input for tunnel safety systems
Borders to permanent or temporary shunting areas	TTD at entry; first section as TTD section coincides with single VSS	Prevents virtual occupation spill over from shunting areas
Locations where SoM or EoM may occur	TTD at mission boundary; TTD section coincides with single VSS	Reliable ETCS initialisation and shutdown
Locations where the impact of a system restart must be low	TTD upstream of critical area	Faster recovery after RBC restart
Locations with frequent change of direction	TTD at platform end (departure side)	Earlier release of downstream infrastructure
Interfaces where non-ERTMS vehicles may enter ETCS L2 HTD trackside	TTD at access point; TTD section coincides with single VSS	Prevents unverified virtual occupation by non-ERTMS vehicles
Locations where the impact of loss of train integrity or communication must be low	TTD near junctions and stations	Limits propagation of ambiguous VSS states
Locations where uncertainty of train position should be reduced	TTD upstream of conflict points	Resets virtual rear end determination

4.1.1. Safety critical aspects

Points (switches)

Switch areas impose specific requirements on the placement of TTD. For safe route setting, route locking, and route release, the interlocking must be able to verify that the switch zone is unoccupied before a conflicting route can be authorised. This requirement for dedicated detection coverage of the switch area applies regardless of the signalling system, since in all cases the interlocking relies on trackside detection information to confirm that the switch area is free before setting a route. Pachl

(2021) specifies the general rules governing which switch detection sections may be combined into a single track section and which must remain separated, depending on the switch geometry and the conflicting routes involved. TTD therefore bounds the switch zone on both sides of the turnout so that the interlocking can reliably determine the occupancy state of the switch area, creating a protected detection zone that fully encloses the switch, as illustrated in Figure 4.2. The mathematical point in the figure is the theoretical intersection of the gauge faces at the crossing and serves as the geometric reference point of the turnout. In Dutch practice, TTD boundaries are placed within 6 m upstream of the frontside switch and at least 5 m downstream of the backside switch, plus the axle counter offset defined at the start of this chapter (ProRail, 2023). The red dashed boundaries in the figure indicate this protected zone. Unlike level crossings, movable bridges, or overhead catenary overlap zones, trains are generally permitted to stop on switches. Consequently, no extended downstream clearance zone is required purely to prevent stopping on the infrastructure element. For switches, TTD boundaries bound the protected detection zone so that the interlocking can determine whether the switch area is occupied. Unlike level crossings and movable bridges, this does not imply that the design procedure necessarily treats the switch zone as a single TTD = VSS section.

Level crossings and movable bridges

Level crossings and movable bridges require reliable detection of approaching, occupying, and clearing trains in order to safely activate, maintain, and release the protection of the infrastructure element. This study therefore treats these locations as protected infrastructure zones that require dedicated train detection coverage (Pachl, 2021). In Dutch practice, the protected zone extends approximately 25 m on each side of the infrastructure element, and the TTD boundaries enclosing it are placed at these limits for infrastructure release, as illustrated in Figure 4.2. Within this protected zone, this study configures the TTD section to coincide with a single VSS, so no additional VSS boundaries subdivide the protected zone internally. This ensures that one unambiguous occupancy state represents the level crossing or movable bridge for protection activation and release (EEIG, 2022).

Because a train may not come to a stand within the protected zone, no block boundary is configured inside this section. The next block boundary on the downstream side, which is the first boundary that can contribute to capacity, is placed a Train Length Dependent (TLD) clearance beyond the protected zone. In the Dutch context this corresponds to 270 m, plus the 9 m axle counter offset. An axle counter placed closer than this may still serve infrastructure release, but it offers no capacity benefit as a block boundary, since a train cannot be authorised to stop there in normal operation.

Overhead catenary overlap zones

Overhead catenary overlap zones separate traction power phases or systems and must be traversed without electrically bridging adjacent supply sections. Trains should therefore not stop within these zones (European Commission, 2014). In Dutch signalling practice, a TTD section is installed in the middle of the overlap zone to enforce this constraint: if the signalling system detects a train within the overlap, it can prevent subsequent trains from entering. This TTD section also serves a maintenance purpose. When one side of the overlap zone is electrically isolated for work on the overhead line equipment, the detection section allows infrastructure managers to monitor whether trains or vehicles are approaching the work area from the energised side (ProRail, 2025b). As with level crossings and movable bridges, a train may not come to a stand within the overlap or its downstream protected zone of approximately 25 m. The next block boundary on the downstream side, which is the first one that contributes to capacity, is therefore placed a TLD clearance (270 m in the Netherlands) plus the 9 m axle counter offset beyond the protected zone. An axle counter placed closer than this offers no capacity benefit, since a train cannot stop there in any case.

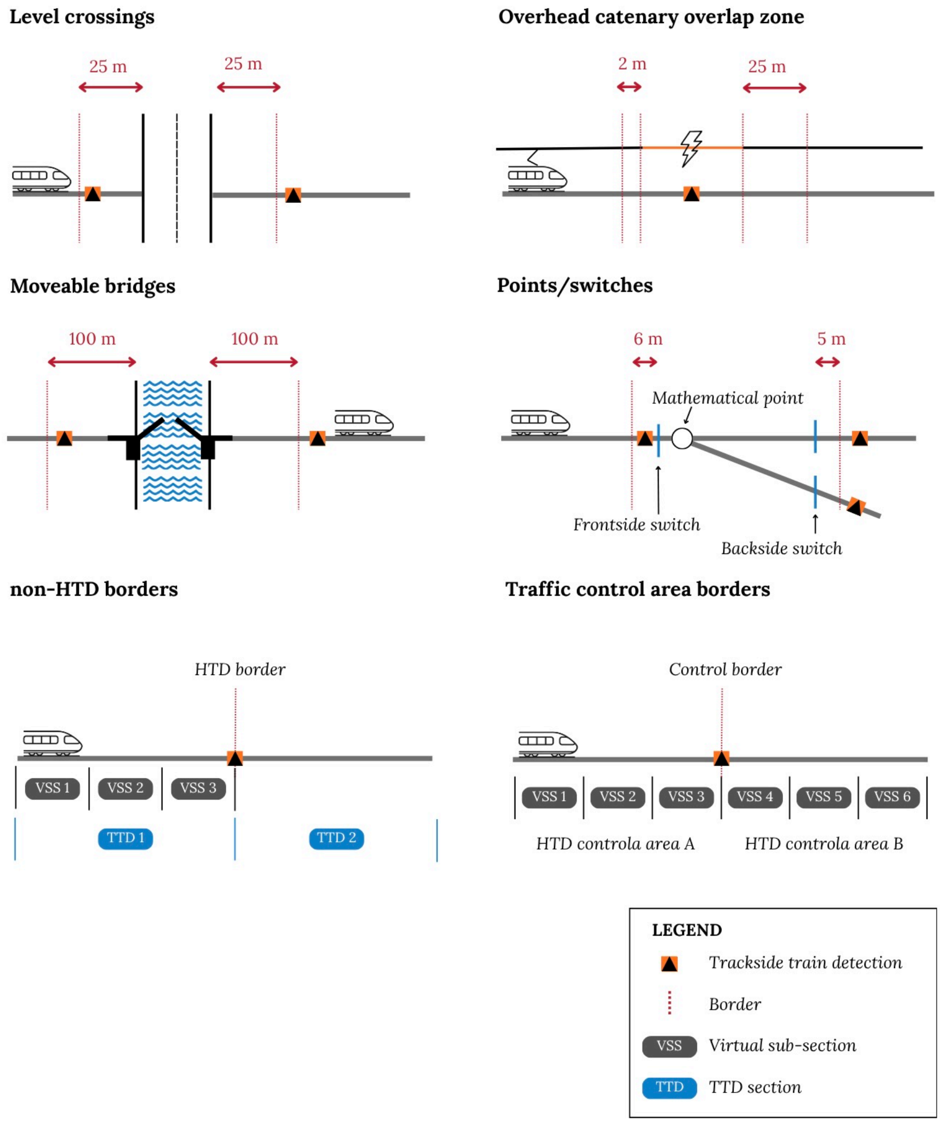

Figure 4.2: Overview of the TTD boundary placement requirements.

Borders to non-ETCS L2 HTD trackside

Transitions between ETCS L2 HTD equipped infrastructure and conventional non-ETCS L2 HTD trackside sections represent safety-critical interfaces for train detection. Within an ETCS L2 HTD area, sections behind a train may be confirmed free and released using onboard train integrity information and train position reports. However, this assumption cannot hold once trains enter infrastructure that does not support ETCS L2 HTD operation. For this reason, the transition between ETCS L2 HTD and non-ETCS L2 HTD trackside is typically aligned with a TTD boundary, as illustrated in Figure 4.2. This ensures that track occupancy can be reliably determined using physical detection and that the transfer between ETCS L2 HTD-based virtual release and conventional train detection can be handled safely.

4.1.2. Operationally critical aspects

Borders between traffic control areas

Borders between traffic control areas are operationally critical interfaces under ETCS L2 HTD. Within a physical TTD section, an unknown VSS state may arise during degraded situations, for example after train reset, loss of train integrity information, communication loss, or system recovery. If a TTD section crosses the boundary between two traffic control areas, this unknown VSS state may propagate into the adjacent control area until it reaches a free TTD boundary. Introducing a TTD boundary at the border between traffic control areas limits this propagation, as illustrated in Figure 4.2. The boundary therefore acts as an operational firewall between control domains. This is not primarily a safety requirement, but an operational and implementation requirement. If an unknown VSS state propagates into an adjacent control area, it may create a non-standard signalling or traffic control indication for the neighbouring dispatcher, requiring additional coordination between traffic controllers and increasing workload during degraded operation. In addition, allowing VSS state propagation across traffic control area boundaries would make implementation in the traffic control systems more complex, since the virtual occupation logic would have to be exchanged and interpreted consistently across control domains. A TTD boundary at the control boundary therefore localises degraded VSS states within one control area, simplifies dispatcher responsibility, and reduces the complexity of implementing ETCS L2 HTD logic in the traffic control system.

4.1.3. General recommendations

The general recommendations in Table 4.3 are derived from the HTD engineering considerations on the use of TTD information for transitions, degraded operation, SoM and EoM, non-equipped trains, and position uncertainty (Aoun, 2025; EEIG, 2022, 2026).

Tunnels with safety systems

Tunnels equipped with integrated safety systems (such as fire detection, ventilation control, evacuation zones, and emergency stop circuits) may benefit from a physically confirmed track occupancy state. Within ETCS L2 HTD operation, a VSS may remain in an ambiguous or unknown state after loss of train integrity confirmation or communication failure. In tunnel environments, where safety systems rely on reliable knowledge of train presence, such uncertainty may be undesirable. For this reason, it is often recommended to place a TTD boundary directly before and after the tunnel. Within the tunnel, additional TTD and VSS boundaries may still be required for capacity reasons or because the tunnels can contain other infrastructure elements such as switches or catenary overlap zones that impose their own placement requirements.

Borders to permanent or temporary shunting areas

Shunting areas form a transition between ETCS L2 HTD controlled area and regions where trains may not report integrity, may reverse direction, or may operate without ETCS supervision. When entering ETCS L2 HTD trackside from such areas, the integrity status of a train may initially be uncertain until trackside detection has confirmed the rear of the train clear. For this reason, it is generally recommended that the boundary between ETCS L2 HTD trackside and a shunting area coincides with a TTD boundary. In such configurations, the first section after entering the mainline may effectively function as a section where the TTD section coincides with a single VSS, since virtual release based on onboard information may not yet be reliable. This arrangement helps prevent the propagation of uncertain or virtual occupancy states from the shunting area into the mainline infrastructure.

Locations where Start of Mission (SoM) or End of Mission (EoM) may occur

SoM and EoM procedures require a clearly identifiable train position to correctly initiate or terminate ETCS supervision. When relying solely on VSS states, some uncertainty in the confirmed release state may remain, particularly after degraded situations or communication interruptions. For this reason, it is generally recommended that locations where SoM or EoM may occur coincide with a TTD boundary. In such configurations, the corresponding section may effectively function as a section where the TTD section coincides with a single VSS, providing a physically confirmed reference point that supports reliable ETCS initialisation or termination.

Locations where the impact of a system restart should be limited

After an RBC restart or system recovery, VSS states may temporarily revert to ambiguous or unknown. Consequently, if no physical segmentation is available, these uncertain states may propagate over large portions of the line. For this reason, it is often recommended to introduce a TTD boundary upstream of areas where the operational impact of such situations should remain limited. The TTD section can constrain the uncertainty zone to a short, physically verified segment, thereby supporting faster recovery of normal operation.

Locations with frequent change of direction

Station environments require particular attention in ETCS L2 HTD configurations. Around platforms, trains frequently stop, change direction, or initiate new operational missions. In such situations, virtual occupation states generated by the ETCS L2 HTD virtual section logic may propagate further along the track than operationally desirable. For this reason, it is often recommended to introduce a physical TTD boundary at the end of the platform in the direction of departure. This detection boundary provides an independent confirmation that the train has cleared the station track and can limit the propagation of virtual occupation states. As a result, downstream infrastructure is less likely to remain constrained by virtual occupation states longer than operationally necessary, which supports operational stability around station areas.

Interfaces where non-ETCS vehicles may enter ETCS L2 HTD trackside

Maintenance vehicles and other non-ETCS vehicles typically do not report position or train integrity information. Without physical detection, these vehicles could keep VSSs virtually occupied. Therefore, it is generally recommended that entry points where non-ETCS vehicles may access the line coincide with a TTD boundary. In such configurations, the first section inside the ETCS L2 HTD area may effectively function as a section where the TTD section coincides with a single VSS, ensuring that entering vehicles are physically detected before ETCS L2 HTD logic based on onboard information is applied. This helps prevent unverified vehicles from causing virtual occupancy propagation.

Locations where the impact of loss of train integrity or communication should remain limited

Loss of train integrity confirmation or communication may cause VSS states to propagate virtual occupancy states backwards along the route. Without physical segmentation, this propagation may therefore affect a large portion of the infrastructure. To mitigate this effect, it is often recommended to introduce TTD boundaries at locations where the operational impact of such situations should remain limited, typically near junctions, station areas, or complex switch zones. These boundaries can interrupt the propagation of virtual occupancy states and limit the affected zone to a smaller infrastructure segment.

Locations where uncertainty of train position should be reduced

In locations where higher localisation accuracy is desirable, such as merging routes, conflict points, or complex switch areas, virtual localisation may benefit from additional physical detection reference points. For this reason, a TTD boundary may be introduced upstream of such critical areas. The corresponding TTD section can act as a physical reference that helps reset the virtual rear end determination and reduce accumulated positional uncertainty.

Locations where trains not equipped with ERTMS (e.g. yellow fleet) could enter ETCS L2 HTD trackside

In areas with (occasional) entry of non-ETCS vehicles, such as road-rail access points or maintenance sidings, physical detection may be required to maintain consistency of ETCS L2 HTD occupancy states. It is therefore often recommended to place a TTD boundary at the access point. The first section inside the ETCS L2 HTD domain may effectively function as a section where the TTD section coincides with a single VSS, ensuring that any entering vehicle is physically detected before ETCS L2 HTD logic is applied. This prevents untracked vehicles from causing persistent virtual occupation states.

4.1.4. System behaviour constraints

In addition to the categorised design aspects, TTD placement must take into account the behaviour of the ETCS HTD system under degraded conditions. Under conditions such as communication loss between train and RBC, loss of train integrity information, or RBC restart, onboard reports alone can no longer resolve the VSS states in the affected area. When communication is lost, the VSSs occupied by the train are set to unknown and this state propagates to adjacent sections until it reaches a free TTD (EEIG, 2026). A similar propagation mechanism applies after loss of train integrity. In both cases, physical TTD boundaries act as firewalls that stop the propagation of ambiguous occupancy states. After an RBC restart, the trackside has no knowledge of train positions and must rely on TTD information to re-establish a safe state (EEIG, 2022, 2026). The classified aspects in Tables 4.1-4.3

largely capture these considerations: the safety-critical requirements already impose TTD boundaries at junctions and other conflict points, while the general recommendations address placement upstream of critical areas and near stations to support faster recovery. Although Table 4.3 lists the placement of a TTD boundary downstream of platforms in the departure direction as a general recommendation, it also serves to limit the impact of degraded situations on station areas by providing a physical reset point for uncertain VSS states.

4.2. Block layout requirements under ETCS L2 HTD

Under ETCS HTD, MBs rather than lineside signals indicate block boundaries. The EUG considers the installation of such MBs at VSS or TTD boundaries optional, although not installing them may impact the operational rules (EEIG, 2022). A TTD or VSS boundary without an SMB is used only for infrastructure release, not as a block boundary. Once an SMB is added, the same boundary also becomes a block boundary at which an EoA is placed. A block boundary is a point at which a train may legitimately come to a stand, so it must not be positioned where stopping is not permitted, nor within an infrastructure element that must remain fully protected during train movements. A release-only boundary, such as those enclosing the protected zones described above, carries no such restriction: it serves only to report whether a train is present within the element, for example on a movable bridge, without allowing it to stop there. The following subsections summarise the block boundary restrictions, which carry over from conventional signalling practice and remain directly applicable under ETCS HTD.

1. Level crossings and movable bridges

Level crossings and movable bridges define protected infrastructure zones where trains must not stop under normal operating conditions. In conventional signalling, Dutch design rules prohibit placing block boundaries within these protected zones (ProRail, 2025b). For virtual block design under ETCS HTD, the same restriction applies to any VSS boundary configured as a marker board location: the corresponding EoA must not fall within the protected part of a level crossing or bridge. In conventional lineside signalling, the next block boundary downstream of such elements is placed only after the entire train has cleared the protected zone. This requires a TLD clearance, which in the Dutch context is the applicable design value of 270 m (as indicated in Figure 4.3). Under ETCS, this principle is retained, with an ETCS SMB replacing the lineside signal. Under ETCS HTD, the same restriction applies to any VSS or TTD boundary used as a block boundary with a SMB. A boundary at a VSS position requires only virtual detection, whereas a boundary at a TTD position additionally requires a physical detection system.

2. Overhead catenary overlap zones

Overhead catenary overlap zones enable pantograph transitions between power supply sections, and a stationary train must not electrically bridge them (European Commission, 2014). Consequently, conventional block boundaries are never placed within these zones (ProRail, 2025b). Under ETCS HTD, the same rule applies to VSS boundaries configured as marker board locations: an EoA within the overlap zone could create a virtual stopping point inside the electrically sensitive area, which is operationally prohibited. The same TLD downstream clearance as described for level crossings and movable bridges applies.

3. Points (switches)

Switch areas impose specific constraints on block boundary placement. Automatic block signals cannot protect sections containing movable track elements or locations where conflicting train movements may occur (Pachl, 2021). Consequently, block boundaries are not located within the switch zone itself, which must instead be protected by controlled signals governed by an interlocking system. For virtual block design, the same applies to VSS boundaries configured as marker board locations: an EoA should not fall within the switch area, where a train standstill could interfere with conflicting movements or complicate degraded mode behaviour. A physical TTD boundary may still exist within or around the switch for interlocking and route release. The design treats such detection as a fixed infrastructure requirement rather than a capacity-driven choice, and therefore excludes it from the layout optimisation. Once the switch area is fully cleared, block boundaries may be placed downstream without an additional clearance distance (ProRail, 2023).

4. Station platforms

In conventional practice, station platform areas do not have strict, harmonised rules for block boundary placement. Instead, operational needs, such as routing flexibility and capacity considerations, determine platform segmentation, while respecting general safety and sighting constraints (Pachl, 2021; Vignali et al., 2020; Wagter, 2019). Under ETCS L2 HTD, VSS boundaries can be placed within platform areas to support operational objectives such as short headways, turn back operations, or limiting the propagation of virtual occupancy states. However, as in conventional practice, they must be placed at locations where trains may safely stop and where they do not interfere with sighting, interlocking release, or platform specific operational rules. The flexibility of VSS segmentation allows for a more capacity-optimised design, but the underlying infrastructure constraints remain unchanged.

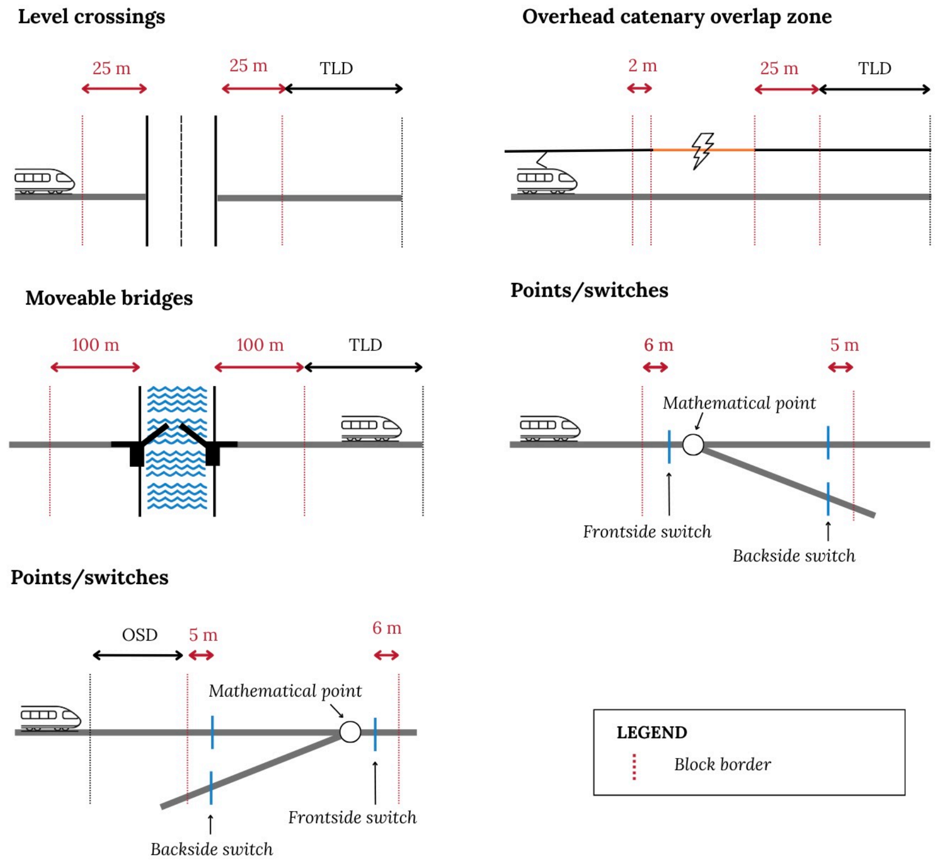

Figure 4.3: Overview of the block boundary placement restrictions.

4.3. Implications for TTD layout optimisation

The requirements and restrictions identified in this chapter together define the constrained design domain within which the optimal TTD layout must be determined. The safety-critical and operationally critical aspects impose fixed TTD locations that cannot be removed or relocated, while the general recommendations identify additional locations where physical detection is beneficial but not strictly required. The release section restrictions determine the feasible VSS grid and possible physical TTD positions, while the block boundary restrictions determine which of those boundaries may be configured as operational MA limits. Within this constrained design domain, the remaining degrees of freedom concern the number and placement of TTD boundaries between the constrained mandatory locations, including station approaches and open-line parts of the corridor. Chapter 5 develops a methodology to optimise this layout, balancing infrastructure cost against operational capacity under mixed-traffic conditions.

5 Optimisation framework for TTD layout design under ETCS HTD

Chapter 4 classified the design requirements that constrain the ETCS HTD layout and translated them into mandatory, forbidden, and optional boundary positions. Within this constrained design domain, the remaining problem is to determine which optional TTD boundaries should be retained, removed, or added during the staged migration to ETCS HTD, over which an increasing share of the fleet becomes TIM-equipped. This chapter develops the corresponding optimisation framework. It derives the analytical headway and cost objectives, constructs the feasible design domain, formulates the staged layout and deployment decisions as a bi-objective MIP, and defines how the resulting alternatives are selected and evaluated. Figure 5.1 summarises the complete procedure.

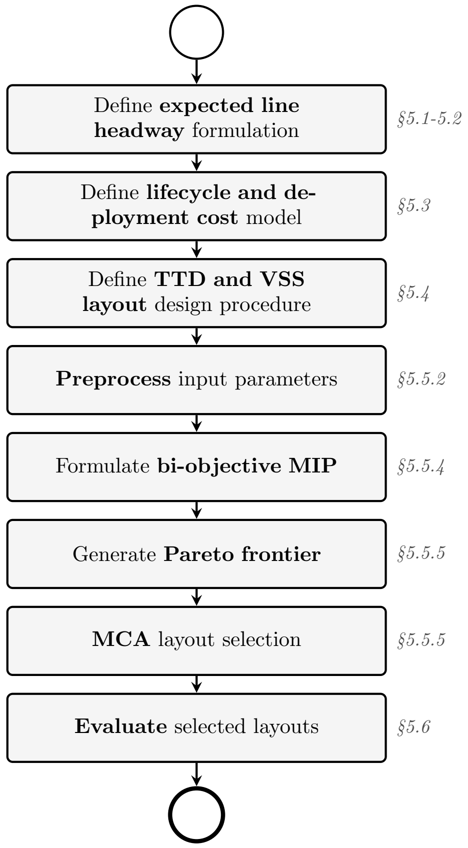

Figure 5.1: Overview of the design framework for optimal TTD boundary placement.

The methodological architecture follows directly from the structure of the migration problem. Optimising each infrastructure stage independently would neglect transition costs and could select a locally attractive layout that requires unnecessary placement or removal in later stages. Prescribing the TIM trajectory beforehand would similarly remove the economic interaction between onboard deployment and residual TTD adaptation. Conversely, relying exclusively on microscopic simulation would provide detailed operational evaluation but would not offer a tractable way to search the large combinatorial space of staged TTD layouts. The framework therefore combines joint staged optimisation with an analytical capacity surrogate and uses microscopic simulation independently to evaluate the operational consequences of the selected alternatives.

5.1. Analytical headway formulation under ETCS L2

This section develops the analytical headway expressions that the optimisation model uses. It starts from blocking time theory for ETCS L2 fixed block operation and then derives the line headway principle that the remainder of the chapter uses. The blocking time principle extends to mixed-traffic corridors by evaluating the timing components for each relevant leader-follower scenario.

5.1.1. Blocking time and minimum line headway

The minimum headway between two following trains depends on the blocking time of the block sections along the line (I. A. Hansen & Pachl, 2014; Pachl, 2021). Under ETCS L2 fixed block operation, the blocking time t_{blocking} of a block section is the time interval during which that block section is reserved exclusively for one train and is therefore blocked to all other trains (Pachl, 2021). This blocking time is usually considerably longer than the time the train physically occupies the block section: the physical occupation is only the running time component t_{run}, whereas the blocking time also includes the route setting, reacting, approach, clearing, and release times around it. As Figure 5.2 illustrates, the blocking time of a block section decomposes into six components:

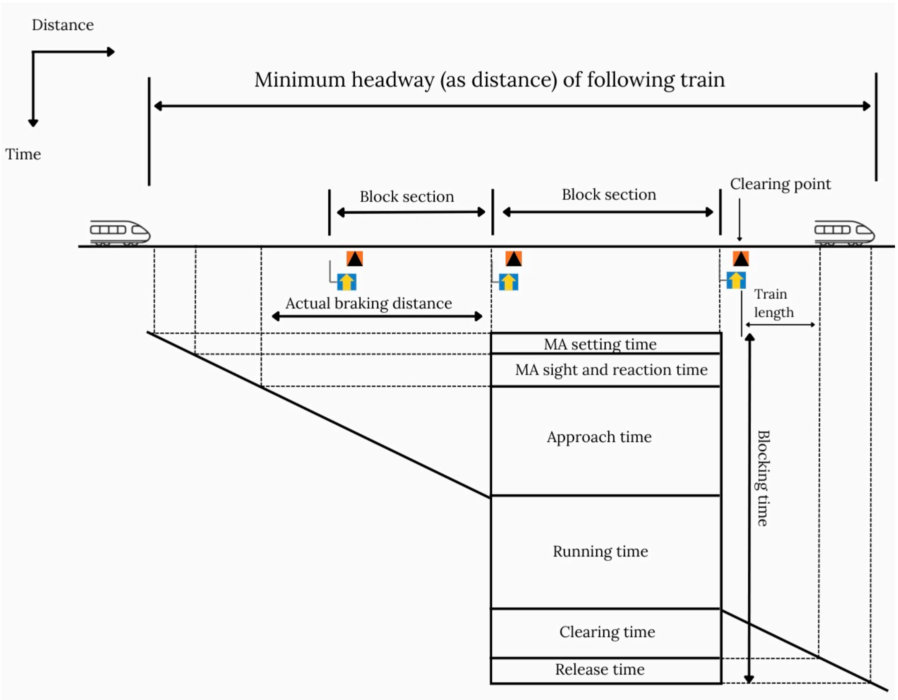

Figure 5.2: Blocking time of a block section under ETCS L2 fixed block operation, based on Pachl (2021).

$$t_{\text{blocking}} = t_{\text{setting}} + t_{\text{sr}} + t_{\text{approach}} + t_{\text{run}} + t_{\text{clear}} + t_{\text{release}}. \tag{5.1}$$

Following the standard blocking time formulation for ETCS L2 with fixed block operation (I. A. Hansen & Pachl, 2014; Pachl, 2021), the six blocking time components are defined as:

- t_{setting}: the route setup time. Besides the processing time required by the RBC to issue a renewed movement authority, it includes the time to move, set, and lock the switches along the route.
- t_{sr}: the time that the driver needs to see and react to the movement authority once it has been notified and acknowledged.
- t_{approach}: the time margin associated with the braking time to the block section entry. In the remainder of the formulation, w_p denotes this component, where p indexes a possible TTD boundary position on the VSS grid and a block section runs from an entry x_p to an exit x_q.

- t_{run}: the running time through the relevant block section.
- t_{clear}: the time for the rear of the train to pass the clearing point, the TTD boundary at which the trackside train detection system registers the block section as vacated.
- t_{release}: the release time to release the block section.

With this decomposition, the blocking times of a train can be plotted along the corridor as a blocking time stairway (Pachl, 2021). The minimum line headway is obtained by shifting the stairway of the following train towards that of the leader until the two just touch without overlapping. The block section at which this limiting interaction occurs is the critical block section, as illustrated in Figure 5.3.

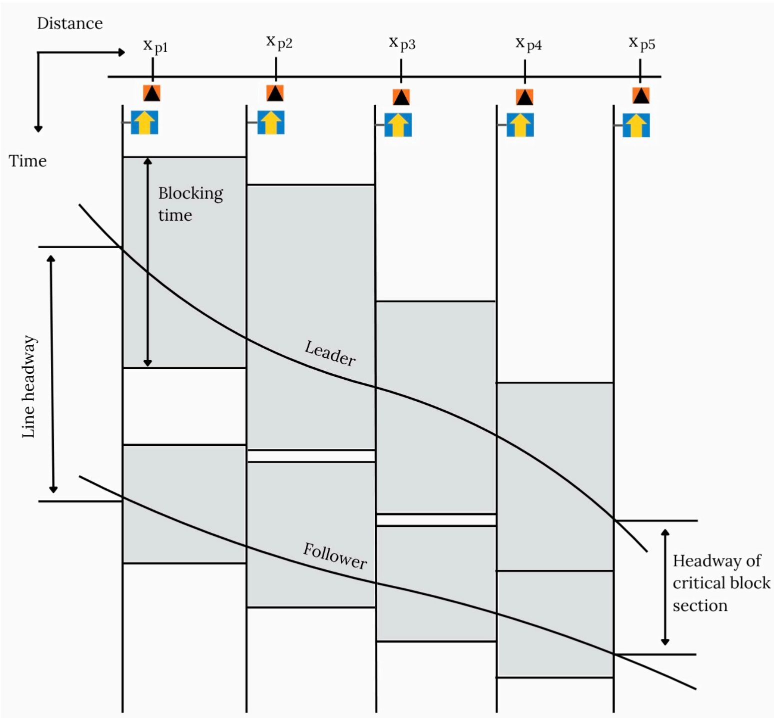

Figure 5.3: Blocking time stairways of a leading and a following train, showing the critical block section that determines the minimum line headway, based on Pachl (2021).

The minimum line headway is therefore the largest section-specific headway along the corridor:

$$H_{\text{line}} = \max_{\text{block sections } (p,q)} h_{pq}, \tag{5.2}$$

where h_{pq} is the minimum headway at block section (p, q). Consequently, capacity depends not only on the number of block sections but also on the locations of their boundaries, because only a binding or near-binding block section can determine the line headway.

For homogeneous traffic, a consecutive leader-follower pair has identical running time profiles and therefore identical blocking time stairways. For heterogeneous traffic, differences in running and stopping patterns shift the stairways relative to each other along the corridor. To capture this interaction, the model extends the blocking time components of Equation (5.1) with a cumulative running time difference, developed in the next subsection.

5.1.2. Minimum headway components

The minimum headway h_{pq} at a block section $[x_p, x_q]$ is obtained by splitting the six blocking time components between the two trains (leader and follower), assigning each component to the train whose movement governs it, following Figure 5.2. The following train contributes the components: the setting and reaction times, grouped as the movement authority time $t_{\text{MA,follower}} = t_{\text{setting}} + t_{\text{sr}}$, and the approach time $w_{p,\text{follower}}$ at the block section entry x_p. The leading train contributes the components: the running time through the block section, the clearing time $c_{q,\text{leader}}$ at the block section exit x_q, and the release time $t_{\text{release,leader}}$.

Homogeneous leader-follower pairs share the same running time profile, so these local components alone determine the minimum headway. For heterogeneous pairs, however, the two trains may already have accumulated a running time difference upstream of the block section entry. This offset shifts the two blocking time stairways relative to each other and must be added as Δ_p. For a block section $[x_p, x_q]$, the minimum headway of a leader-follower pair is therefore

$$h_{pq} = \Delta_p + t_{\text{MA,follower}} + w_{p,\text{follower}} + \big(t_{\text{leader}}(x_q) - t_{\text{leader}}(x_p)\big) + c_{q,\text{leader}} + t_{\text{release,leader}}, \tag{5.3}$$

where $t_{\text{leader}}(x)$ and $t_{\text{follower}}(x)$ denote the cumulative running times of the leader and follower measured from a common reference point at the corridor entry, $\big(t_{\text{leader}}(x_q) - t_{\text{leader}}(x_p)\big)$ is the running time of the leading train through the block section, and Δ_p is the cumulative running time difference at the entry:

$$\Delta_p = t_{\text{leader}}(x_p) - t_{\text{follower}}(x_p). \tag{5.4}$$

This offset is zero for homogeneous pairs, so the minimum headway reduces to the local blocking time components, and non-zero for heterogeneous pairs. A positive value means the leader has accumulated running time relative to the follower before reaching x_p, which increases the required separation, while a negative value means the leader runs ahead, which reduces it.

The detailed computation of five of the six components, namely the three constant components (the setting, sight and reaction, and release times), the approach time from the follower's braking curve, and the clearing time from the leader's length, is given in Appendix C.1. A one at a time sensitivity analysis of the headway components, reported in Appendix C.2, confirms that the running time dominates the headway. The simplified computation of the remaining components, in particular the constant-deceleration approximation of the approach time, therefore has limited influence on which TTD layout the optimisation selects. The cumulative running time difference Δ_p is defined above. The sixth component, the running time t_{run}, takes a different form under ETCS HTD than under ETCS L2 and is therefore developed in Section 5.2.

5.2. Analytical headway formulation under ETCS HTD

This section adapts the blocking time and line headway framework of Section 5.1 to ETCS HTD, following the release and operational block hierarchy established in Chapter 4. When the leading train reports confirmed integrity, infrastructure behind it can be released at VSS level and the following train can make use of the applicable virtual block boundaries. Without confirmed integrity, release remains dependent on the physical TTD section and the following train remains governed by the applicable ETCS L2 fixed block boundaries. The remainder of this section defines the two release regimes, introduces the expected line headway weighted by the TIM ratio, and derives the corresponding minimum headway expressions.

5.2.1. Release regimes and section level approximation

The blocking time t_{blocking} of Equation (5.1) takes a different form for the two release regimes, because the running time component t_{run} differs depending on the integrity status of the leading train. These regimes distinguish fast from slow track-releasing pairs, where fast and slow refer to the release mechanism, not to the train type or running speed (EEIG, 2026):

- **Fast track-releasing pairs** are pairs in which the leading train reports confirmed integrity. Infrastructure behind the leader can then be released at VSS level, while the movement authority of the follower remains limited by the applicable operational virtual block boundaries.
- **Slow track-releasing pairs** are pairs in which the leading train cannot provide confirmed integrity. Infrastructure behind the leader can then be released only at TTD-section level, while the movement authority of the follower remains limited by the applicable ETCS L2 fixed block boundaries.

By the standard blocking time definition, line headway is evaluated over operational block sections (I. A. Hansen & Pachl, 2014; Pachl, 2021). This study instead uses the relevant release section as an analytical capacity surrogate, because the optimisation changes the physical TTD layout rather than the operational block layout. Equation (5.3) is therefore evaluated over one VSS for a fast track-releasing pair and over the active TTD section for a slow track-releasing pair.

The movement authority and release times remain unchanged. The cumulative running time difference and follower approach time are evaluated at the release section entry, the leader running time through the release section, and the clearing time at its exit. Shortening a release section therefore reduces the leader running time through it, while the boundary-dependent terms may also vary with boundary location. This formulation links the selected TTD boundaries directly to the analytical capacity objective.

Let $\mathcal{P}$ denote the set of consecutive VSSs and $\mathcal{S}$ the set of active TTD sections. Because every TTD boundary coincides with a VSS boundary, a VSS can be indexed by its left boundary p, whereas a TTD section is indexed by both endpoints (p, q), with $q \geq p + 1$.

Applying Equation (5.3) to these two release sections gives the per-section minimum headways $h_{\text{VSS},p}$ and $h_{\text{TTD},pq}$ for fast and slow track-releasing pairs, respectively.

Applying the line headway of Equation (5.2) to each regime yields the fast and slow track-releasing line headways:

$$F = \max_{p \in \mathcal{P}} h_{\text{VSS},p}, \tag{5.5}$$

$$G = \max_{(p,q) \in \mathcal{S}} h_{\text{TTD},pq}. \tag{5.6}$$

Here F is the layout-independent fast-release line headway determined by the fixed VSS grid, whereas G is the maximum over the active TTD sections and is the layout-dependent component.

Because the line headway is determined by the maximum section headway, adding a TTD boundary does not necessarily reduce G. A boundary improves the analytical line headway only when it subdivides a TTD section whose headway is binding, or sufficiently close to binding, in this maximum. Subdividing a non-binding section leaves the line headway unchanged.

During migration, the TIM ratio ρ is the proportion of leader-follower pairs that operate as fast track-releasing pairs, with the remaining $1 - \rho$ operating as slow pairs. Equivalently, ρ represents the share of relevant leading train events for which confirmed train integrity is available. When it is interpreted as a fleet deployment share, the model assumes that this capability is distributed proportionally across the leader service types in the traffic mix, so that a single ratio applies to every scenario. Assuming the release status of a leader remains unchanged over its corridor traversal, the expected line headway weighted by the TIM ratio is

$$H^{\rho}_{\text{line}} = \rho F + (1 - \rho) G. \tag{5.7}$$

This deterministic weighting captures the first-order effect of partial TIM deployment while avoiding a full stochastic model of train by train availability.

The formulation above applies to one leader-follower scenario. On a mixed-traffic corridor, the headway is evaluated separately for several representative scenarios, each defined by a pair of consecutive train services with its own speed profile, stopping pattern, formation length, and cumulative running time difference Δ_p. A homogeneous corridor reduces to a single scenario, for which $\Delta_p = 0$. The scenario-specific expected line headways are weighted by their operational frequency so that the aggregated headway reflects the actual traffic composition rather than assigning equal importance to frequently and rarely occurring train pairs. This aggregation is formalised in Section 5.5.4.

The key approximation in this analytical formulation is that headway is evaluated over release sections rather than exclusively over operational block sections. Operationally, however, a release boundary is not necessarily a movement authority boundary. As Chapter 4 shows, some mandatory release boundaries at bridges, level crossings, and catenary overlaps cannot be used to terminate a movement authority. A following train therefore cannot always exploit a shorter release section to the same extent as assumed by the analytical formulation. The analytical line headway should consequently be interpreted as an optimistic infrastructure-based surrogate for comparing TTD layouts. The RailSys consistency checks use the actual operational block configuration to assess the consequence of this approximation, with the results reported in Chapter 7.

The analytical line headway derived above forms the first objective of the bi-objective optimisation formulated in Section 5.5.4. The following section develops the second and final objective. Operational stability, the third performance criterion, is not included in the objective function and is instead evaluated separately in RailSys.

5.3. Lifecycle and deployment cost model

This section defines the second objective of the bi-objective optimisation: the cost associated with the TTD layout and equipping rolling stock with TIM. The cost objective function combines infrastructure costs with a rolling stock retrofit component. Capacity and cost are the two criteria that the R2DATO D37.1 deployment strategy framework identifies as central to ETCS HTD deployment decisions (Aoun, 2026a), which motivates their use as the two competing objectives of this study.

The model includes only cost components that differ between TTD layout and deployment alternatives and therefore affect the optimisation. It omits costs that are required for a full ETCS HTD implementation but are not differentiated between the alternatives considered in this study, such as the RBC, generic ETCS trackside equipment, and the baseline VSS configuration. These fixed background costs would shift the total rollout cost by an approximately constant amount, so the comparative cost-minimisation objective excludes them.

These cost components are incurred over the course of the migration, and the framework distinguishes between infrastructure stages and migration years because the two deployment processes occur at different temporal scales. Physical changes to the TTD layout require trackside interventions and are therefore assumed to occur only at discrete migration stages. Equipping rolling stock with TIM can progress more gradually and is therefore represented annually. A stage consequently spans one or more years during which the TTD layout remains fixed, while the TIM ratio may change from year to year.

Because these costs fall at different points across the migration horizon, they cannot be compared directly on an undiscounted basis. The model therefore converts all costs to a common present-value basis using a social discount rate, so that alternative deployment strategies with different cost timings can be compared consistently.

The total discounted cost, formalised in Section 5.5.4, comprises four components:

- **Transition cost.** Placing a new TTD boundary, represented by an axle counter, and removing an existing one each incur a one-off cost. The model charges these transition costs once per stage transition, discounted to the first year of the stage, whenever the layout changes between two consecutive stages. It treats the unit costs for axle counter placement and removal as assumed planning parameters, and the case study specifies them.
- **Horizon maintenance.** Every active TTD boundary incurs a recurring annual maintenance cost. Summed over the migration years and discounted per year, this captures the upkeep of the boundaries in service during the planning horizon. The case study specifies the annual maintenance unit cost.
- **Terminal maintenance.** The layout in service at the end of the migration horizon continues to require maintenance over its remaining asset lifetime. A discounted terminal maintenance term accounts for this residual cost so that the optimiser does not favour boundaries merely because they are introduced late in the evaluated horizon.
- **Rolling stock retrofit cost.** Increasing the TIM ratio requires additional rolling stock capable of providing the integrity confirmation required for VSS level release. The model represents this through an approximate corridor-attributed retrofit cost per train. Actual retrofit costs depend on rolling stock type and onboard architecture, so this value is treated as a planning approximation. The purchase cost of replacement rolling stock is excluded because fleet replacement is a broader investment decision rather than an incremental TIM retrofit cost.

The cost objective function supports two planning scopes that share the same mathematical structure and differ only in interpretation. In the ETCS HTD migration scope, the objective combines the infrastructure lifecycle cost with the approximate corridor-attributed TIM retrofit cost, representing a system level trade-off between capacity improvement and the combined infrastructure and rolling stock investment needed to realise it. In the narrower infrastructure scope, the model sets the TIM retrofit cost component to zero, so the TIM ratio is still chosen by the optimiser but carries no retrofit cost and enters only through its effect on the expected line headway.

The notation table (Section 5.5.3) lists the cost parameters, the discount rate, and the asset lifetime, and Chapter 6 specifies them for the case corridor.

With both the capacity and the cost of a layout now quantified, the next section constructs the design domain, namely the VSS grid and the possible TTD boundary positions over which the optimisation chooses.

5.4. Design procedure for VSS grid construction and TTD boundary placement

Before introducing the optimisation model, this section defines the feasible VSS and TTD design domain. Mandatory and forbidden positions follow from Chapter 4. The VSS grid is then constructed first, after which the optimisation selects physical TTD boundaries only from the resulting VSS boundary positions. Figure 5.4 summarises the four-step procedure.

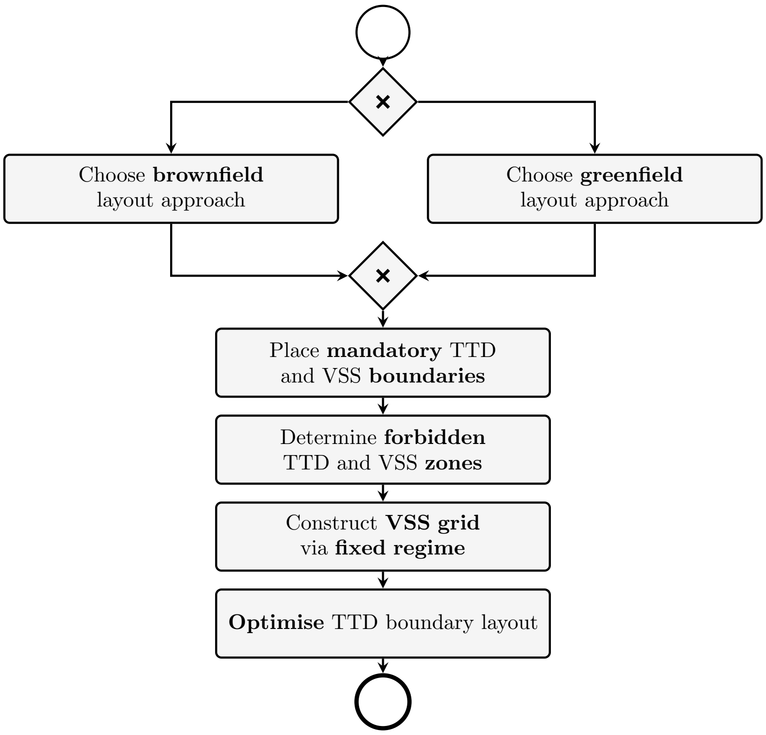

Figure 5.4: Design procedure for TTD and VSS layout construction.

Step 1: Place mandatory TTD and VSS boundaries

Infrastructure and operational requirements impose positions where both a TTD and a VSS boundary must be present. These mandatory positions form the set $\mathcal{M}$, remain fixed at every migration year, and follow the safety critical classification established in Chapter 4. In this study, the corridor start and end points are chosen to coincide with mandatory positions from $\mathcal{M}$ near the respective station nodes. These fixed endpoints provide boundary conditions for the optimisation model, ensuring that all TTD sections and VSSs that enter the headway and cost calculations lie fully within the study corridor.

Step 2: Determine forbidden zones

Certain corridor segments are unsuitable for boundary placement, as Chapter 4 identifies. These positions form the forbidden set $\mathcal{F}$, with $\mathcal{F} \cap \mathcal{M} = \varnothing$, and neither TTD nor VSS boundaries may lie within them. The mandatory overrules forbidden rule excludes positions from $\mathcal{F}$ that geometrically fall inside an extended forbidden zone but carry a mandatory classification. Together with $\mathcal{M}$, the forbidden set defines the feasible design domain for the optimisation.

Step 3: Construct the VSS grid in the remaining corridor space

The procedure fills the corridor segments between mandatory and forbidden positions with VSS boundaries at a fixed interval d_{VSS}. This study adopts the fixed short block regime as defined in the R2DATO D37.2 framework (Aoun, 2025). This regime is characterised by short VSSs of fixed length and is appropriate when capacity improvement and fast recovery after disruptions are the primary design objectives (Aoun, 2025). A smaller d_{VSS} yields a finer VSS boundary layout and therefore a larger design domain for the TTD boundary optimisation, which can improve the best achievable headway. However, a finer grid also increases the number of binary decision variables and therefore the computational cost of the optimisation model. The choice of d_{VSS} within the range defined by the D37.2 framework therefore represents a trade-off between design domain granularity and computational tractability. Chapter 6 reports the specific value this study uses.

Where the distance between two consecutive breakpoints is not an exact multiple of d_{VSS}, the procedure divides the available span into equally sized sub-sections. Specifically, for a free span of width w between two breakpoints, the procedure divides the span into $n = \max\bigl(1, \operatorname{round}(w/d_{\text{VSS}})\bigr)$ equal parts and rounds the resulting interior boundary coordinates to integer metres, while the two span endpoints remain fixed. Rounding the boundary coordinates rather than the individual section lengths keeps every boundary within the span and prevents a rounding residual from accumulating across the subdivision, so that all VSSs within a span are approximately d_{VSS}. The resulting VSS grid defines the possible TTD boundary positions for placement in Step 4.

Step 4: Optimise the TTD boundary layout

Given the VSS grid from Step 3 and the mandatory and forbidden sets from Steps 1-2, the optimisation selects which feasible VSS boundaries should also function as physical TTD boundaries. Every selected physical boundary therefore lies on the predefined VSS grid (Aoun, 2025; EEIG, 2026).

The framework does not require the VSS grid itself to be generated by the fixed short block regime used here. Once a feasible VSS grid is fixed, the optimisation determines the number and placement of physical TTD boundaries within that release structure. The next section formalises this boundary selection problem as a MIP.

5.5. Optimisation model for TTD boundary placement

The optimisation problem combines two decisions that develop over the same migration horizon but at different time scales. As illustrated in Figure 5.5, the TIM deployment level can change from year to year, while the physical TTD layout changes only at discrete migration stages. Throughout the migration, the VSS grid remains fixed. The model therefore determines which positions on this grid function as physical TTD boundaries in each stage while simultaneously determining the progression of TIM deployment.

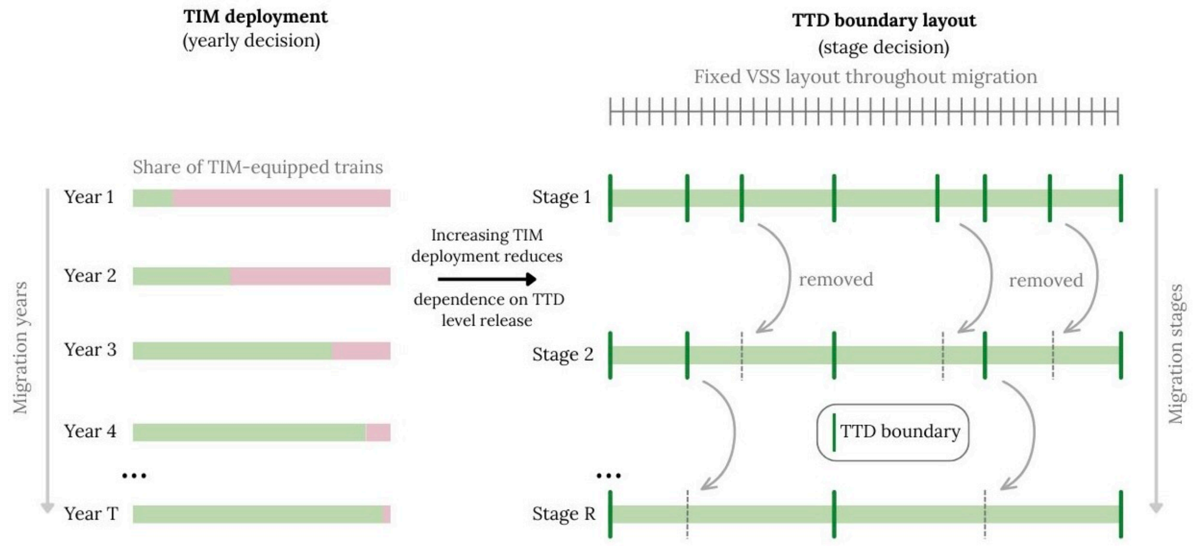

Figure 5.5: Illustrative example of the two coupled decisions the optimisation model determines jointly: the yearly TIM deployment trajectory (left) and the staged physical TTD boundary layout on the fixed VSS grid (right).

The two decisions are coupled because increasing TIM deployment allows a larger share of relevant train movements to use VSS level release, reducing their dependence on physical TTD level release. The nominal capacity value of additional physical boundaries can therefore change as deployment progresses, even though keeping them in service still incurs infrastructure cost. Figure 5.5 illustrates boundary removal between stages, although the model can retain, remove, or add optional boundaries wherever permitted by the design constraints.

The available layout choices also depend on the migration starting condition. In a brownfield application, the optimisation starts from the inherited physical TTD layout and determines which existing boundaries should be retained or removed and where additional boundaries may be introduced. In a greenfield application, no inherited physical layout constrains the design, so the TTD boundary locations are selected from scratch within the feasible design domain. The initial layout vector $\hat{y}$ represents this distinction mathematically.

This coupled and time-dependent boundary selection problem is formulated as a bi-objective MIP. The expected line headway and discounted cost formulations developed in the preceding sections link the layout and deployment decisions to the two objectives. Minimising headway alone would provide no incentive to remove boundaries that no longer improve the binding headway, whereas minimising cost alone could remove boundaries that still contribute materially to capacity. The formulation therefore minimises the maximum expected line headway and the total discounted lifecycle and deployment cost over the migration horizon.

The migration stages are therefore optimised jointly, so that layout performance in each stage is evaluated together with the cost of transitioning between successive layouts.

5.5.1. Modelling scope and assumptions

The optimisation is interpreted under the following modelling assumptions:

- The VSS grid is fixed before optimisation, and the model selects the physical TTD boundaries from the feasible positions on this grid.
- Train running characteristics are precomputed and treated as independent of the selected physical TTD layout. Analytical headway is evaluated over effective release sections as the capacity surrogate developed in Section 5.2, without timetable allowances.
- Partial TIM deployment is represented by a deterministic, non-decreasing ratio selected from a discrete set in each migration year. When interpreted as fleet deployment, this ratio is assumed to

apply proportionally across the relevant leader service types, and the integrity status of a leader remains unchanged during its corridor traversal.

- The cost objective includes only cost components that differ between the alternatives considered.
- Operational stability is outside the optimisation objective and constraints and is evaluated separately in RailSys.

5.5.2. Precomputed parameters

The analytical headway calculations are performed separately for each representative ordered leader-follower scenario, because the train combinations differ in running and stopping characteristics. These scenarios form the set $\mathcal{K}$, indexed by k, and the optimisation is built over the scenarios with positive traffic weight, $\mathcal{K}^+ = \{k \in \mathcal{K} : \pi_k > 0\}$, because zero-weight scenarios affect neither the objective nor the selected layout. All timing quantities that are independent of the selected TTD layout are calculated before optimisation.

Because the VSS grid is predefined, the selected physical layout affects only which candidate slow track-releasing sections become active and therefore the resulting slow track-releasing line headway G_k^r. The resulting precomputed quantities are summarised in Table 5.1.

Table 5.1: Parameters precomputed for each representative leader-follower scenario before optimisation.

Parameter	Meaning	Unit
$w_{k,p}$	Follower approach time at boundary position x_p for scenario k.	[s]
$\delta_{k,p}$	Leader running time through VSS $[x_p, x_{p+1}]$ for scenario k.	[s]
$c_{k,p}$	Leader clearing time at x_p for scenario k.	[s]
$\Delta_{k,p}$	Cumulative leader-follower running time offset at x_p for scenario k.	[s]
$h_{\text{VSS},k,p}$	Fast track-releasing section headway for VSS p and scenario k.	[s]
F_k	Fast track-releasing line headway for scenario k.	[s]
$h_{\text{TTD},k,ij}$	Slow track-releasing section headway for feasible candidate TTD section $[x_i, x_j]$, $(i,j) \in A_q$, and scenario k.	[s]
L_k	Certified lower bound on the slow track-releasing line headway for scenario k: the smallest maximum section headway attainable along any feasible sequence of TTD sections between consecutive mandatory boundaries.	[s]

5.5.3. Notation

Table 5.2 summarises the symbols used in the MIP formulation. The precomputed timing parameters $(w, \delta, c, \Delta, h_{\text{VSS}}, F, h_{\text{TTD}})$ enter as constants and are documented in Table 5.1 of Section 5.5.2.

Table 5.2: Notation for the TTD boundary optimisation model.

Symbol	Description	Unit
Sets and indices		
$\mathcal{B}$	Possible TTD boundary positions on the VSS grid, indexed $p = 0, \dots, B$	
$\mathcal{P}$	VSSs $[x_p, x_{p+1}]$ between consecutive positions in $\mathcal{B}$, indexed $p \in \{0, \dots, B-1\}$	
$\mathcal{M} \subseteq \mathcal{B}$	Mandatory TTD boundary positions, with $\{0, B\} \subseteq \mathcal{M}$	
$\mathcal{F} \subseteq \mathcal{B}$	Forbidden TTD boundary positions, with $\mathcal{F} \cap \mathcal{M} = \varnothing$	
$\mathcal{T} = \{0, 1, \dots, T-1\}$	Migration years in the migration horizon, indexed t	
$\mathcal{R} = \{0, 1, \dots, R-1\}$	Migration stages, each covering a contiguous block of years, indexed r	
$\mathcal{K}$	Set of representative ordered leader-follower scenarios, indexed k	
$\mathcal{K}^+ = \{k \in \mathcal{K} : \pi_k > 0\}$	Scenarios with positive traffic weight, over which the optimisation is built	
$L_t \subseteq [0, 1]$	Set of admissible TIM deployment levels in year t (each a ratio in $[0, 1]$), indexed ℓ	
Q	Mandatory intervals, one per pair of consecutive mandatory boundaries, indexed q with endpoints $m_q < n_q$	
$V_q \subseteq \mathcal{B}$	Optional boundary positions strictly inside interval q (forbidden positions excluded)	
$A_q = \{(i, j) : i, j \in \{m_q, n_q\} \cup V_q,\ i < j\}$	Candidate TTD sections (arcs) within interval q, indexed (i, j)	
$r(t) \colon \mathcal{T} \to \mathcal{R}$	Year to stage mapping (non-decreasing, surjective)	
Model parameters		
$\pi_k \in [0, 1]$	Scenario weight (expected share in traffic mix), with $\sum_k \pi_k = 1$	[-]
$\hat{y}_p \in \{0, 1\}$	Initial boundary state before stage 0 (brownfield or greenfield)	[-]
c_{place}	Cost of placing a new TTD boundary	[€]
c_{remove}	Cost of removing an existing TTD boundary	[€]
c_{maint}	Maintenance cost per active TTD boundary per year	[€/yr]
c_{train}	Per-train TIM retrofit cost	[€]
N_{train}	Number of trainsets circulating on the corridor	[-]
d	Social discount rate	[-]
$\alpha_t = (1 + d)^{-t}$	Discount factor for year t	[-]
β_{tail}	Discounted annuity factor for terminal maintenance	[-]
T_{asset}	Expected asset lifetime of a TTD boundary	[yr]
t_r^*	First year of stage r	[yr]
$C_{\max}$	Cost level (ε-constraint upper bound on f_2)	[€]
Decision variables		
$y_p^r \in \{0, 1\}$	1 if a TTD boundary is active at position p at stage r	[-]
$x_{ij}^r \in \{0, 1\}$	1 if section (arc) (i, j) is an active TTD section at stage r	[-]
$b_\ell^t \in \{0, 1\}$	1 if TIM deployment level $\ell \in L_t$ is selected at year t	[-]
$G_k^r \geq 0$	Slow track-releasing line headway for scenario k under the physical TTD layout of stage r	[s]
$H^t \geq 0$	Expected line headway at year t	[s]
$\bar{H} \geq 0$	Maximum expected line headway across all years	[s]
Derived expression		
$\rho_t = \sum_{\ell \in L_t} \ell \, b_\ell^t$	Selected TIM ratio in year t	[-]
Initial conditions		
$y_p^{-1} = \hat{y}_p$	Layout before the migration horizon	[-]
$\rho_{-1} = 0$	TIM ratio before the migration horizon	[-]

5.5.4. MIP formulation

The optimisation has two objectives: the maximum expected line headway $\bar{H}$ and the total discounted lifecycle and deployment cost f_2. The maximum rather than the average expected line headway over the migration horizon is minimised because each migration stage must provide an acceptable level of capacity in its own right. An average objective could allow poor capacity performance in an intermediate stage to be compensated by lower headways in later years with higher TIM deployment. The minimax objective therefore limits the worst capacity condition encountered during the migration.

The bi-objective optimisation problem takes the following form:

$$\min\ (f_1,\ f_2) \tag{5.8}$$

where:

$$f_1 = \bar{H} \tag{5.9}$$

For compactness, the selected TIM ratio is written as

$$\rho_t = \sum_{\ell \in L_t} \ell\, b_\ell^t \quad \forall\, t \in \mathcal{T}, \qquad \rho_{-1} = 0.$$

This is a derived expression rather than an additional decision variable.

The initial conditions $y_p^{-1} = \hat{y}_p$ and $\rho_{-1} = 0$ define the layout and TIM deployment before the migration horizon. They ensure that first-stage placement and removal costs and the first-year retrofit increment are defined relative to the initial state. These are fixed reference values rather than additional decision variables.

$$\begin{aligned} f_2 = & \underbrace{\sum_{r \in \mathcal{R}} \alpha_{t_r^*} \sum_{p \in \mathcal{B}} \Big(c_{\text{place}}\, [y_p^r - y_p^{r-1}]^+ + c_{\text{remove}}\, [y_p^{r-1} - y_p^r]^+ \Big)}_{\text{transition cost}} + \underbrace{\sum_{t \in \mathcal{T}} \alpha_t \sum_{p \in \mathcal{B}} c_{\text{maint}}\, y_p^{r(t)}}_{\text{yearly maintenance}} \\ & + \underbrace{\beta_{\text{tail}}\, c_{\text{maint}} \sum_{p \in \mathcal{B}} y_p^{R-1}}_{\text{terminal maintenance}} + \underbrace{\sum_{t \in \mathcal{T}} \alpha_t\, (\rho_t - \rho_{t-1})\, N_{\text{train}}\, c_{\text{train}}}_{\text{TIM retrofit cost}} . \end{aligned} \tag{5.10}$$

The boundary variables y_p^r indicate which physical TTD boundaries are active, but they do not directly identify which active boundaries are consecutive and therefore form a TTD section. The arc variables x_{ij}^r provide this second piece of information through a *section-arc path formulation*. Within each mandatory interval $q \in Q$, bounded by consecutive mandatory boundaries $m_q < n_q$, a feasible layout is a sequence of active TTD sections that starts at m_q, visits the selected optional boundary positions in geographic order, and ends at n_q. The binary selector x_{ij}^r marks each candidate section $(i, j) \in A_q$ as active or inactive at stage r. Because every TTD boundary coincides with a VSS boundary and no active section may span a mandatory boundary, selecting a layout is equivalent to selecting one directed path from m_q to n_q in each interval. Since the section headway $h_{\text{TTD},k,ij}$ is a precomputed constant, bounding G_k^r by the headways of the active sections is linear and needs no big-M switching.

The same graph yields a layout-independent lower bound on G_k^r. Because every feasible layout is a path from m_q to n_q, the smallest maximum section headway attainable across all such paths bounds G_k^r from below in every stage:

$$L_k = \max_{q \in Q}\ \min_{P \in \mathcal{P}_q}\ \max \left\{ 0,\ \max_{(i,j) \in P} h_{\text{TTD},k,ij} \right\}, \tag{5.11}$$

where $\mathcal{P}_q$ denotes the feasible paths from m_q to n_q within interval q. The bound L_k is precomputed once per scenario by a bottleneck-path dynamic program over the interval graph, detailed in Appendix B.

subject to:

$$\bar{H} \geq H^t \qquad \forall\, t \in \mathcal{T} \tag{C1}$$

$$H^t \geq \sum_{k \in \mathcal{K}^+} \pi_k \left[\rho_t F_k + (1-\rho_t) G_k^{r(t)}\right] \qquad \forall\, t \in \mathcal{T} \tag{C2}$$

$$G_k^r \geq h_{\text{TTD},k,ij}\, x_{ij}^r \qquad \forall\, k \in \mathcal{K}^+,\ r \in \mathcal{R},\ (i,j) \in A_q,\ q \in Q \tag{C3}$$

$$G_k^r \geq L_k \qquad \forall\, k \in \mathcal{K}^+,\ r \in \mathcal{R} \tag{C4}$$

$$\sum_{j:\,(m_q,j)\in A_q} x_{m_q j}^r = 1 \qquad \forall\, q \in Q,\ r \in \mathcal{R} \tag{C5a}$$

$$\sum_{i:\,(i,n_q)\in A_q} x_{i\,n_q}^r = 1 \qquad \forall\, q \in Q,\ r \in \mathcal{R} \tag{C5b}$$

$$\sum_{i:\,(i,p)\in A_q} x_{ip}^r = \sum_{j:\,(p,j)\in A_q} x_{pj}^r \qquad \forall\, p \in V_q,\ q \in Q,\ r \in \mathcal{R} \tag{C5c}$$

$$y_p^r = \sum_{i:\,(i,p)\in A_q} x_{ip}^r \qquad \forall\, p \in V_q,\ q \in Q,\ r \in \mathcal{R} \tag{C5d}$$

$$y_p^r = 1 \qquad \forall\, p \in \mathcal{M},\ r \in \mathcal{R} \tag{C6a}$$

$$y_p^r = 0 \qquad \forall\, p \in \mathcal{F},\ r \in \mathcal{R} \tag{C6b}$$

$$f_2 \leq C_{\max} \tag{C7}$$

$$\sum_{\ell \in L_t} b_\ell^t = 1 \qquad \forall\, t \in \mathcal{T} \tag{C8a}$$

$$\rho_t \geq \rho_{t-1} \qquad \forall\, t \in \mathcal{T} \setminus \{0\} \tag{C8b}$$

$$\begin{aligned}
&y_p^r \in \{0,1\} && \forall\, p \in \mathcal{B},\ r \in \mathcal{R},\\
&x_{ij}^r \in \{0,1\} && \forall\, (i,j) \in A_q,\ q \in Q,\ r \in \mathcal{R},\\
&b_\ell^t \in \{0,1\} && \forall\, \ell \in L_t,\ t \in \mathcal{T},\\
&G_k^r \geq 0 && \forall\, k \in \mathcal{K}^+,\ r \in \mathcal{R},\\
&H^t \geq 0 && \forall\, t \in \mathcal{T},\\
&\bar{H} \geq 0
\end{aligned} \tag{C9}$$

The following paragraphs explain the constraints. Constraint C1 links the yearly expected line headways to the minimax objective by making $\bar{H}$ an upper bound on every yearly headway H^t. In the headway-minimisation pass, minimising $\bar{H}$ tightens it to the maximum yearly expected line headway.

Constraint C2 lower-bounds the yearly headway variable by the expected line headway for the selected layout and deployment level. For every representative ordered leader-follower scenario k, the derived ratio ρ_t weights the layout-independent fast track-releasing headway F_k and the layout-dependent slow track-releasing headway $G_k^{r(t)}$. The scenario weights π_k then aggregate these scenario-specific expected line headways into the yearly expected line headway H^t. Since ρ_t is a linear expression in the binary deployment level variables and $G_k^{r(t)}$ is a decision variable, the term $\rho_t G_k^{r(t)}$ makes C2 bilinear.

Constraint C3 lower-bounds the slow track-releasing line headway by the headway of each active TTD section. When section (i,j) is active ($x_{ij}^r = 1$), the right-hand side equals $h_{\text{TTD},k,ij}$ and forces G_k^r to at least that value. When the section is inactive ($x_{ij}^r = 0$), the right-hand side is zero and the bound is vacuous, so no big-M switching is needed. Because the first objective minimises $\bar{H}$ and thereby drives G_k^r downwards, G_k^r equals the maximum headway over the active sections in any stage and scenario that determines $\bar{H}$. The headway variables G_k^r, H^t, and $\bar{H}$ are therefore epigraph variables that the minimisation tightens to their defining values where they are binding, while they may retain harmless slack in a stage or year that does not determine $\bar{H}$.

Constraint C4 adds the certified lower bound L_k of Equation (5.11). It is redundant at the integer level, because every feasible path attains at least L_k by construction, but it tightens the linear relaxation and accelerates the branch and bound search.

Constraints C5a to C5d impose the path structure that links the layout to the arc selection. Constraint C5a sends exactly one section out of the left mandatory endpoint of each interval, and C5b lets exactly one section enter the right mandatory endpoint. Constraint C5c is flow conservation at every interior optional position, so any section entering a position is balanced by one leaving it, which prevents the path from terminating away from an endpoint. Constraint C5d links the layout to the path: position p is active exactly when the selected path visits it.

Constraints C6a and C6b impose the technical feasibility requirements on the TTD layout. Mandatory boundary positions must be active in every stage, whereas forbidden positions cannot be selected. Because forbidden positions are absent from the interval vertex sets, no arc reaches them, so C6b is enforced structurally by the arc graph as well. The remaining possible TTD boundary positions are optional and are chosen by the optimiser through the path selection.

Constraint C7 imposes the cost level $C_{\max}$ used by the ε-constraint method. It restricts the feasible alternatives to layout and deployment trajectories whose total discounted cost does not exceed the specified level.

Constraints C8a and C8b govern the TIM deployment trajectory. Constraint C8a selects one admissible deployment level $\ell \in L_t$ in each year and thereby defines ρ_t, while C8b requires the deployment ratio to remain non-decreasing over the migration horizon. The initial value $\rho_{-1} = 0$ is used to compute the first-year retrofit increment.

Finally, C9 defines the binary layout, arc, and deployment decisions and the non-negative continuous headway variables.

The formulation is arc-based: the layout couples to the slow track-releasing section headway through the linear arc constraints C3 to C5d, so no layout-dependent maximisation remains. Two nonlinearities remain: the bilinear product $\rho_t G_k^{r(t)}$ in the expected line headway of constraint (C2), and the positive-part transition operators in the infrastructure transition cost component of f_2. The computational implementation solves an objective-equivalent MILP in which the bilinear product is disaggregated by a per-level big-M disjunction driven by the deployment level binaries b_ℓ^t, and the positive-part operators are linearised through auxiliary transition binaries a_p^r, d_p^r.

Appendix B gives the MILP reformulation, including the big-M calibration and LP-relaxation strengthening. The model is implemented in PuLP and solved with Gurobi 12.0 as the branch and bound backend. Appendix A.1 reports the configured MIP optimality tolerance.

5.5.5. Pareto frontier generation and layout selection

Pareto frontier generation

The two optimisation objectives represent a trade-off: lower expected line headway can require more infrastructure or earlier TIM deployment, while lower cost limits the available capacity improvements. A Pareto frontier is used to show this trade-off. Each point on the frontier represents an alternative for which neither headway nor cost can be improved without worsening the other objective.

The frontier is generated using an ε-constraint method with lexicographic refinement. This approach avoids specifying capacity-cost weights before the trade-off is known and can identify non-convex parts of the trade-off that a weighted-sum formulation may miss in mixed-integer problems (Mavrotas, 2009).

The method evaluates a series of maximum cost levels $C_{\max}$. At each cost level, the model first asks a simple question: what is the lowest expected line headway that can be achieved without exceeding this budget? Mathematically, the first pass minimises $f_1 = \bar{H}$ subject to the model constraints and the cost bound $f_2 \leq C_{\max}$. Repeating this calculation for different values of $C_{\max}$ reveals how the achievable headway changes as the available budget increases.

The first pass alone does not necessarily identify the cheapest design for the resulting headway. Several layouts may achieve the same best headway while having different actual costs, because cost is only bounded and is not minimised in this pass. A second, lexicographic pass is therefore performed at each cost level. It minimises f_2 while requiring that $\bar{H}$ does not exceed the headway obtained in the first pass. This retains headway as the first priority and then selects the least-cost solution without worsening the first-pass headway.

After all cost levels have been evaluated, a dominance check removes generated alternatives for which another solution has equal or lower cost and equal or lower headway, with at least one of the two objectives strictly improved. The remaining alternatives form the generated Pareto frontier. Because only a finite set of cost levels is evaluated and the MIP problems are solved within the configured optimality tolerance, this set is formally a discrete approximation of the complete mathematical Pareto frontier. For readability, it is referred to as the Pareto frontier throughout the remainder of this thesis.

MCA layout selection

The Pareto frontier deliberately leaves the trade-off unresolved: no point is objectively preferable without specifying how much importance is assigned to headway relative to cost. The Multi-Criteria Analysis (MCA) step introduces these preferences only after the frontier has been generated. Each alternative combines a staged physical TTD layout with an associated TIM deployment trajectory, with the physical layout forming the main infrastructure design focus of this study.

The MCA-selected alternative is chosen using multi-criteria weights derived from the R2DATO D37.1 deployment strategy framework (Aoun, 2026a). The original D37.1 framework considers a broader set of decision criteria, whereas the analytical selection in this study uses only the two criteria represented directly by the optimisation model: expected line headway and cost. The corresponding D37.1 weights are therefore renormalised to sum to one after the other criteria are excluded.

Because headway and cost have different units and scales, both criteria are first transformed to a common scale using min-max normalisation over the Pareto alternatives. The renormalised stakeholder weights are then applied to these normalised values, and the MCA-selected alternative is the point with the lowest weighted sum (Ehrgott, 2005). The weights therefore affect only the final selection and not the generation of the Pareto frontier. Chapter 6 specifies the weights used in the case study.

Figure 5.6 illustrates the resulting Pareto generation and MCA selection procedure.

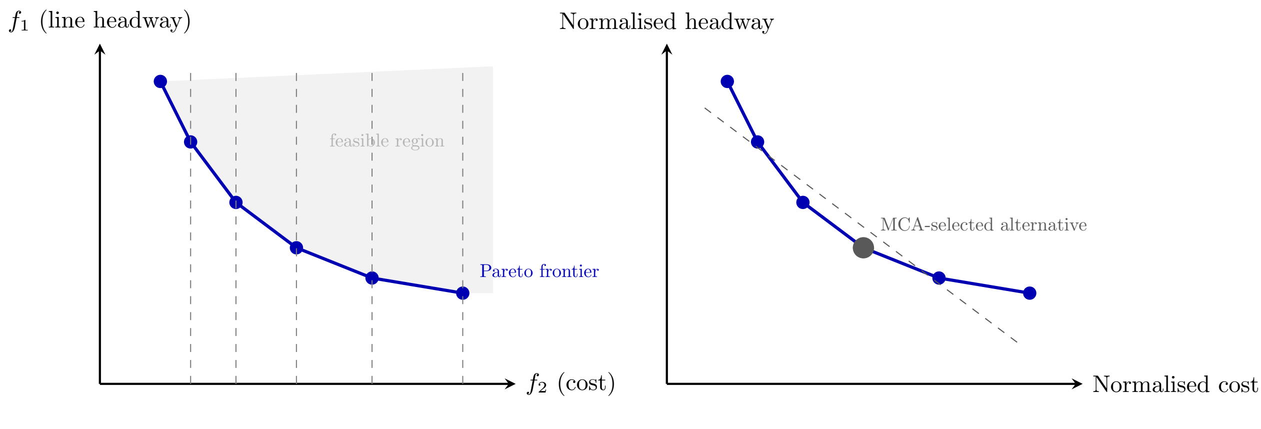

(a) Pareto frontier **(b) MCA layout selection**

Figure 5.6: Schematic illustration of the two-step decision procedure in which the Pareto frontier represents the trade-off between headway and cost (a), after which the D37.1 MCA weights are used to select one alternative (b).

5.5.6. Computational verification and spatial layout check

Before applying the optimisation framework to the case study, two complementary checks are performed. The first examines whether the computational implementation behaves as intended, while the second examines whether the optimisation responds meaningfully to the spatial placement of TTD boundaries.

Computational verification

The computational implementation is first tested on a deliberately small instance for which every feasible layout can be evaluated directly. This provides an independent reference: the MILP should identify the same best solution as the exhaustive evaluation of all feasible alternatives. After optimisation, the resulting headway and cost are also recalculated independently from the selected layout and deployment decisions, and the returned solution is checked against the mathematical constraints. Additional checks confirm that the numerical bounds used in the MILP do not exclude valid solutions and that the preprocessed quantities passed to the model are internally consistent. Finally, the solver output is inspected to determine whether accepted solutions satisfy the configured optimality tolerance. These checks assess internal model correctness rather than operational validity. Detailed procedures and results are reported in Appendix C, in particular the validation suite in Section C.5.

Spatial layout check

The second check examines whether boundary location matters independently of boundary count. Three layouts are compared under identical traffic, infrastructure, and TIM conditions. The existing layout provides the reference. The second retains the same number of active boundaries but allows their locations to be optimised, isolating the effect of placement. The third allows both the number and locations to vary and minimises cost while requiring analytical headway to remain within a predefined equivalence tolerance of the second layout. This comparison therefore distinguishes the effect of boundary location from the effect of boundary quantity. The results are reported in Section 7.1.

5.6. Evaluation approach

This section defines how the optimisation outputs are evaluated through controlled case study comparisons, analytical-microscopic consistency checks, and KPI assessment. Chapter 6 specifies the corresponding configurations and RailSys setup, while Chapter 7 reports the results.

5.6.1. Case study comparisons

The optimisation framework is applied separately to three case study types that differ in two structural characteristics: the initial TTD layout and the traffic composition. The types are a brownfield heterogeneous case, which represents the main operational setting, a greenfield heterogeneous case, and a greenfield homogeneous case. Comparing the brownfield and greenfield heterogeneous types examines the influence of the inherited brownfield detection asset base, while comparing the two greenfield types examines the influence of traffic heterogeneity under the same greenfield starting condition and common deployment assumptions. For each type the framework selects a staged TTD layout, and the selected layouts, together with their analytical expected line headway and discounted cost, are compared across types. These selected layouts subsequently enter the KPI evaluation.

The microscopic evaluations are performed in RailSys (Rail Management Consultants International GmbH, 2022). Chapter 6 specifies the corresponding infrastructure and timetable configurations.

5.6.2. Analytical-microscopic consistency

The analytical headway formulation is compared with RailSys at two levels. First, a reference condition consistency check compares the analytical and microscopic line headway for the existing brownfield TTD layout without VSS level release. This assesses the analytical approximation before the optimised layouts are interpreted.

Second, after the staged layouts have been selected, the analytical expected line headways of the heterogeneous T1 and T2 layouts are compared with the corresponding microscopic RailSys line headways under partial TIM deployment. This assesses whether the analytical approximation remains representative as both the residual TTD layout and the release regime change during migration. The homogeneous T3 configuration remains part of the capacity consumption evaluation but is not included in this additional headway consistency comparison.

In both checks, the microscopic line headway is obtained from UIC 406 compression without timetable allowances. For a periodic traffic pattern, the compressed occupation time T_{occ} contains N_{paths} repeating headway intervals, one per train path, giving

$$\bar{h}_{\text{RailSys}} = \frac{T_{\text{occ}}}{N_{\text{paths}}}.$$

This quantity can be compared directly with the analytical expected line headway, which likewise excludes timetable allowances. These comparisons assess analytical-microscopic consistency and do not constitute a separate KPI.

For partial TIM deployment, the fast and slow RailSys results are combined using the same TIM ratio and traffic-scenario weighting as in the analytical formulation. The same weighting procedure is subsequently used for the KPI 2 capacity calculation.

5.6.3. KPI evaluation

The selected layouts are evaluated against the three KPIs identified in Section 3.4: lifecycle and deployment cost, nominal capacity consumption, and operational stability. Cost follows directly from the optimisation output, capacity is evaluated for all three case study types, and operational stability is assessed for the primary T1 case against baseline B1.

KPI 1: discounted lifecycle and deployment cost

KPI 1 is the total discounted lifecycle and deployment cost obtained from the optimisation objective f_2 and therefore requires no microscopic simulation.

KPI 2: capacity

KPI 2 measures nominal capacity consumption using UIC 406 timetable compression in RailSys. Unlike the no-allowance consistency check, the applicable timetable allowances are added to the compressed occupation time. Capacity consumption is then expressed as a percentage of the one-hour evaluation period, following the principle shown in Figure 5.7.

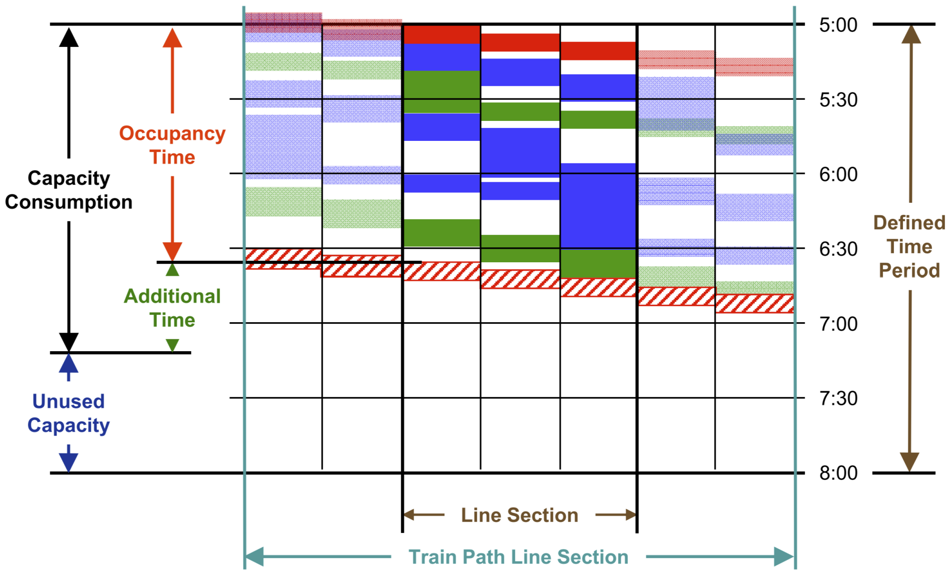

Figure 5.7: UIC 406 capacity consumption concept, distinguishing occupation time, additional time, and unused capacity within the defined evaluation period (International Union of Railways, 2013).

KPI 3: operational stability

KPI 3 evaluates operational stability ex post through the capacity margin and deterministic disturbance recovery.

Capacity margin. The capacity margin is the unused share of the evaluated hour and therefore the complement of the KPI 2 capacity consumption. It provides a static indication of the timetable space available to absorb perturbations, but does not show how an actual delay propagates.

Deterministic disturbance. Operational stability is additionally tested by applying the same predefined disturbance to each evaluated layout. Keeping the disturbed train, location, and delay magnitude fixed ensures that differences in recovery reflect the evaluated configuration. Because delay propagation depends on which individual trains can report confirmed integrity, concrete trainset-level TIM assignments are required rather than the aggregated ratio used for nominal capacity. Chapter 6 specifies the disturbance and these assignments. Recovery is evaluated through recovery time, the number of delayed trains, and total delay. The disturbance result is evaluated through three indicators:

- *Recovery time*: the elapsed time between the onset of the station delay and the moment at which all downstream arrival delays fall below the recovery threshold.
- *Number of delayed trains*: the number of downstream trains whose arrival delay exceeds the recovery threshold.
- *Total delay*: the sum of the downstream arrival delay accumulated over all trains and all evaluated stops during the disturbance simulation, expressed in seconds.

Together, the capacity margin and the deterministic delay recovery distinguish between potential and realised operational stability: the margin indicates whether the layout leaves enough timetable space to absorb perturbations, while the disturbance test shows whether this holds under a controlled disruption.

6 Case study: SAAL corridor

This chapter applies the optimisation and evaluation framework of Chapter 5 to a section of the SAAL corridor in the Netherlands. It first describes the case study corridor and its current operation, then defines the case study design together with the fixed block baseline configurations. It then specifies the shared case study parameters, deployment horizon, and cost assumptions, before describing the consistency checks and the RailSys evaluation configurations used to produce the results in Chapter 7.

6.1. Case study corridor and current operation

The case study uses a 14.269 km section of the SAAL corridor between Weesp and the approach to Almere Centrum. This section is based on the SAAL case study considered in R2DATO WP37 and forms the common corridor setting for the analyses in this chapter (Aoun, 2024, 2026a). It contains the intermediate stations Almere Poort (AMPO) and Almere Muziekwijk (ALMM) and carries a heterogeneous mix of Intercity (IC) and Stop Train (ST) services. The case study baseline represents this corridor under ETCS L2 fixed block signalling. The ST services stop at both intermediate stations, whereas the IC services run non-stop through the section. During the peak hour the corridor follows a repeating quarter-hour pattern of one IC and three ST services (IC-ST-ST-ST), giving 16 train paths per hour. These are assumed to be operated by eight unique trainsets on a half-hour circulation, so each trainset runs two paths per hour.

Figure 6.1 shows the IC and ST maximum speed profiles together with the infrastructure elements relevant to TTD boundary placement. The shaded areas indicate catenary overlap and switch zones, while AMPO and ALMM mark the two intermediate station locations. The different stopping patterns are visible in the speed profiles, with the ST services stopping at both stations and the IC services continuing through the corridor.

The infrastructure elements shown in Figure 6.1 are translated into mandatory and forbidden boundary positions using the design requirements of Chapter 4. The orange markers along the lower axis indicate the resulting mandatory TTD positions used by the optimisation model. For readability, the figure shows only these mandatory positions. The initial brownfield layout contains further TTD positions in addition to the mandatory ones, and the complete brownfield TTD layout is specified in Section 6.3 and Appendix A.3.

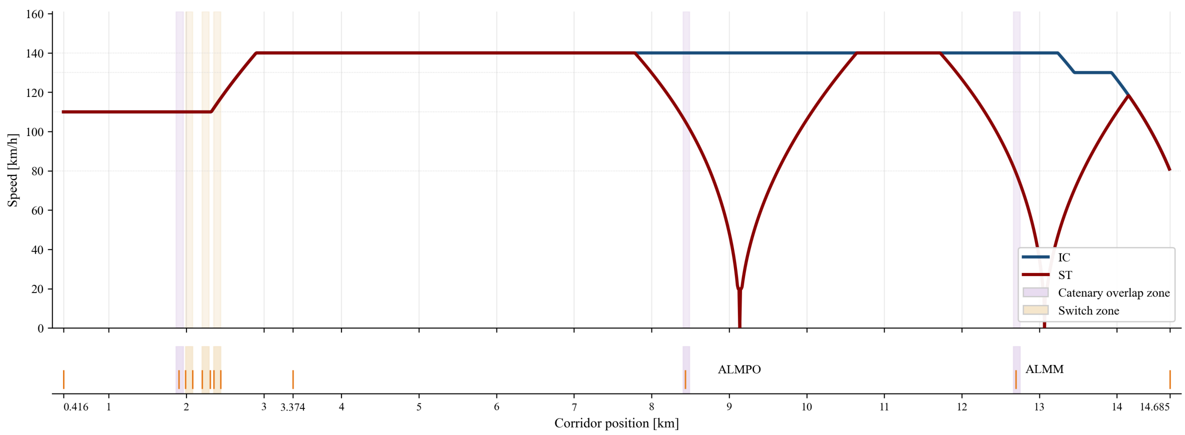

Figure 6.1: Speed profiles and mandatory TTD positions for the case study corridor section.

The existing brownfield TTD layout and heterogeneous IC/ST traffic therefore define the primary brownfield starting condition for the case study. The physical corridor geometry is kept fixed, while the initial infrastructure condition and traffic composition are varied to examine their influence on the layouts selected by the optimisation model.

6.2. Case study design

Based on this starting condition, three case study types are defined. T1 retains both the existing brownfield TTD layout and the heterogeneous IC/ST traffic and therefore forms the primary operational case. T2 replaces the brownfield layout with a greenfield layout while retaining the same traffic, so that comparison with T1 examines the influence of the inherited brownfield detection asset base. T3 retains the greenfield setting of T2 but replaces the mixed traffic with homogeneous IC-only traffic, so that comparison with T2 examines the influence of traffic heterogeneity under the same greenfield starting condition and common deployment assumptions. All three types use the same corridor section and shared case study assumptions unless explicitly stated otherwise. Table 6.1 summarises the three types.

Table 6.1: Case study types considered in this study.

Type	Initial TTD layout	Traffic setting
T1	Brownfield	Heterogeneous
T2	Greenfield	Heterogeneous
T3	Greenfield	Homogeneous

For each case study type, the optimisation model is solved over the multi-stage migration horizon defined in Section 6.4, producing a selected TTD layout for each stage. For the staged configurations, *Ti-Sj* denotes type i evaluated at migration stage j. In addition to the three optimised case study types, two fixed block reference configurations are used in the microscopic evaluation. B1 is the heterogeneous ETCS L2 fixed block baseline for T1 and T2, while B2 is the corresponding homogeneous baseline for T3. Table 6.4 in Section 6.7 gives the full configuration overview, including the staged T1-Sj to T3-Sj layouts.

For each case study type, one alternative is selected from the Pareto frontier using the MCA procedure of Section 5.5.5. The Cost (21 %) and Capacity (25 %) weights of the R2DATO D37.1 deployment strategy framework (Aoun, 2026a) are renormalised over the two optimisation criteria, giving $w_C = 0.457$ and $w_H = 0.543$. The same weights are applied to all three case study types to provide a consistent comparative selection rule.

6.3. Shared case study parameters

For the analytical model, each service class is represented by one deterministic train profile. The heterogeneous T1 and T2 cases use the leader-follower proportions implied directly by the repeating IC-ST-ST-ST pattern described in Section 6.1, while T3 contains only the homogeneous IC→IC scenario.

Table 6.2 summarises the main input parameters for the case study. The structural differences between the three types are defined in Table 6.1.

Table 6.2: Main input parameters for the SAAL case study.

Parameter	Case study value
Corridor and infrastructure	
Section length	14.269 km
Intermediate stations	AMPO km 9.132, ALMM km 13.068
Mandatory TTD positions	12
Forbidden infrastructure zones	9
Nominal VSS spacing d_{VSS}	50 m (Aoun, 2025)
Rolling stock and traffic	
Acceleration	0.5 m/s^2 (Scheepmaker et al., 2020)
Maximum train length	340 m (IC), 270 m (ST)
Station dwell time	30 s
MA processing time t_{MA}	10 s (Jansen, 2019; Vergroesen, 2020)
Release delay t_{release}	2 s (Versluis et al., 2025)
Peak corridor frequency	16 train paths/h
Service pattern	IC-ST-ST-ST
Corridor fleet N_{train}	8 unique trainsets
Migration	
Planning horizon T	8 years
Layout stages R	4 stages of 2 years
Terminal TIM deployment ρ	0.95 in year 8
Cost and lifecycle	
Axle counter placement	€ 50,000 (assumed)
Axle counter removal	€ 5,000 (assumed)
Axle counter maintenance	€ 5,000/year (assumed)
TIM retrofit	€ 84,420/train (Aoun, 2026a)
Social discount rate d	2.8 % (Werkgroep discontovoet, 2025)
Axle counter asset lifetime	25 years

The case study uses a nominal VSS spacing of $d_{\mathrm{VSS}} = 50$ m. This corresponds to the lower end of the approximately 50-200 m range of the fixed short block regime adopted in Section 5.4 from the R2DATO D37.2 methodology (Aoun, 2025). The finest spacing is selected because each VSS boundary forms a possible location for a physical TTD boundary. A finer VSS grid therefore provides the optimiser with a larger set of possible locations and greater flexibility in selecting the TTD layout, at the expense of a larger optimisation problem.

Applying the design procedure of Section 5.4 to the case study corridor yields 12 mandatory TTD positions. T1 starts from an inherited brownfield initial layout of 44 TTD boundaries (12 mandatory positions plus 32 selectable inherited axle counters within the corridor), whereas T2 and T3 use a greenfield initial state ($\hat{y} = 0$). The resulting design domain contains 278 possible TTD boundary positions for T1 and 281 for T2 and T3. Excluding the nine forbidden positions leaves 269 feasible positions for T1 and 272 for T2 and T3, each of which carries a retention variable in every migration stage.

6.4. TIM deployment and migration horizon

Alongside the spatial design domain, the case study assumes an eight-year migration horizon over which the fleet is progressively equipped with TIM, reflecting the planning assumption that rolling stock is retrofitted gradually rather than all at once. Rolling stock TIM deployment may change annually, whereas modifications to the physical TTD layout are assumed to take place once every two years, reflecting the assumption that trackside infrastructure changes require longer planning and implementation lead times. The horizon is therefore divided into four two-year infrastructure stages, within which the TTD layout remains fixed while the TIM deployment level may change from year to year. For readability, migration years and infrastructure stages are numbered 1-8 and 1-4 in this chapter, corresponding to the model indices $t = 0, \ldots, 7$ and $r = 0, \ldots, 3$, respectively.

Within this temporal structure, the model selects the TIM deployment level from a discrete admissible set in each year, as listed in Table 6.3. The sets use 5 percentage-point increments to provide flexibility in retrofit timing while limiting the number of binary level selection variables. The admissible ranges increase over the migration horizon, and the non-decreasing deployment constraint prevents the selected deployment level from decreasing over time.

Table 6.3: Discrete TIM deployment level sets per year for the case study.

Migration year	Deployment levels L_t	$\lvert L_t \rvert$
1	$\{0.05,\ 0.10,\ 0.15\}$	3
2	$\{0.10,\ 0.15,\ 0.20\}$	3
3	$\{0.20,\ 0.25,\ 0.30,\ 0.35\}$	4
4	$\{0.30,\ 0.35,\ 0.40,\ 0.45,\ 0.50\}$	5
5	$\{0.45,\ 0.50,\ 0.55,\ 0.60,\ 0.65\}$	5
6	$\{0.60,\ 0.65,\ 0.70,\ 0.75,\ 0.80\}$	5
7	$\{0.75,\ 0.80,\ 0.85,\ 0.90\}$	4
8	$\{0.95\}$	1
Total binary level selection variables		30

The final year is fixed at $\rho = 0.95$, representing near-full rather than complete TIM deployment and allowing for a small share of rolling stock that may remain unequipped at the end of the migration. Temporary unavailability of confirmed train integrity reporting on equipped trains is an operational condition and is not represented by a reduction in the deployment level.

6.5. Lifecycle and deployment cost assumptions

The infrastructure unit costs in Table 6.2 are order of magnitude planning assumptions rather than project-specific procurement values. Cost results are therefore interpreted as relative lifecycle comparisons rather than exact investment budgets.

Alongside these infrastructure costs, a per-train TIM retrofit cost of €84,420 represents the rolling stock component, based on the R2DATO D37.1 system level estimate (Aoun, 2026a). This value is necessarily approximate because retrofit costs vary by rolling stock type, onboard architecture, and procurement strategy. The value is therefore used as a corridor-attributed planning parameter for comparing layout alternatives, not as a fleet-specific procurement estimate. The sensitivity analysis in Chapter 7 uses the alternative settings of €0 and €42,500 per train only.

Eight unique trainsets represent the corridor fleet. The evaluated peak pattern contains 16 train paths per hour, but the assumed half-hour circulation means that each trainset appears twice within the hourly pattern. The model therefore scales the TIM retrofit cost by the number of unique trainsets rather than by the number of train paths. At the terminal deployment target of 0.95, this corresponds to an undiscounted corridor-attributed retrofit cost of €641,592.

Beyond the migration horizon, the terminal maintenance term accounts for the fact that axle counters active at the end of the 8-year planning horizon continue to incur maintenance costs during their remaining asset lifetime. With a 2.8 % social discount rate and an assumed axle counter asset lifetime of 25 years, the discounted terminal maintenance factor equals 11.029. At a maintenance cost of €5,000 per active boundary per year, this gives a residual maintenance liability of €55,144 per boundary active in the final stage.

Finally, a full-retention T1 reference, in which all 44 inherited TTD boundaries remain active throughout the migration, is additionally evaluated for the lifecycle cost comparison.

Computational implementation details, model dimensions, solver settings, and the cost-level grid used to generate the Pareto frontiers are reported in Appendix A.1.

6.6. Consistency-check configurations

The configurations below instantiate the two checks defined in Chapter 5: the reference-condition analytical-microscopic consistency check of Section 5.6.2 and the spatial layout check of Section 5.5.6. Their numerical outcomes are reported in Chapter 7.

Reference condition comparison. The reference condition check uses the existing brownfield layout of 44 active TTD boundaries under the heterogeneous case study traffic. The analytical representation is evaluated at the conventional ETCS L2 physical release condition with no VSS level release, that is $\rho = 0$. The microscopic comparator is B1, the heterogeneous ETCS L2 fixed block baseline, with the RailSys line headway derived from the compressed occupation without timetable allowances so that it is comparable with the analytical value.

Spatial experiment. The spatial experiment compares three layouts on the same corridor, heterogeneous traffic, $\rho = 0$, infrastructure geometry, possible TTD design domain, and analytical headway formulation. Layout 1 is the existing 44-boundary brownfield layout. Layout 2 retains 44 active TTD boundaries with their locations optimised for minimum expected line headway. Layout 3 allows both the number and locations to vary, minimising physical placement cost subject to reproducing the Layout 2 analytical headway within the defined equivalence tolerance. All non-layout inputs are held constant, so the first comparison isolates the spatial response at a fixed detection quantity and the second tests whether comparable analytical capacity performance is attainable at lower physical detection cost.

Selected-layout consistency scope. The additional analytical-microscopic consistency check under partial TIM deployment is reported for the selected heterogeneous T1 and T2 layouts. The homogeneous T3 layout remains part of the microscopic capacity consumption evaluation but is not included in this headway consistency comparison.

6.7. RailSys case study setup

This section specifies the RailSys configurations used to apply the evaluation approach of Section 5.6 to the case study corridor. Nominal capacity is evaluated for T1, T2, and T3 against their traffic-matched ETCS L2 fixed block baselines, while deterministic disturbance recovery is evaluated for the primary T1 case against B1.

The RailSys evaluation distinguishes between the signalling and release configuration, the physical TTD layout, and the traffic setting. The brownfield VSS grid, which includes the existing axle counter positions, is used for T1, whereas the greenfield VSS grid is used for T2 and T3. The fixed block baselines B1 and B2 contain no VSS level release. Table 6.4 summarises the configurations, using the notation of Section 6.2.

Table 6.4: RailSys configurations used for the microscopic evaluation.

Config.	System	TTD layout	Traffic	Evaluation
B1	ETCS L2 fixed block	Brownfield	Heterogeneous	Capacity baseline for T1 and T2, and operational stability baseline for T1
B2	ETCS L2 fixed block	Brownfield	Homogeneous	Capacity baseline for T3
T1-Sj	ETCS HTD	T1 selected stage-j layout	Heterogeneous	Capacity and operational stability
T2-Sj	ETCS HTD	T2 selected stage-j layout	Heterogeneous	Capacity
T3-Sj	ETCS HTD	T3 selected stage-j layout	Homogeneous	Capacity

6.7.1. Timetable and deployment trajectory

The T1 and T2 RailSys timetables reproduce the heterogeneous operating pattern defined in Section 6.1, while T3 uses the homogeneous IC-only traffic setting defined in Section 6.2. KPI 3 is evaluated only for Type 1, the primary operational case with the current brownfield infrastructure and heterogeneous traffic. Extending it to the counterfactual Types 2 and 3 would require additional trainset-level TIM assignments and simulations without adding to the operational stability assessment.

For the capacity evaluation, the fast and slow track-releasing cases are aggregated using the deterministic TIM ratio and scenario weighting of Section 5.6.3. Each selected stage layout is evaluated at the end of stage TIM ratio of its MCA-selected deployment trajectory, reported with the results in Chapter 7. Because the TTD layout remains fixed within each two-year stage while deployment can change annually, the final-year ratio of each stage represents the operational state immediately before the next infrastructure modification. A full deployment reference at $\rho = 1.00$, representing full VSS level release within the same timetable and VSS configuration, is additionally evaluated for each traffic setting. For the disturbance recovery evaluation, a concrete trainset-level TIM assignment represents the T1 trajectory, because fractional post-processing cannot reconstruct delay propagation (Section 6.7.2).

6.7.2. Disturbance scenario for operational stability (KPI 3)

The KPI 3 disturbance is a 15 min standstill of the first ST service at AMPO. The test compares the T1 selected staged ETCS HTD layout with the ETCS L2 fixed block baseline. The comparison assesses whether recovery remains comparable to the fixed block baseline under the defined disturbance. Operational stability under the deterministic disturbance is evaluated using the three recovery indicators defined in Section 5.6.3: recovery time, number of delayed trains, and total delay. Recovery time is measured as the elapsed time until all downstream arrival delays fall below 60 s. Because TIM capability is a rolling stock property, the assignment operates at trainset level: the evaluated hour contains 16 train paths run by 8 unique trainsets, with trainset U_i operating paths T_i and T_{i+8}, so equipping U_i makes both its paths integrity reporting.

Each stage is evaluated under a constructed best-case and worst-case TIM assignment, fixed before the recovery outcome is observed. These cases represent the two assignment extremes considered in this study rather than an exhaustive optimisation over all possible trainset-level assignments. The best-case assignment equips the disturbed trainset U_2 first and then the trainsets immediately behind it, so the delayed train can release VSSs behind itself and reduce the blocking of following trains. The worst-case assignment instead equips the least helpful trainsets first, the preceding trainset U_1 and then the rear of the hourly sequence working forward, so U_2 and its immediate followers stay unequipped longest, exactly where VSS level release would most relieve the standstill. Table 6.5 lists both assignments.

Because TIM capability is a rolling stock property, KPI 3 uses integer trainset counts. The equipped share must therefore be a multiple of $1/8 = 0.125$, so the trainset-level assignment provides a discrete operational representation of the continuous T1 end of stage TIM ratio selected by the optimisation model. The T1 TIM ratio column of Table 6.5 reports this end of stage ratio, and the effective column reports the discrete multiple of 1/8 produced by the assignment. For stage 2, the $3/8 = 0.375$ assignment slightly over-represents the selected T1 end of stage ratio of 0.30 and should therefore be interpreted as a discrete operational representation rather than an exact match.

Table 6.5: Best-case and worst-case TIM assignments for the T1 KPI 3 delay-recovery simulations, with effective deployment ratios constrained to multiples of 1/8 and the T1 ratio column reporting the selected end of stage deployment ratio represented by each assignment.

Stage	T1 TIM ratio	Effective TIM ratio	Best-case assignment	Worst-case assignment
1	0.15	0.125	U_2	U_1
2	0.30	0.375	U_2, U_3, U_4	U_1, U_8, U_7
3	0.65	0.625	U_2, U_3, U_4, U_5, U_6	U_1, U_8, U_7, U_6, U_5
4	0.95	1.000	$U_1, \ldots, U_8$	$U_1, \ldots, U_8$

Together, the case study types, fixed block baseline configurations, and RailSys scenarios define the evaluation cases used in Chapter 7.

7 Results

This chapter reports the case study results for the staged TTD layouts selected during migration from ETCS L2 to HTD on a section of the SAAL corridor between Weesp and Almere Centrum. After assessing analytical-microscopic consistency and the spatial response of the optimisation, the results are organised around three main findings: the value of inherited physical detection, the traffic-dependent spatial value of residual TTD boundaries, and the relationship between nominal cost-capacity performance and operational stability. All monetary values are case study estimates based on the unit cost assumptions of Chapter 6 and are intended for relative comparison rather than as project level investment estimates.

7.1. Analytical-microscopic consistency and spatial response

Before interpreting the case study results, the analytical model is compared with a microscopic RailSys simulation of the current ETCS L2 brownfield layout, following the reference condition analytical-microscopic consistency check defined in Section 6.6. At $\rho = 0$, all relevant leader-follower movements depend on TTD level release, representing conventional ETCS L2 fixed block operation. For the existing layout of 44 TTD boundaries, the model yields an expected line headway of 155.15 s, compared with 162.00 s in RailSys (Appendix D.1). The resulting difference of 4.2 % indicates close agreement at the reference condition and supports the use of the analytical headway model for optimisation. For the selected heterogeneous T1 and T2 migration layouts, the analytical and RailSys line headways remain within approximately 1-5 s across all four stages, extending this consistency to partial TIM deployment (Appendix D.1). A spatial experiment then examines the spatial optimisation logic, following the setup defined in Section 6.6 (Figure 7.1). The figure compares three layouts at the reference condition: the current brownfield layout of 44 boundaries ($\bar{H} = 155.15$ s), a layout that keeps the boundary count fixed at 44 but optimises their locations for minimum expected line headway ($\bar{H} = 151.98$ s), and a least-cost layout constrained to reproduce that optimised headway within the equivalence tolerance, which is achieved with only 26 boundaries ($\bar{H} = 156.42$ s). Optimising the locations at a fixed count therefore lowers the expected line headway, and the optimised headway can be reproduced within the tolerance using substantially fewer boundaries. The experiment shows that boundary placement, and not only boundary count, governs the analytical capacity performance, and that comparable performance can be attained at lower physical detection cost.

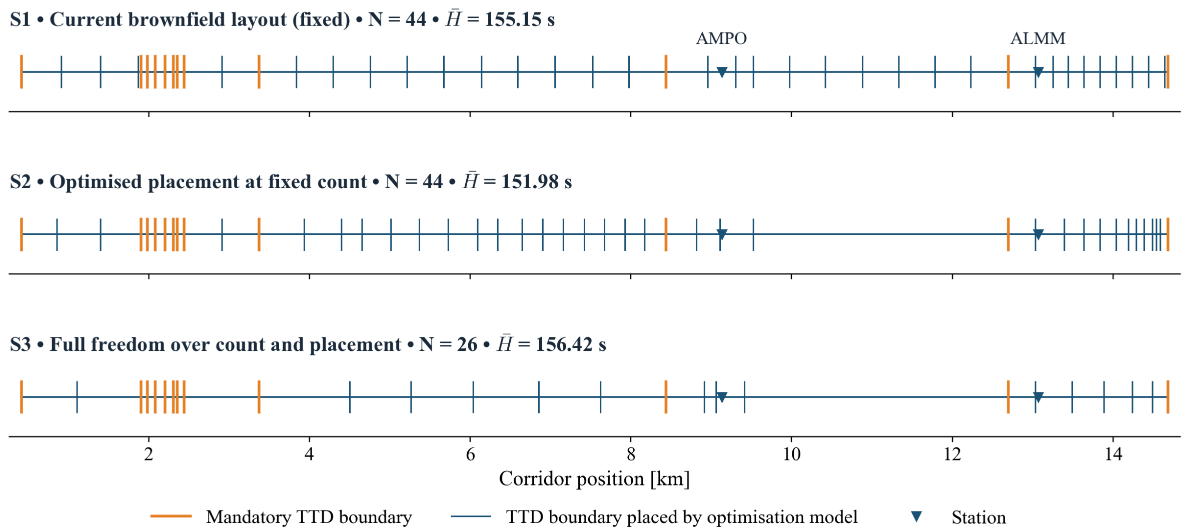

Figure 7.1: Comparison of the current, optimised 44-boundary, and least-cost TTD layouts on the SAAL corridor at $\rho = 0$.

7.2. Infrastructure inheritance and staged asset reuse

The effect of infrastructure inheritance is examined by comparing T1 and T2 under the same heterogeneous traffic and TIM deployment assumptions. T1 starts from the existing ETCS L2 brownfield layout with 44 inherited TTD boundaries, whereas T2 uses a greenfield starting condition in which all required physical detection must be newly installed. This comparison isolates the effect of the infrastructure starting condition on the cost of achieving comparable nominal capacity performance. The selected alternatives achieve almost identical maximum expected line headways, with 162.18 s for T1 and 162.58 s for T2, while discounted lifecycle and deployment cost increases from € 2.03 M to € 2.82 M (Figure 7.2). A similar cost difference is observed across the Pareto frontiers, indicating that inherited infrastructure primarily reduces the cost of attaining a given headway rather than changing the attainable headway itself.

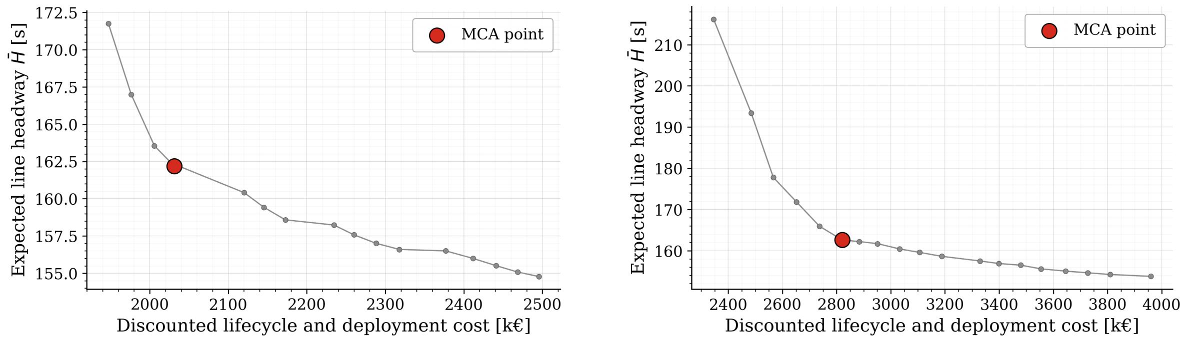

(a) T1 brownfield migration.

(b) T2 greenfield deployment.

Figure 7.2: Pareto frontiers for the brownfield T1 and greenfield T2 migration approaches.

The selected layouts also remain spatially similar (Figure 7.3). T1 retains 21, 20, 16, and 12 active boundaries across the four stages, compared with 19, 18, 16, and 12 for T2. In both cases, optional boundaries are concentrated around AMPO and ALMM, and both trajectories converge to the same 12 mandatory positions in the final stage. The exact per-stage positions for all three case study types are reported in Appendix A.4. The lower cost of T1 therefore arises mainly because useful inherited boundaries can remain in service without a new placement cost, whereas equivalent boundaries in T2 must first be installed.

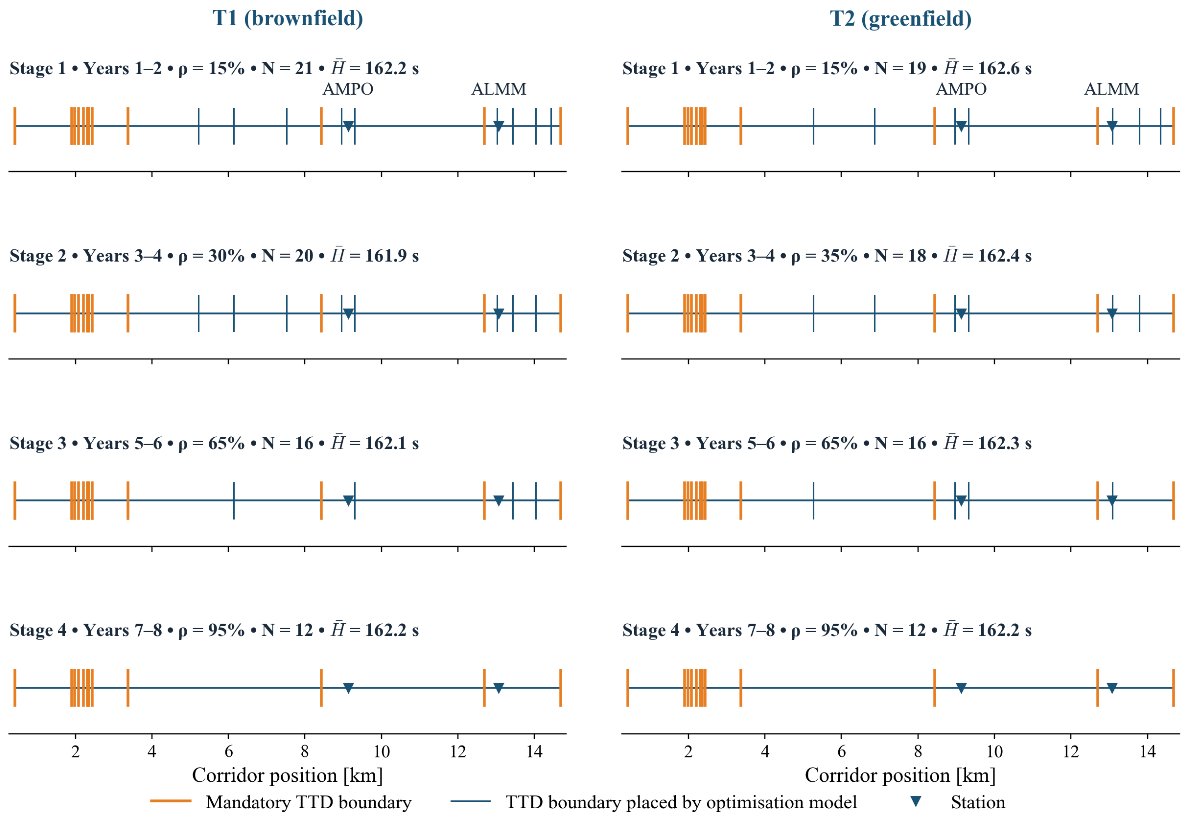

Figure 7.3: Selected staged TTD layouts for the brownfield T1 and greenfield T2 migration approaches.

Across the migration stages, the selected trajectories predominantly reduce residual physical detection rather than repeatedly relocating boundaries. This pattern is consistent with the declining nominal capacity value of additional physical subdivision as a larger share of relevant train interactions uses VSS level release. The magnitude of the brownfield cost advantage depends on the case study unit cost assumptions, but the comparison shows that inherited TTD assets can retain economic value throughout the staged migration. The coupling between fleet deployment and residual physical detection is also visible in the T1 TIM retrofit cost sensitivity. Retrofit costs of €0, €42 500, and €84 420 per train produce selected boundary trajectories of $\{30, 25, 20, 14\}$, $\{25, 23, 18, 13\}$, and $\{21, 20, 16, 12\}$, respectively (Appendix A.5). Changing the retrofit cost therefore shifts the joint trade-off between fleet deployment and residual physical detection, showing that under a combined TIM and infrastructure deployment perspective, the selected TTD migration path depends partly on the economics of TIM deployment.

7.3. Traffic-dependent spatial value of TTD boundaries

The effect of traffic composition is examined by comparing the two greenfield cases, which share the same infrastructure starting condition and the same freedom to optimise TTD boundary locations. T2 represents heterogeneous traffic, whereas T3 represents homogeneous traffic. The substantially lower expected line headway of T3 (94.41 s compared with 162.58 s for T2) primarily reflects the lower attainable headway under homogeneous traffic and should therefore not be interpreted as superior optimisation performance. The more informative difference is spatial (Figure 7.4). Under heterogeneous traffic, optional TTD boundaries are concentrated more strongly around the station areas, whereas the homogeneous layout distributes residual detection more evenly along the corridor and retains additional boundaries on the open line between AMPO and ALMM. The same spatial pattern persists across the other generated non-dominated alternatives, indicating that the difference is systematic rather than specific to the selected alternatives.

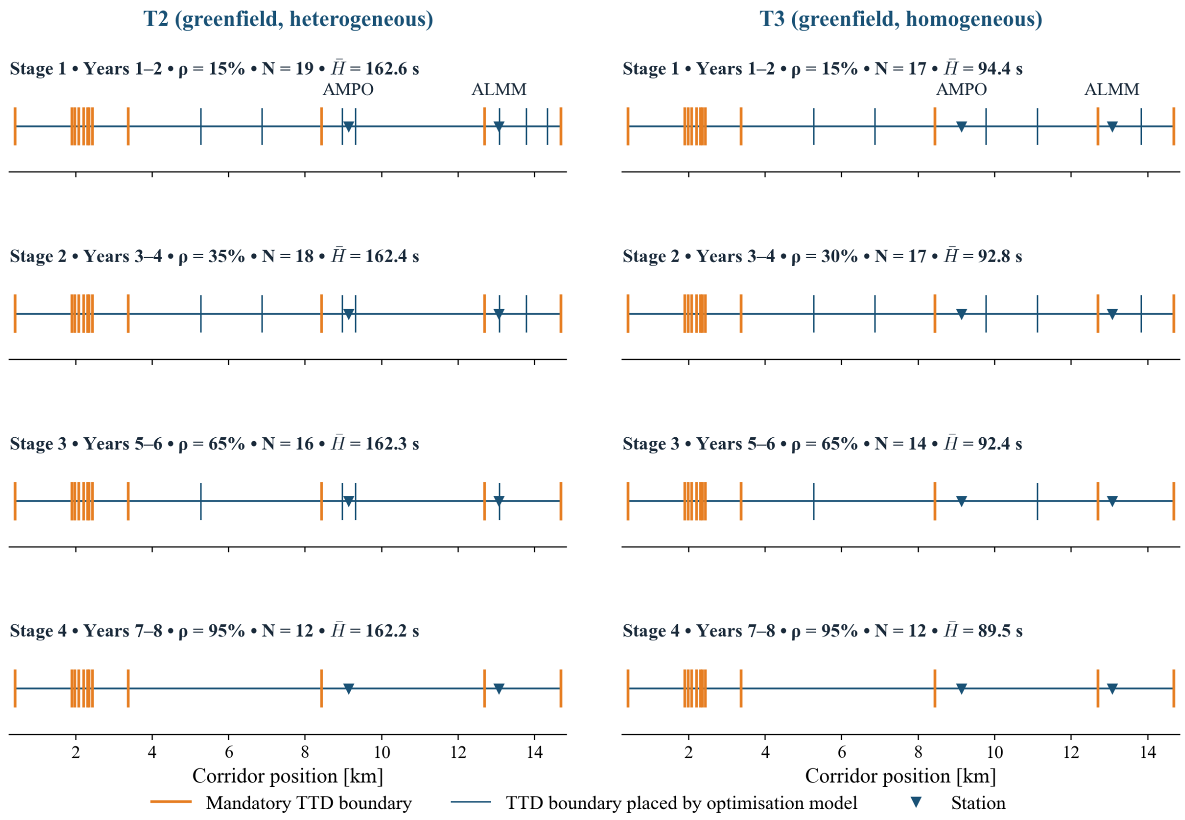

Figure 7.4: MCA-selected staged TTD layouts across the four migration stages for T2 (greenfield, heterogeneous) and T3 (greenfield, homogeneous).

This difference follows from the cumulative running time differences between heterogeneous train pairs. Differences in braking, stopping, dwell, and acceleration cause stopping and non-stopping trajectories to diverge along the corridor, making some TTD sections more restrictive than others. A TTD boundary therefore has greatest nominal capacity value where it subdivides a section governing a critical heterogeneous leader-follower interaction. In the homogeneous T3 case, the identical intercity services have no cumulative running time difference within the case study corridor, so the spatial value of residual detection is distributed more evenly. This contrast is particularly visible on the open-line section between AMPO and ALMM. T2 retains no optional boundary within this section, whereas T3 does. For the heterogeneous traffic pattern, additional subdivision of this section does not improve the binding headway sufficiently to justify its cost, and the same absence of optional subdivision occurs in the heterogeneous T1 migration. This conclusion is specific to the nominal cost-capacity objective. The operational relevance of this long TTD section under disturbed operation is revisited in Section 7.4.3. Overall, the comparison shows that the value of a TTD boundary depends not only on its location, but also on the train interactions occurring at that location.

7.4. Capacity, cost, and operational stability

The MCA-selected migration paths are now evaluated against lifecycle cost, nominal capacity consumption, and operational stability. Each selected path combines a staged TTD boundary trajectory, reported in Sections 7.2 and 7.3 and in Appendix A.4, with an MCA-selected TIM deployment trajectory. Table 7.1 lists the end of stage TIM ratios of the T1, T2, and T3 trajectories used in the microscopic evaluation. The three types share the stage-1, stage-3, and stage-4 ratios and differ only in stage-2, where T1 and T3 reach 0.30 while T2 reaches 0.35.

Table 7.1: End of stage TIM ratios of the T1, T2, and T3 MCA-selected trajectories, used in the RailSys capacity evaluation.

Layout stage	Migration years	T1	T2	T3
1	Years 1-2	0.15	0.15	0.15
2	Years 3-4	0.30	0.35	0.30
3	Years 5-6	0.65	0.65	0.65
4	Years 7-8	0.95	0.95	0.95

7.4.1. Capacity consumption

The capacity results show a common migration pattern. In the early stages, the selected HTD layouts consume more capacity than their traffic-matched ETCS L2 fixed block baselines. As TIM deployment increases, this difference becomes smaller, and by the final stage all three selected layouts reach a capacity consumption close to or slightly below their corresponding baseline.

Table 7.2 reports the UIC 406 capacity consumption (KPI 2) for the four migration stages. The heterogeneous T1 and T2 layouts are compared with B1, which consumes 93.7 % of the evaluated hour. T1 and T2 consume up to approximately 2.5 percentage points more capacity than B1 in stages 1-3, but decrease to 93.2 % and 93.0 % in stage 4. The homogeneous T3 layout follows the same overall pattern relative to B2. Capacity consumption is between 61.4 % and 63.2 % in stages 1-3 compared with 55.8 % for B2, before decreasing to 55.2 % in stage 4. All selected layouts remain within the one-hour evaluation period in every stage.

Table 7.2: RailSys UIC 406 capacity consumption for the selected layouts and their traffic-matched ETCS L2 fixed block baselines B1 and B2 across the four migration stages.

Configuration	Traffic	S1	S2	S3	S4
B1 (heterogeneous fixed block baseline)	Heterogeneous	93.7 %			
B2 (homogeneous fixed block baseline)	Homogeneous	55.8 %			
T1	Heterogeneous	96.0 %	95.3 %	95.1 %	93.2 %
T2	Heterogeneous	96.2 %	95.7 %	95.5 %	93.0 %
T3	Homogeneous	63.2 %	61.4 %	62.3 %	55.2 %

As TIM deployment increases across the migration stages, the remaining dependence on TTD level release decreases and the capacity consumption approaches the traffic-matched baseline. For reference, full TIM deployment ($\rho = 1.00$) gives a capacity consumption of 90.7 % for the heterogeneous traffic setting and 52.9 % for the homogeneous setting. These values provide the case study full-VSS-release reference for the evaluated timetables and approximately 50 m VSS configurations. The stage 4 values are therefore already relatively close to this reference.

The only non-monotonic pattern occurs for T3 between stages 2 and 3, where capacity consumption increases slightly from 61.4 % to 62.3 %. Over the same transition, the selected physical layout is reduced from 17 to 14 active TTD boundaries while the TIM ratio increases from 0.30 to 0.65. The selected migration path therefore accepts a small increase in capacity consumption in exchange for a sparser physical layout within the joint lifecycle cost-capacity trade-off.

Figure 7.5 provides a spatial example of the capacity result for T1 stage 2. The binding critical block sections occur mainly on the approaches to AMPO and ALMM and at the corridor entry. These locations lie close to several optional TTD boundaries retained by the selected T1 layout. No binding contact point occurs within the long open-line section between AMPO and ALMM. For this nominal compressed timetable, additional physical subdivision of that section therefore does not determine the line headway.

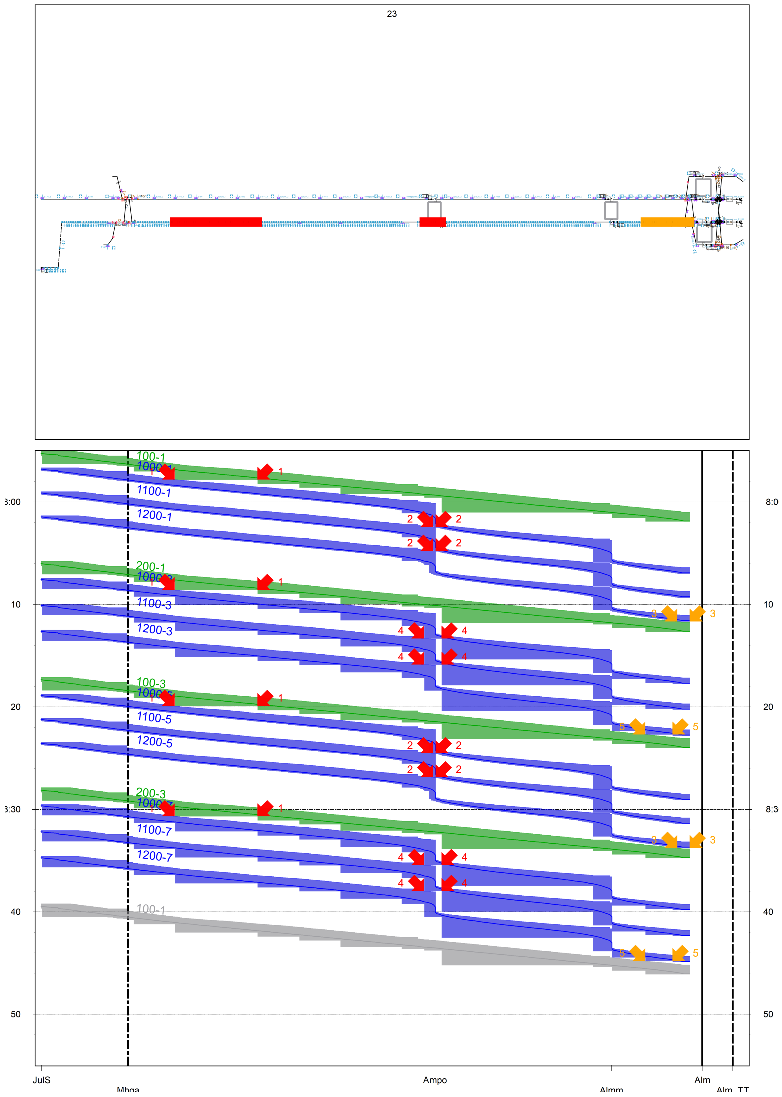

Figure 7.5: RailSys UIC 406 compressed timetable of the T1 stage 2 layout, with red arrows marking the binding critical block sections that set the line headway, using the same compression procedure applied to every case study type and migration stage.

7.4.2. Lifecycle cost relative to full retention

KPI 1 compares the lifecycle cost of the selected T1 migration path with a full-retention reference that introduces ETCS HTD on the same 50 m VSS grid and under the same TIM deployment trajectory, but removes no physical detection and keeps all 44 inherited TTD boundaries active throughout the migration.

Under the case study cost assumptions, the selected trajectory has a discounted lifecycle and deployment cost of approximately €2.03 M. The full-retention reference costs approximately €4.61 M. The selected migration path therefore reduces the modelled discounted cost by approximately €2.58 M, or 56 %.

This lower cost comes with a higher maximum expected line headway. The headway increases from 154.51 s for full retention to 162.18 s for the selected trajectory, a difference of 7.67 s, or approximately 5 %.

The comparison therefore shows that staged reduction of the inherited TTD layout can substantially lower lifecycle cost relative to full retention, while requiring only a comparatively small increase in maximum expected line headway.

7.4.3. Operational stability under deterministic disturbance

For the primary T1 case, the complementary capacity margin provides the static component of the operational stability assessment (KPI 3). The available margin increases from 4.0 % in stage 1 to 4.7 %, 4.9 %, and 6.8 % in stages 2-4, compared with 6.3 % for the B1 fixed block baseline. The selected layout therefore provides less nominal spare capacity than the baseline during the first three migration stages, while the final stage slightly exceeds the baseline margin. This static margin does not by itself determine realised disturbance recovery. The deterministic disturbance defined in Section 6.7.2 is then used to assess realised delay recovery for the T1 migration stages under the constructed best-case and worst-case TIM assignments. Figure 7.6 reports the resulting recovery time, number of delayed trains, and total delay. In the early stages, recovery is materially slower than the ETCS L2 fixed block baseline. In stage 1, recovery takes 44.23-48.00 min against 33.00 min for the baseline, with 11-12 delayed trains against 8 and 11 674-12 741 s of total delay against 8 786 s. Stage 2 improves but remains worse than the baseline on all three indicators. By stage 3, recovery time and the number of delayed trains have effectively converged to the baseline, while total delay remains scenario-dependent, with the best-case at 8 381 s below and the worst-case at 9 975 s about 13.5 % above the baseline. In stage 4, all three indicators equal or improve on the baseline, with a total delay of 8 140 s (7.4 % below baseline). Across the migration stages, both TIM deployment and the residual TTD layout change between stages, so the staged comparison alone does not isolate the contribution of either mechanism to the recovery improvement.

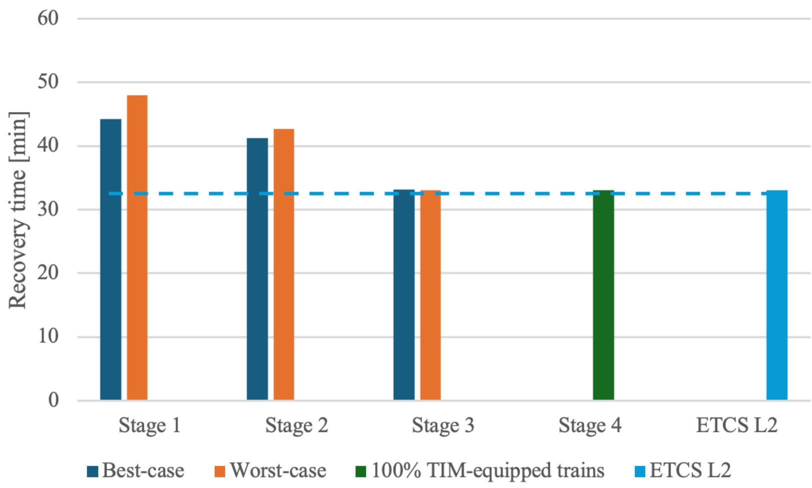

(a) Recovery time [min].

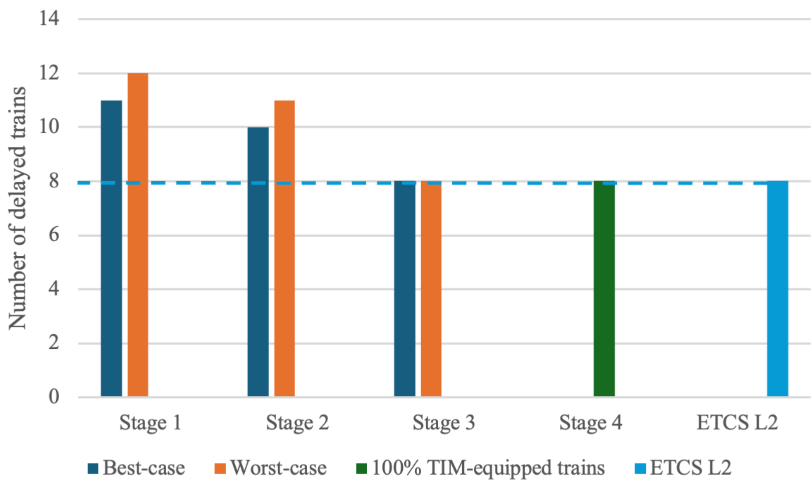

(b) Delayed trains.

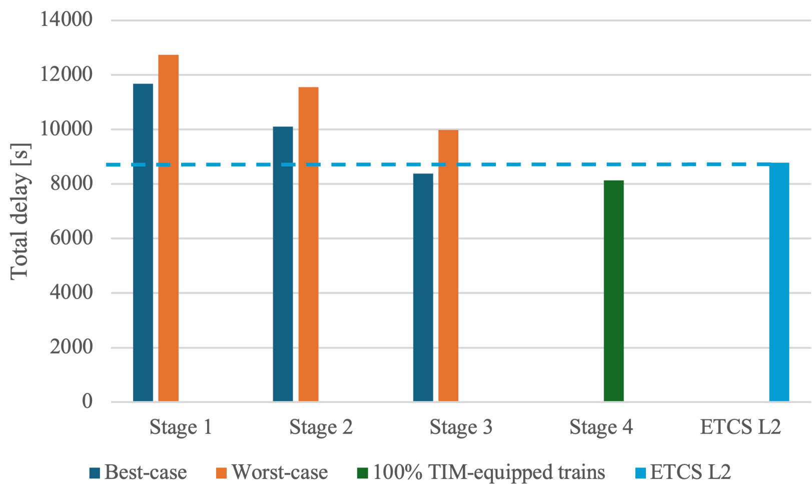

(c) Total delay [s].

Figure 7.6: Disturbance recovery indicators after the T1 15 min standstill at AMPO, across the four migration stages under best-case and worst-case TIM assignments.

The time-distance diagrams provide a spatial view of this disturbance propagation for stage 2 (Figure 7.7). The corresponding diagrams for all four migration stages under both TIM assignments are reported in Appendix D.2. Red occupation bands indicate trains delayed by more than 60 s. In both TIM assignments, delays propagate from the disturbed area around AMPO towards ALMM and remain visible for several subsequent train movements. The worst-case assignment produces a more persistent delayed pattern, consistent with the higher recovery indicators reported above.

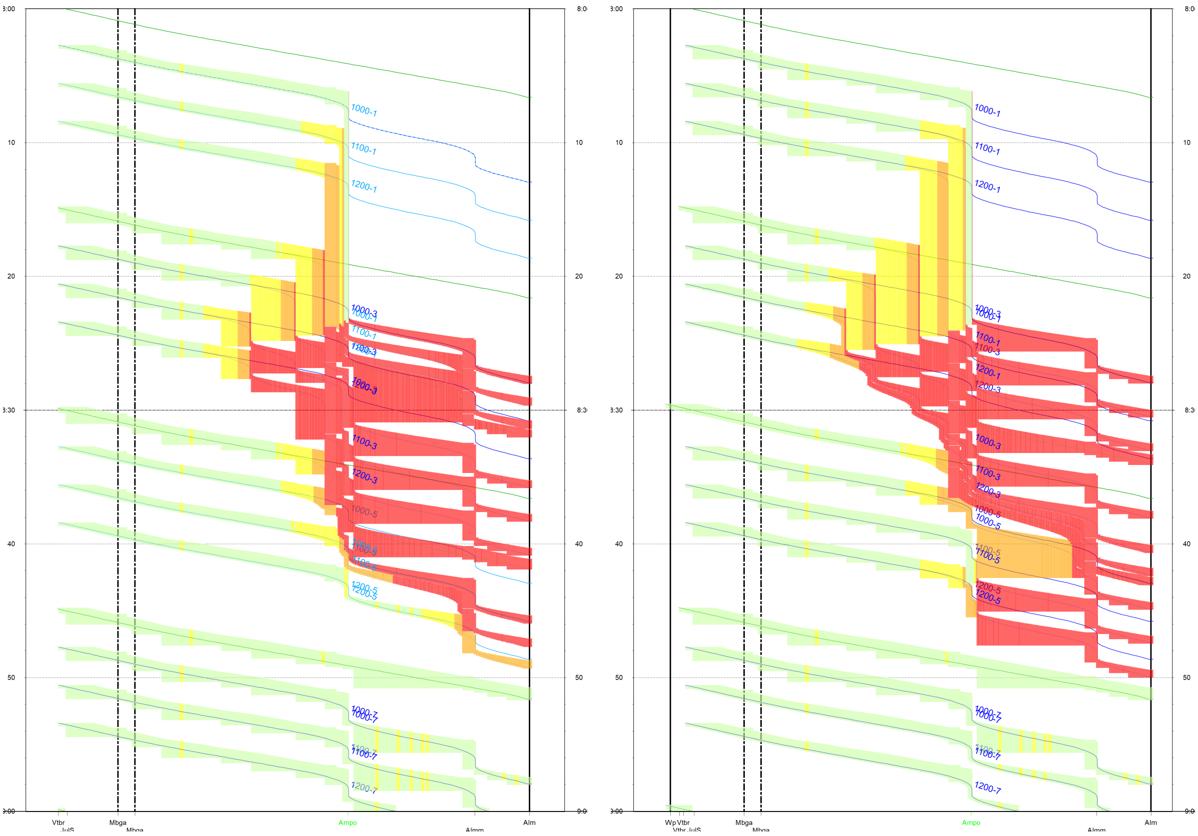

(a) Best-case TIM assignment. **(b)** Worst-case TIM assignment.

Figure 7.7: RailSys time-distance diagrams for the T1 stage 2 disturbance under the best-case and worst-case TIM assignments, with red occupation bands indicating trains delayed by more than 60 s.

The early-stage recovery penalty points to an operational role of physical detection that is not represented by the nominal cost-capacity objective. The long open-line section between AMPO and ALMM is not binding under nominal UIC 406 compression and is therefore left without optional subdivision in the selected T1 layout. Under disturbed operation, however, this long TTD section appears to become operationally relevant while a substantial share of movements remains dependent on TTD level release. An ex-post RailSys diagnostic further examines this mechanism. In this diagnostic, traffic, TIM assignment, and all other infrastructure conditions are held fixed. Retaining one or two additional TTD boundaries within the long AMPO-ALMM section restores the evaluated low-TIM disturbance recovery approximately to the B1 baseline level, while nominal UIC 406 capacity consumption remains unchanged. The diagnostic therefore shows that a boundary can have limited value for the binding nominal headway while still contributing to disturbance recovery. Nominal capacity value and operational stability value are consequently distinct spatial properties of residual physical detection, motivating further consideration of an operational stability safeguard in Chapter 8.

8 Discussion

This chapter interprets the main findings of Chapter 7 and discusses their implications and limitations for TTD layout design under ETCS HTD.

8.1. Interpretation of the main findings

The spatial experiment shows that boundary location matters independently of boundary count. Optimising the locations of 44 boundaries reduces the analytical expected line headway from 155.15 s to 151.98 s, while comparable performance within the predefined equivalence tolerance is obtained with 26 boundaries. A TTD boundary therefore has nominal capacity value when it subdivides a binding or near-binding slow track-releasing section. This finding is consistent with the distinction between critical and stretchable zones in the equi-block framework (Quaglietta, 2014). The present study provides evidence that the same spatial principle applies to residual TTD under ETCS HTD, while also showing that a location that is stretchable from a nominal-capacity perspective may still contribute to disturbance recovery. Consequently, nominally non-critical TTD sections should not be assumed to be operationally irrelevant.

The T1-T2 comparison shows that infrastructure inheritance mainly changes the cost of attaining comparable nominal performance. The selected brownfield and greenfield alternatives have nearly identical maximum expected line headways, while discounted cost is approximately € 2.03 M for T1 and € 2.82 M for T2. Useful inherited boundaries therefore constitute an economic migration asset. The retrofit-cost sensitivity further shows that TIM deployment and TTD adaptation are coupled decisions: changing the economics of onboard deployment changes the selected infrastructure trajectory. This extends previous HTD simulation studies that vary train-integrity availability for predefined infrastructure configurations (Jansen, 2019; Knutsen et al., 2024) by treating the infrastructure trajectory itself as endogenous.

Traffic composition changes the spatial value of TTD as well. Under heterogeneous traffic, cumulative leader-follower running time differences concentrate useful boundaries around specific corridor locations, whereas identical IC services in T3 distribute this value more evenly. Existing HTD studies mainly evaluate TIM deployment, VSS length, and predefined infrastructure configurations (Aoun, 2024; Knutsen et al., 2024). The present results add that the spatial value of residual TTD is itself traffic-dependent and decreases from a nominal capacity perspective as more movements use VSS-level release. Residual TTD optimisation is therefore most consequential during migration, before VSS-level release becomes dominant.

8.2. Nominal capacity and operational stability

The RailSys evaluation shows that nominal capacity and operational stability are distinct design criteria. The selected T1 layouts provide less capacity margin and slower disturbance recovery than B1 during the early migration stages. With traffic, TIM assignment, and all other infrastructure conditions held fixed, retaining one or two additional TTD boundaries within the long AMPO-ALMM section restores approximately baseline-level recovery without changing nominal UIC 406 capacity consumption. Locations that are critical for nominal capacity and disturbance recovery therefore need not coincide. This does not contradict previous nominal HTD capacity findings, but identifies an additional spatial criterion that becomes relevant when TTD boundaries are actively removed during migration. The finding is limited to T1 and the evaluated deterministic disturbance and therefore does not establish a general maximum TTD-section length, reinforcing the need to distinguish the comparative validity of the optimisation framework from the broader operational validity of its selected layouts.

8.3. Methodological limitations and layout selection

Computational verification establishes that the implemented MILP represents the analytical formulation, while comparison with RailSys supports the use of the headway model for comparative layout screening. Further limitations are the single corridor, deterministic train profiles, aggregate representation of partial TIM deployment through ρ, T1-only disturbance evaluation, and planning-level cost assumptions.

A further approximation concerns the relation between infrastructure release and operational block separation. Standard blocking time theory evaluates headway over operational block sections, whereas the analytical model evaluates the applicable VSS or TTD release section so that TTD placement directly affects the capacity objective. This is optimistic because a release boundary is not necessarily a movement authority boundary. A more operationally faithful formulation could retain the actual fixed and virtual block sections for the blocking time calculation while using the selected VSS and TTD layout to determine their release conditions. TTD boundaries required for detection but not corresponding to operational block boundaries would remain represented in the engineering and lifecycle-cost model, while affecting capacity only through operational block release. RailSys partly addresses this approximation by evaluating the selected layouts with the actual operational block configuration.

The MCA-selected alternative is likewise a decision-support projection rather than a uniquely preferred solution. Frontier-specific normalisation means that identical weights do not represent the same absolute euro-headway trade-off across T1-T3. In practical applications, explicit capacity and operational stability constraints within a minimum-cost formulation may therefore be more transparent when stakeholder requirements are known.

8.4. Implications and future research

For infrastructure managers, the findings suggest a four-step migration process: establish the mandatory residual TTD layer, value optional boundaries under the expected traffic mix and TIM trajectory, compare cost and nominal capacity over the complete migration horizon, and subject proposed removals to an operational stability screen. An optional boundary should therefore be removed only when it is not required by engineering constraints, provides insufficient lifecycle-adjusted nominal capacity value, and does not materially support disturbance recovery. Inherited TTD is consequently better viewed as a migration asset whose value is reassessed over time than as legacy infrastructure to be removed as early as possible.

The findings motivate three hypotheses for further research: a corridor-specific deployment level may exist beyond which optional TTD subdivision has negligible nominal capacity value, the locations with greatest nominal capacity and disturbance-recovery value may systematically differ, and the preferred timing of TTD removal may depend on the cost and timing of TIM deployment rather than only its final target. These hypotheses should be tested across additional corridors, traffic patterns, and disturbance scenarios.

Further work should also examine discrete trainset-level deployment, temporary integrity unavailability, and the direct representation of operational stability within the optimisation. In particular, the disturbance results motivate investigating whether a maximum TTD-section length should be imposed as an operational-stability safeguard during migration, particularly at low TIM deployment levels. The present study does not establish a general threshold, so such a constraint would require calibration across different corridors, traffic conditions, TIM assignments, and disturbance scenarios before being incorporated into the optimisation. A further methodological extension should couple the VSS and TTD release hierarchy explicitly to the operational block hierarchy, calculating blocking time over the actual fixed and virtual block sections while using the selected release layout to determine when these blocks can be released. This could ultimately be extended towards joint optimisation of TTD placement, VSS spacing, and usable movement authority boundaries.

9 Conclusion

This thesis addressed how TTD boundary layouts under ETCS HTD can be designed to balance capacity, infrastructure cost, and migration feasibility under mixed-traffic conditions. Existing block-layout methods and HTD studies were found to optimise conventional or virtual block structures or evaluate predefined TTD layouts, leaving the staged adaptation of residual TTD during partial TIM deployment unresolved. The engineering requirements were therefore translated into mandatory, forbidden, and optional TTD boundary positions that define the admissible design domain, while the operational block hierarchy determines where movement authorities may terminate.

Within this domain, the developed analytical model shows that the nominal capacity value of a TTD boundary depends on its location, traffic composition, and TIM deployment rather than on boundary count alone. The T1-T2 comparison shows that infrastructure inheritance mainly affects the cost of achieving comparable performance: the selected brownfield and greenfield alternatives reach maximum expected line headways of 162.18 s and 162.58 s, respectively, while discounted cost increases from approximately €2.03 M to €2.82 M when the layout must be built from a greenfield starting condition. The T2-T3 comparison further shows that traffic composition changes the spatial value of residual detection. Under heterogeneous traffic, useful boundaries concentrate more strongly around specific corridor locations, whereas the homogeneous T3 layout distributes them more evenly along the corridor. Increasing TIM deployment further reduces dependence on TTD-level release and therefore the nominal value of optional subdivision.

The staged bi-objective optimisation jointly selects TTD layouts and annual TIM deployment. For the primary T1 case, the selected trajectory retains 21, 20, 16, and 12 boundaries over four stages and converges to the 12 mandatory positions, with a maximum expected line headway of 162.18 s and discounted cost of approximately €2.03 M. Relative to retaining all 44 inherited boundaries under the same TIM trajectory, this reduces modelled cost by approximately €2.58 M, or 56 %, while increasing maximum expected line headway by 7.67 s, approximately 5 %. The spatial experiment further confirms that comparable nominal performance can be obtained with substantially fewer boundaries when they are located effectively.

Microscopic evaluation shows, however, that nominal cost-capacity performance does not guarantee equivalent operational stability. The selected T1 layouts recover more slowly than the ETCS L2 fixed block baseline during the low-TIM stages, while retaining one or two additional boundaries in the long AMPO-ALMM TTD section restores approximately baseline-level recovery without changing nominal UIC 406 capacity consumption. The resulting design guideline is therefore to assess optional TTD removal sequentially against engineering requirements, lifecycle-adjusted nominal capacity value, and operational stability.

Taken together, the results show that residual TTD under ETCS HTD should be designed as a staged, location-dependent migration resource coupled to TIM deployment. Engineering requirements define the mandatory residual layer, optimisation identifies where optional boundaries remain justified by lifecycle cost and nominal capacity, and operational evaluation determines whether their removal remains acceptable from a disturbance-recovery perspective.

References

Anthropic. (2026). Claude. https://claude.ai

Aoun, J. (2024). *Rail to digital automated up to autonomous train operation, d37.3 – requirements specification to fa1 wp8/9* (D37.3, R2DATO Work Package 37) (Deliverable R2D-WP37-D-PRO-001-03). Europe's Rail Joint Undertaking. https://rail-research.europa.eu/pages/fp2-r2dato/deliverables

Aoun, J. (2025). *D37.2 - methodology for optimised placement of (virtual) blocks* (Project Deliverable No. D37.2). Europe's Rail.

Aoun, J. (2026a). *D37.1 - standardisation and deployment* (Project Deliverable No. D37.1). Europe's Rail.

Aoun, J. (2026b). *Rail to digital automated up to autonomous train operation, d37.4 – determining etcs htd/hl3 capacity impact analysis using simulations* (D37.4, R2DATO Work Package 37) (Deliverable R2D-WP37-D-PRO-017-04). Europe's Rail Joint Undertaking. https://rail-research.europa.eu/pages/fp2-r2dato/deliverables

DeepL SE. (2026). Deepl translator. https://www.deepl.com/translator

Delft High Performance Computing Centre (DHPC). (2024). DelftBlue Supercomputer (Phase 2).

EEIG, E. U. G. (2022). *Guideline htd engineering* (tech. rep. No. 21E087-3). ERTMS Users Group EEIG. https://ertms.be/wp-content/uploads/simple-file-list/21E087-3_Guideline-HTD-Engineering.pdf

EEIG, E. U. G. (2026). *Ertms/etcs hybrid train detection (htd) principles* (tech. rep. No. 16E0422). ERTMS Users Group EEIG. https://ertms.be/wp-content/uploads/2026/02/16E0422-_HTDclean.pdf

Ehrgott, M. (2005). *Multicriteria optimization* (2nd). Springer. https://doi.org/10.1007/3-540-27659-9

Engels, S., Peham, T., Przigoda, J., Przigoda, N., & Wille, R. (2025). Design tasks and their complexity for the european train control system with hybrid train detection. *EURO Journal on Transportation and Logistics*, 100161. https://www.cda.cit.tum.de/files/eda/2025_eurojtl_etcs_design_tasks_and_complexity.pdf

Engels, S., Peham, T., & Wille, R. (2023). A symbolic design method for etcs hybrid level 3 at different degrees of accuracy. *23rd Symposium on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2023)*, *115*, 6:1–6:17. https://doi.org/10.4230/OASIcs.ATMOS.2023.6

Engels, S., & Wille, R. (2024). Towards an optimization pipeline for the design of train control systems with hybrid train detection. *24th Symposium on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2024)*. https://doi.org/10.4230/OASIcs.ATMOS.2024.12

European Commission. (2014). Commission regulation (eu) no 1301/2014 on the technical specification for interoperability relating to the energy subsystem of the rail system in the union [Consolidated version available via ERA]. https://www.era.europa.eu/domains/technical-specifications-interoperability/energy-tsi_en

Commission Implementing Regulation (EU) 2023/1695 on the technical specification for interoperability relating to the control-command and signalling subsystems of the rail system in the European Union, 2023, https://eur-lex.europa.eu/eli/reg_impl/2023/1695

European Commission. (2025). What is ertms and how does it work? https://transport.ec.europa.eu/transport-modes/rail/ertms/what-ertms-and-how-does-it-work_en

European Commission, Directorate-General for Mobility and Transport. (2026). *Third ERTMS work plan: ERTMS deployment is progressing, but more must be done.* https://transport.ec.europa.eu/news-events/news/third-ertms-work-plan-ertms-deployment-progressing-more-must-be-done-2026-02-23_en

European Union Agency for Railways. (2022). *Harmonised ERTMS marker boards* (tech. rep.). European Union Agency for Railways. https://www.era.europa.eu/system/files/2022-11/Harmonised%5C%20ERTMS%5C%20Marker%5C%20Boards.pdf

European Union Agency for Railways. (2023). *Braking curves*. https://www.era.europa.eu/domains/european-rail-traffic-management-system/braking-curves

European Union Agency for Railways, UNISIG, & ERTMS Users Group EEIG. (2023a, July 5). *ERTMS/ETCS glossary of terms and abbreviations* (tech. rep. No. SUBSET-023) (CCS TSI Appendix A mandatory specification). European Union Agency for Railways. https://www.era.europa.eu/system/files/2023-09/index003_-_SUBSET-023_v400.pdf

European Union Agency for Railways, UNISIG, & ERTMS Users Group EEIG. (2023b, July 5). *ERTMS/ETCS system requirements specification: Chapter 2, basic system description* (tech. rep. No. SUBSET-026-2) (CCS TSI Appendix A mandatory specification). European Union Agency for Railways. https://www.era.europa.eu/era-folder/1-ccs-tsi-appendix-mandatory-specifications-etcs-b4-r1-rmr-gsm-r-b1-mr1-frmcs-b0-ato-b1

Flammini, F. (2013). Automatic train protection systems. *Industrial Engineering & Management*, *02*(05). https://doi.org/10.4172/2169-0316.1000120

Hansen, D., Leuschel, M., Körner, P., Krings, S., Naulin, T., Nayeri, N., Schneider, D., & Skowron, F. (2020). Validation and real-life demonstration of etcs hybrid level 3 principles using a formal b model. *International Journal on Software Tools for Technology Transfer*. https://doi.org/10.1007/s10009-020-00551-6

Hansen, I. A., & Pachl, J. (2014). *Railway timetabling & operations*. Eurailpress.

International Union of Railways. (2013). *Uic code 406: Capacity*.

Jansen, J. M. (2019). *Ertms/etcs hybrid level 3: A simulation-based impact assessment for the dutch railway network* [Master's thesis, Delft University of Technology]. https://repository.tudelft.nl/record/uuid:639015a7-d633-4b73-b782-9b9a153eae11

Knutsen, D., Fu, J., & Olsson, N. O. E. (2024). Capacity evaluation of ERTMS/ETCS hybrid level 3 using simulation methods. *Journal of Rail Transport Planning & Management*, *30*, 100449. https://doi.org/10.1016/j.jrtpm.2024.100449

Knutsen, D., Olsson, N. O. E., Fu, J., & Rosberg, T. (2025). Capacity evaluation of ertms/etcs htd and moving block. *Journal of Rail Transport Planning & Management*, *33*, 100506. https://doi.org/10.1016/j.jrtpm.2025.100506

Koopman - van Welij, D. (2026). *Rail to digital automated up to autonomous train operation, d32.5 - dato capacity and impact simulations report* (D32.5, R2DATO Work Package 32) (Deliverable R2D-WP32-D-PRO-019-04). Europe's Rail Joint Undertaking.

Mavrotas, G. (2009). Effective implementation of the ε-constraint method in multi-objective mathematical programming problems. *Applied Mathematics and Computation*, *213*(2), 455–465. https://doi.org/10.1016/j.amc.2009.03.037

OpenAI. (2026). Chatgpt. https://chat.openai.com

Pachl, J. (2021). *Railway signalling principles* [Accessed: 2025-03-08]. Jörn Pachl. http://www.joernpachl.de/eBook%20RSP.pdf

ProRail. (2023). *Ontwerpvoorschrift baan en bovenbouw wissels en kruisingen ovs00056-6.1 (v009)* (tech. rep.). ProRail.

ProRail. (2025a). Ertms: Het digitale spoorplatform. https://www.prorail.nl/programmas/ertms

ProRail. (2025b). *Ontwerp voorschrift ertms ovs60040-2 (v005)* (tech. rep.). ProRail.

ProRail. (2025c). *Programma hoogfrequent spoorvervoer (phs)*. https://www.prorail.nl/programmas/programma-hoogfrequent-spoorvervoer

ProRail. (2025d). Prorail: Investeringen blijven achter, focus op onderhoud. https://www.prorail.nl/nieuws/prorail-investeringen-blijven-achter-focus-op-onderhoud

Quaglietta, E. (2014). A simulation-based approach for the optimal design of signalling block layout in railway networks. *Simulation Modelling Practice and Theory*. https://www.sciencedirect.com/science/article/pii/S1569190X13001664

Rail Management Consultants International GmbH. (2022). *Railsys* (Version 2022). https://rmcon-int.de/en/release-railsys-2022-2/

Scheepmaker, G. M., Willeboordse, H. Y., Hoogenraad, J. H., Luijt, R. S., & Goverde, R. M. P. (2020). Comparing train driving strategies on multiple key performance indicators. *Journal of Rail Transport Planning & Management*, *13*, 100163. https://doi.org/10.1016/j.jrtpm.2019.100163

The Railway Technical Website. (2023). Signalling. https://www.railway-technical.com/signalling/

UNIFE. (2024). *Ertms/etcs levels* (tech. rep.). UNIFE. https://www.ertms.net/wp-content/uploads/2024/03/ERTMS-ETCS-LEVELS_Updated-factsheet_2024.pdf

Vergroesen, R. (2020). *Ertms/etcs hybrid level 3 and ato: A simulation-based capacity impact study for the dutch railway network* [Master's thesis, Delft University of Technology].

Versluis, N. D., Pellegrini, P., Quaglietta, E., Goverde, R. M. P., & Rodriguez, J. (2025). Impact of track discretisation on conflict detection and resolution under etcs with onboard train integrity monitoring. *Journal of Rail Transport Planning & Management*. https://doi.org/10.1016/j.jrtpm.2025.100533

Vignali, V., Cuppi, F., Lantieri, C., Dimola, N., Galasso, T., & Rapagnà, L. (2020). A methodology for the design of sections block length on etcs l2 railway networks. *Journal of Rail Transport Planning & Management*. https://doi.org/10.1016/j.jrtpm.2019.100160

Wagter, W. H. (2019). Block sections around stations under etcs: The effect of the block layout on the headways at stations [s1822497]. https://essay.utwente.nl/80829/

Werkgroep discontovoet. (2025). *Rapport werkgroep discontovoet 2025* (Report). Ministry of Finance. https://www.rwseconomie.nl/documenten/2026/februari/20/rapport-werkgroep-discontovoet-2025

A Additional optimisation data

This appendix provides supplementary computational and optimisation data supporting the case study results of Chapter 7. It documents the computational setup, the generated Pareto frontiers, the initial physical TTD reference layouts, the selected boundary positions, and the TIM retrofit cost sensitivity.

A.1. Computational setup and solution settings

The case study optimisation is implemented in PuLP and solved with Gurobi 12.0 using the objective-equivalent MILP formulation of Appendix B. The T1, T2, and T3 Pareto frontiers were generated on the DelftBlue supercomputer (Delft High Performance Computing Centre (DHPC), 2024), with each cost-level instance solved using a relative MIP optimality-gap tolerance of 1 %.

Table A.1 gives the dimensions of the solved T1 MILP, which is representative of the computational scale of the case study runs. The section-arc variables x^r_{ij} account for most of the binary variables.

Table A.1: Dimensions of the solved MILP for the T1 case study instance.

Dimension	Count
Total variables	41,388
Binary variables	41,366
Layout binaries y^r_p	1,076
Arc-selection binaries x^r_{ij}	38,108
Transition binaries a^r_p, d^r_p	2,152
Level-selection binaries b^t_ℓ	30
Continuous variables (G^r_k, H^t, $\bar{H}$, C)	22
Constraints	123,039

The Pareto frontier is generated by varying the cost cap $C_{\max}$ over a discrete grid spanning the full-retention and mandatory-only cost extremes and applying the two-pass lexicographic procedure of Section 5.5.5 at each cost level. For the TIM retrofit cost sensitivity, the complete T1 frontier is re-solved at €0, €42 500, and €84 420 per train because changing the retrofit cost can alter both the deployment trajectory and the residual physical TTD layout.

The implementation satisfies the computational verification checks defined in Section 5.5.6. Detailed verification, achieved-gap, and formulation-strength diagnostics are reported in Appendix C, in particular Sections C.5 and C.4.4.

A.2. Pareto frontiers by case study type

Figure A.1 collects the three discrete Pareto frontiers reported in Chapter 7. Panel (a) is T1 (brownfield, heterogeneous traffic), panel (b) is T2 (greenfield, heterogeneous traffic) and panel (c) is T3 (greenfield, homogeneous IC-only traffic). In each panel, the MCA-selected alternative is marked in red and the remaining non-dominated alternatives are drawn in grey. All alternatives are non-dominated within the configured 1 % relative MIP optimality gap tolerance.

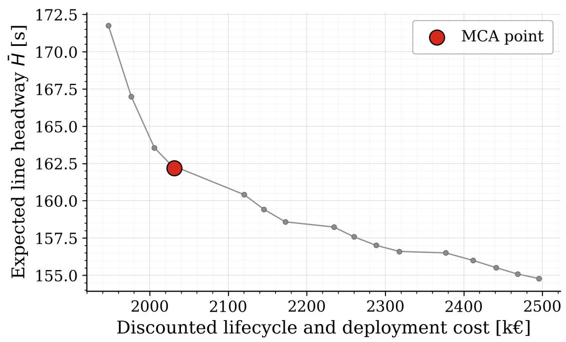

(a) T1: brownfield, heterogeneous traffic

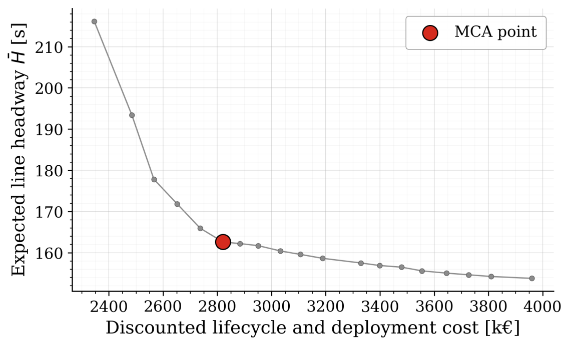

(b) T2: greenfield, heterogeneous traffic

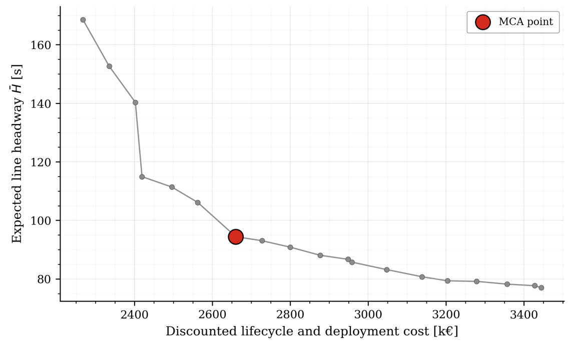

(c) T3: greenfield, homogeneous IC-only traffic

Figure A.1: Pareto frontiers for the three case study types, with the MCA-selected alternative marked in red in each panel.

A.3. Initial physical TTD reference layouts

This section documents the physical TTD data used by the three case study types before the optimisation. T1 starts from the current brownfield layout listed in Table A.2, which contains 44 inherited TTD boundary positions. T2 and T3 start from an empty initial layout ($\hat{y} = 0$). The 12 mandatory boundary positions listed in Table A.3 are then enforced from stage 1 onward by the mandatory-position constraint of the formulation. The comparison between T1 and T2 therefore examines the effect of the inherited detection asset base under the same heterogeneous traffic setting.

Table A.2: Brownfield initial TTD layout used for T1 (44 boundaries).

Brownfield boundary positions [km]
0.416, 0.909, 1.399, 1.878, 1.904, 1.987, 2.080, 2.206, 2.308, 2.354, 2.442, 2.911, 3.374, 3.834, 4.294, 4.756, 5.214, 5.672, 6.142, 6.594, 7.056, 7.521, 7.974, 8.436, 8.958, 9.305, 9.521, 9.971, 10.420, 10.880, 11.330, 11.780, 12.230, 12.696, 12.933, 13.251, 13.438, 13.638, 13.838, 14.038, 14.238, 14.438, 14.638, 14.685

Table A.3: Mandatory TTD boundary positions enforced from stage 1 onward, applied to T1, T2, and T3.

Mandatory boundary positions [km]
0.416, 1.904, 1.987, 2.080, 2.206, 2.308, 2.354, 2.442, 3.374, 8.436, 12.696, 14.685

A.4. Selected physical TTD layouts by case study type

This section lists the physical TTD boundary positions of the MCA-selected layout for every case study type reported in Chapter 7. Positions are given in kilometres from the corridor origin. Mandatory boundaries are included in every layout. For each type, the active physical TTD boundaries per migration stage are given as a table and shown as a layout figure. The tables report the boundaries after the lexicographic Pass B cost clean-up step, so within-stage year to year layout variation is absent by construction.

A.4.1. T1 (brownfield, heterogeneous traffic)

The T1 MCA-selected layout achieves $\bar{H} = 162.18\,\text{s}$ at a discounted cost of €2,031,448 with per-stage active boundary counts $\{21, 20, 16, 12\}$. The migration trajectory progressively removes optional boundaries as TIM deployment increases and converges to the 12 mandatory positions in the final stage. In the early stages the selected optional boundaries cluster around the approach to AMPO, the station areas around AMPO and ALMM, and the downstream side of ALMM.

Table A.4: T1 MCA-selected layout, brownfield initial layout with heterogeneous traffic (discounted cost €2,031,448, $\bar{H} = 162.18\,\text{s}$).

Stage	End of stage ρ	Active boundary positions [km]
1	0.15	0.416, 1.904, 1.987, 2.080, 2.206, 2.308, 2.354, 2.442, 3.374, 5.214, 6.142, 7.521, 8.436, 8.958, 9.305, 12.696, 13.030, 13.438, 14.038, 14.438, 14.685 (21)
2	0.30	0.416, 1.904, 1.987, 2.080, 2.206, 2.308, 2.354, 2.442, 3.374, 5.214, 6.142, 7.521, 8.436, 8.958, 9.305, 12.696, 13.030, 13.438, 14.038, 14.685 (20)
3	0.65	0.416, 1.904, 1.987, 2.080, 2.206, 2.308, 2.354, 2.442, 3.374, 6.142, 8.436, 9.305, 12.696, 13.438, 14.038, 14.685 (16)
4	0.95	0.416, 1.904, 1.987, 2.080, 2.206, 2.308, 2.354, 2.442, 3.374, 8.436, 12.696, 14.685 (12, mandatory only)

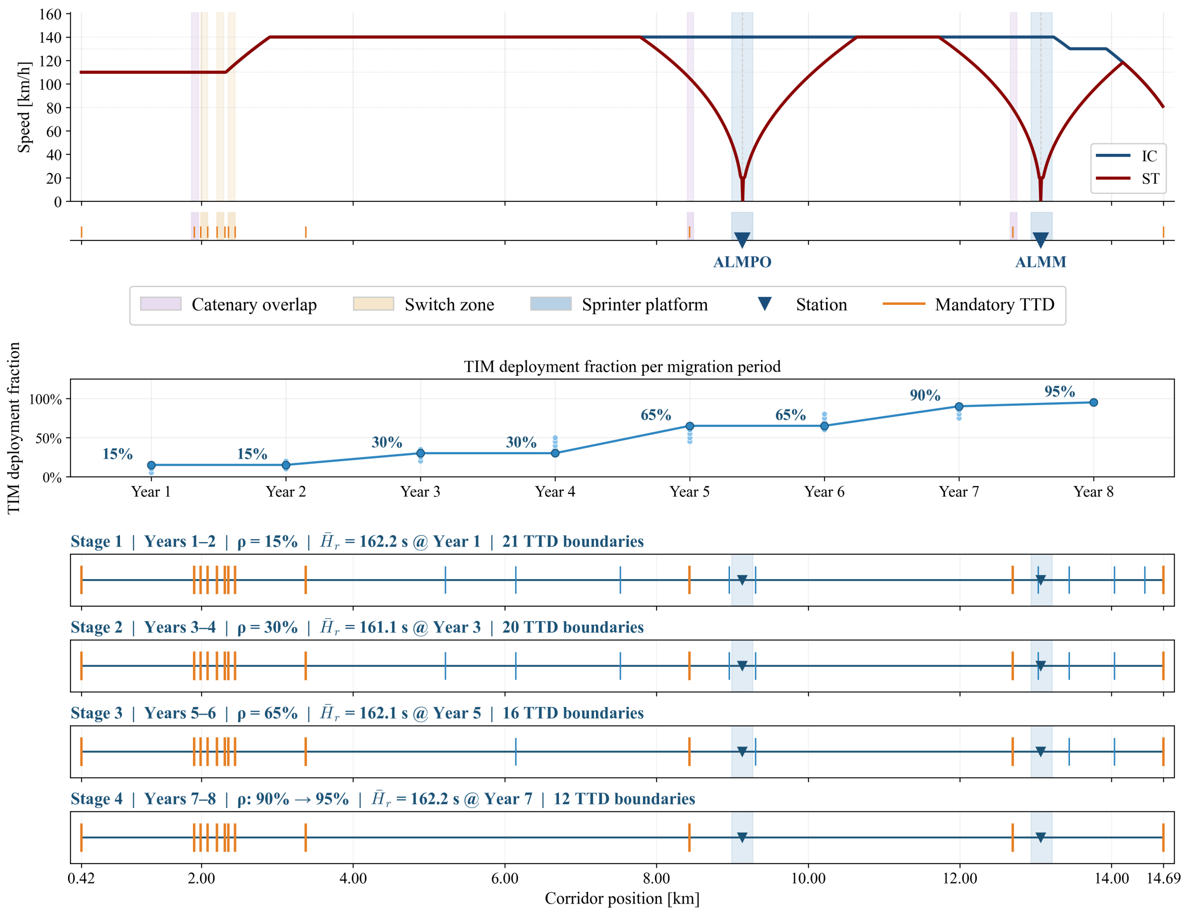

Figure A.2: Selected per-stage physical TTD layout for T1: brownfield initial layout with heterogeneous IC/ST traffic.

A.4.2. T2 (greenfield, heterogeneous traffic)

T2 replaces the brownfield initial layout of T1 with an empty initial layout while keeping the same heterogeneous IC/ST traffic setting. The MCA-selected layout achieves $\bar{H} = 162.58\,\text{s}$ at a discounted cost of €2,820,816 with per-stage active boundary counts $\{19, 18, 16, 12\}$. The selected headway is essentially unchanged compared to T1 (162.58 s versus 162.18 s), while the discounted cost is about €790 k higher because every optional boundary must be newly installed rather than retained from the inherited layout. The selected optional boundaries concentrate near AMPO (8.966 km and 9.316 km) and ALMM (13.080 km, 13.782 km, 14.334 km).

Table A.5: T2 MCA-selected layout, greenfield initial layout with heterogeneous traffic (discounted cost €2,820,816, $\bar{H} = 162.58\,\text{s}$).

Stage	End of stage ρ	Active boundary positions [km]
1	0.15	0.416, 1.904, 1.987, 2.080, 2.206, 2.308, 2.354, 2.442, 3.374, 5.268, 6.863, 8.436, 8.966, 9.316, 12.696, 13.080, 13.782, 14.334, 14.685 (19)
2	0.35	0.416, 1.904, 1.987, 2.080, 2.206, 2.308, 2.354, 2.442, 3.374, 5.268, 6.863, 8.436, 8.966, 9.316, 12.696, 13.080, 13.782, 14.685 (18)
3	0.65	0.416, 1.904, 1.987, 2.080, 2.206, 2.308, 2.354, 2.442, 3.374, 5.268, 8.436, 8.966, 9.316, 12.696, 13.080, 14.685 (16)
4	0.95	0.416, 1.904, 1.987, 2.080, 2.206, 2.308, 2.354, 2.442, 3.374, 8.436, 12.696, 14.685 (12, mandatory only)

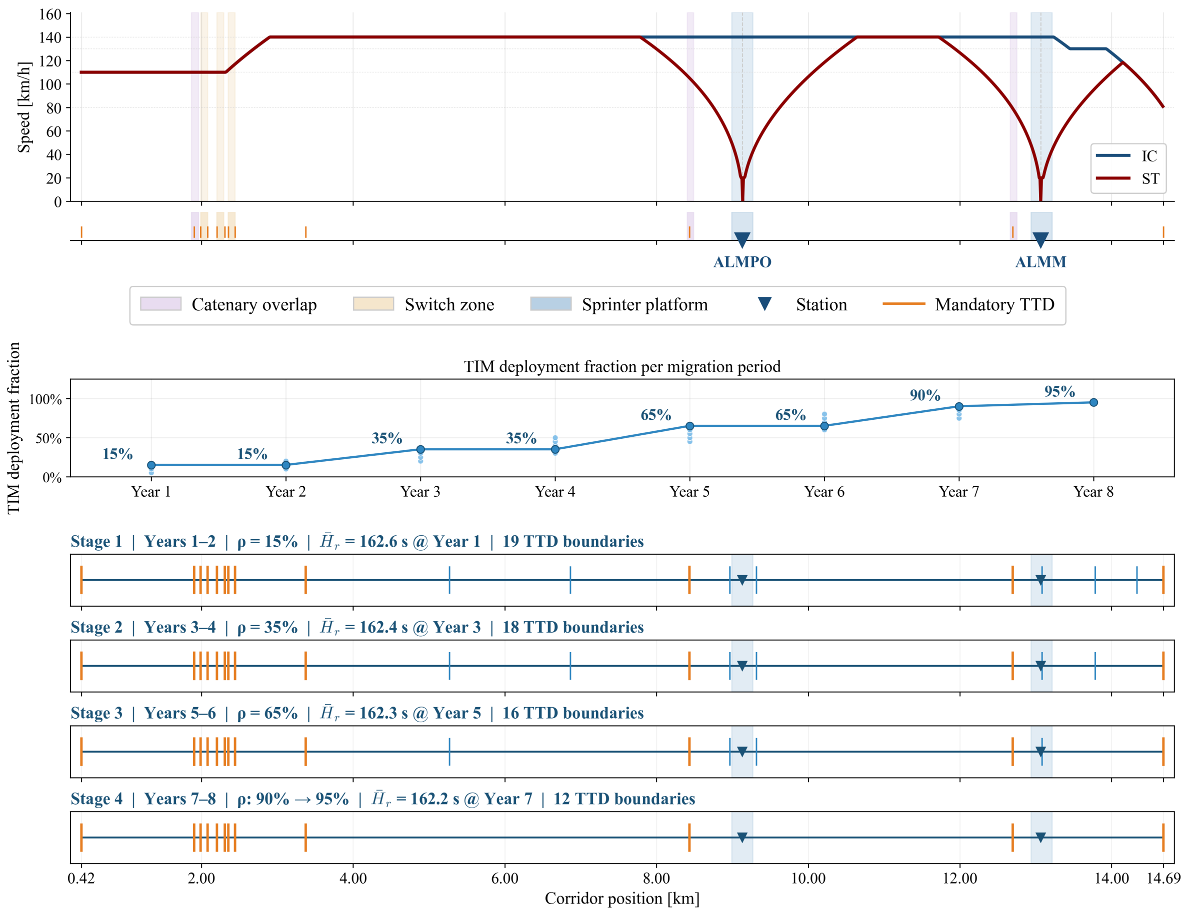

Figure A.3: Selected per-stage physical TTD layout for T2: greenfield initial layout with heterogeneous IC/ST traffic.

A.4.3. T3 (greenfield, homogeneous IC-only traffic)

T3 keeps the greenfield initial layout of T2 and switches the traffic to homogeneous IC-only. The MCA-selected layout achieves $\bar{H} = 94.41\,\text{s}$ at a discounted cost of €2,660,346 with per-stage active boundary counts $\{17, 17, 14, 12\}$. Under the same greenfield starting condition and common infrastructure cost structure as T2, the comparison examines the influence of traffic composition. The T3 layout distributes the optional boundaries more evenly along the corridor and includes open-line positions in the AMPO-ALMM section (9.765 km and 11.115 km) that do not appear in T2.

Table A.6: T3 MCA-selected layout, greenfield initial layout with homogeneous IC-only traffic (discounted cost €2,660,346, $\bar{H} = 94.41\,\text{s}$).

Stage	End of stage ρ	Active boundary positions [km]
1	0.15	0.416, 1.904, 1.987, 2.080, 2.206, 2.308, 2.354, 2.442, 3.374, 5.268, 6.863, 8.436, 9.765, 11.115, 12.696, 13.832, 14.685 (17)
2	0.30	0.416, 1.904, 1.987, 2.080, 2.206, 2.308, 2.354, 2.442, 3.374, 5.268, 6.863, 8.436, 9.765, 11.115, 12.696, 13.832, 14.685 (17)
3	0.65	0.416, 1.904, 1.987, 2.080, 2.206, 2.308, 2.354, 2.442, 3.374, 5.268, 8.436, 11.115, 12.696, 14.685 (14)
4	0.95	0.416, 1.904, 1.987, 2.080, 2.206, 2.308, 2.354, 2.442, 3.374, 8.436, 12.696, 14.685 (12, mandatory only)

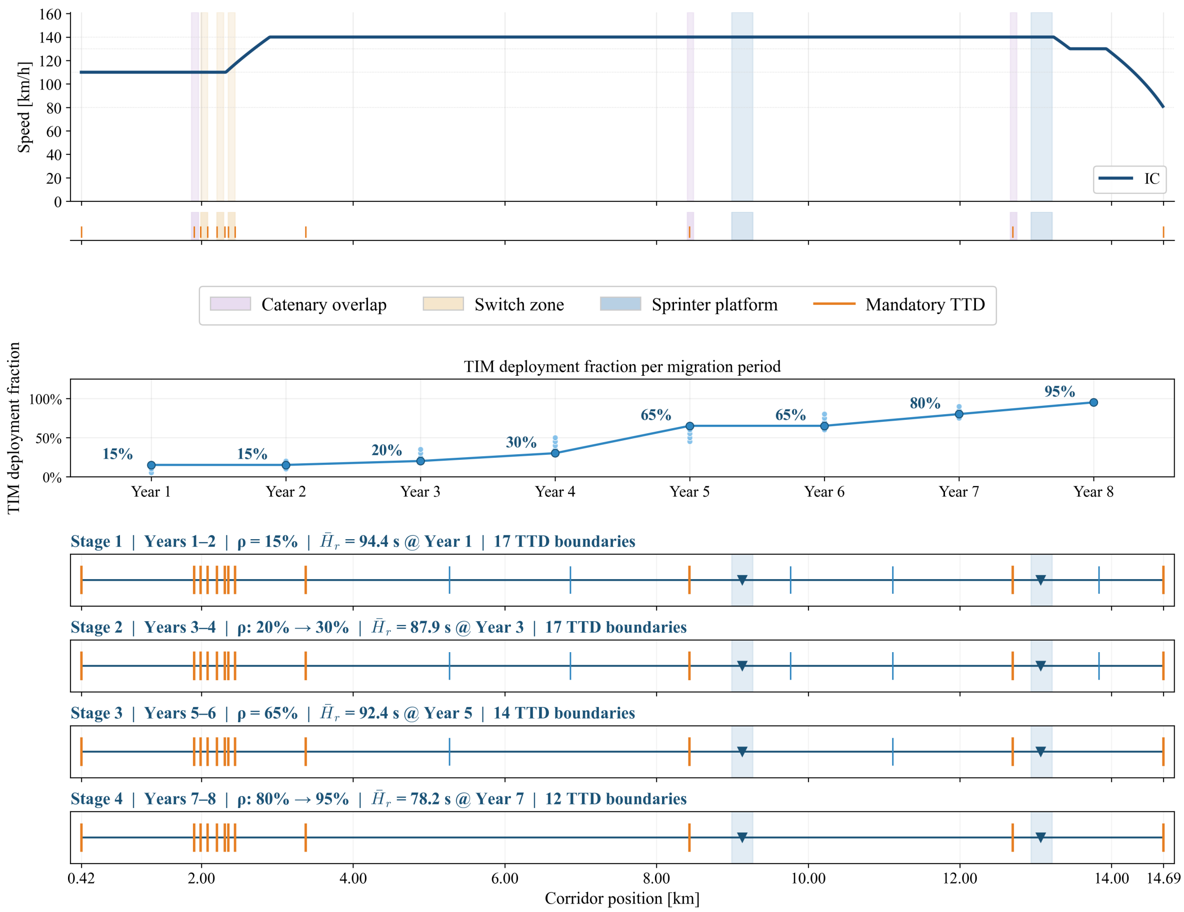

Figure A.4: Selected per-stage physical TTD layout for T3: greenfield initial layout with homogeneous IC-only traffic.

A.5. T1 TIM retrofit cost sensitivity

The T1 Pareto sweep was repeated for retrofit costs of €0, €42 500, and the €84 420 baseline to examine how the joint cost-capacity trade-off between fleet deployment and residual physical detection responds to the retrofit cost assumption. Table A.7 reports the resulting MCA-selected objective values. The shift in the MCA-selected point also produces a different selected TTD layout: Figure A.5 shows the selected layouts for the €0 and €42 500 settings, while the €84 420 baseline layout is the one already shown in Figure A.2. A lower retrofit cost brings TIM deployment forward and retains denser residual physical detection during the early migration stages.

Table A.7: T1 MCA-selected objective results under three TIM retrofit cost settings.

c_{train} [€]	$\bar{H}$ [s]	Cost [M€]
0	156.60	1.74
42 500	158.59	1.89
84 420	162.18	2.03

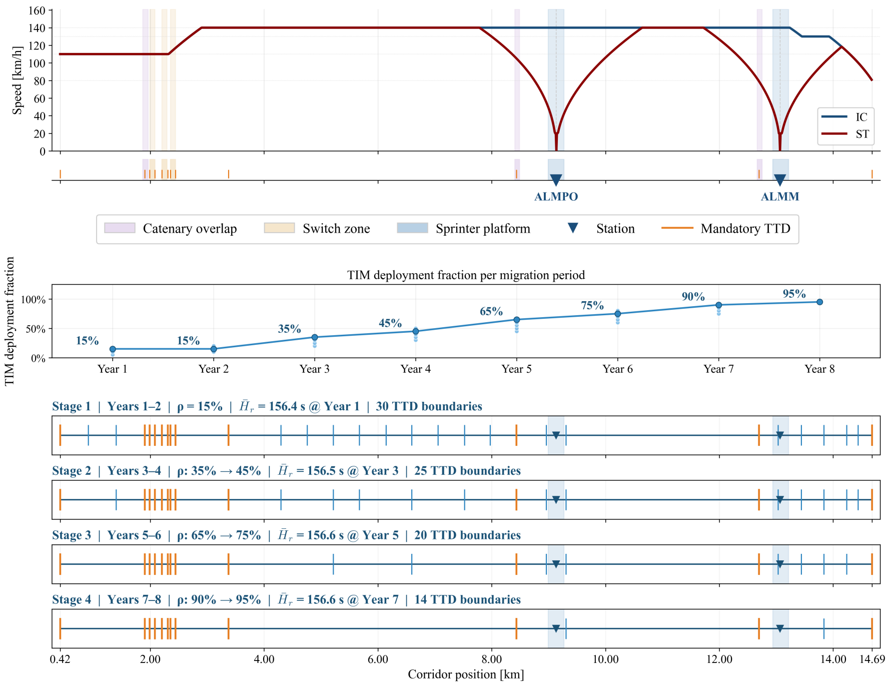

(a) $c_{\text{train}} = 0$

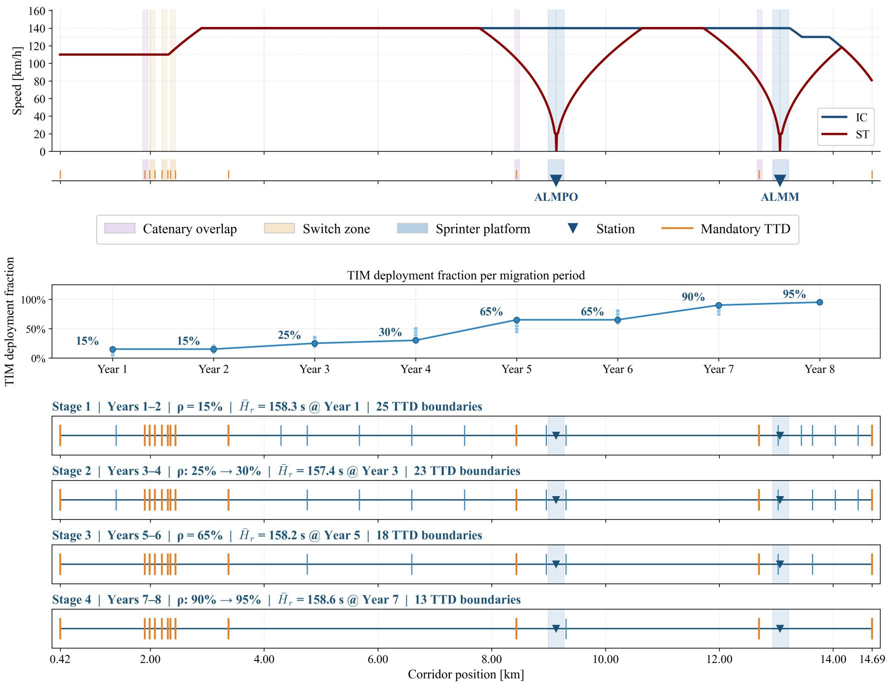

(b) $c_{\text{train}} = 42\,500$

Figure A.5: Selected per-stage physical TTD layouts for T1 under the €0 and €42 500 TIM retrofit cost settings; the €84 420 baseline layout is shown in Figure A.2.

B MILP reformulation

The section-arc representation of the MIP in Section 5.5.4 is already linear, so no further reformulation of the physical layout structure is required. Two nonlinearities remain: the positive-part operators in the infrastructure transition cost and the product between the selected TIM ratio and the stage-dependent slow track-releasing headway. This appendix gives the objective-equivalent MILP solved in the computational implementation, developing the formulation, the big-M calibration, and the certified lower bound. The numerical verification of the implemented model is reported separately in Appendix C. Only notation additional to Chapter 5 is introduced below.

B.1. MILP formulation

Additional notation. Table B.1 lists the symbols additional to the model notation of Section 5.5.3. The layout, arc, deployment, and headway variables (y_p^r, x_{ij}^r, b_ℓ^t, G_k^r, H^t, $\bar{H}$) are unchanged and are not repeated here.

Table B.1: Notation additional to Chapter 5 used in the MILP reformulation.

Symbol	Meaning
$\mathcal{B}^{\text{feas}} = \mathcal{B} \setminus \mathcal{F}$	Feasible boundary positions declared in the production MILP (forbidden positions carry no variable)
$a_p^r \in \{0,1\}$	Boundary added at position p on entry to stage r
$d_p^r \in \{0,1\}$	Boundary removed at position p on entry to stage r
$C \geq 0$	Total discounted lifecycle and deployment cost, $C = f_2$
$M_{G,k}$	Conservative scenario-specific bound on the slow track-releasing line headway, used for deactivation
$M_{t,\ell}^H$	Expected line headway big-M for year t and deployment level ℓ
$\hat{H}$	Pass-A incumbent headway used as the Pass-B cap

Objectives and cost definition. For compactness, write

$$\rho_t = \sum_{\ell \in L_t} \ell \, b_\ell^t, \qquad y_p^{-1} = \hat{y}_p, \qquad \rho_{-1} = 0.$$

The total discounted lifecycle and deployment cost is the linear expression

$$\begin{aligned} C = & \sum_{r \in \mathcal{R}} \alpha_{t_r^*} \sum_{p \in \mathcal{B}^{\text{feas}}} \left(c_{\text{place}} \, a_p^r + c_{\text{remove}} \, d_p^r \right) \\ & + \sum_{t \in \mathcal{T}} \alpha_t \sum_{p \in \mathcal{B}^{\text{feas}}} c_{\text{maint}} \, y_p^{r(t)} + \beta_{\text{tail}} \, c_{\text{maint}} \sum_{p \in \mathcal{B}^{\text{feas}}} y_p^{R-1} \\ & + \sum_{t \in \mathcal{T}} \alpha_t \left(\rho_t - \rho_{t-1} \right) N_{\text{train}} \, c_{\text{train}}. \end{aligned} \tag{B.1}$$

The two-objective problem is solved by the ε-constraint method as two single-objective passes over the common constraints below.

Pass A. min $\bar{H}$ subject to $C \leq C_{\max}$ and the common constraints.

Pass B. min C subject to $\bar{H} \leq \hat{H}$ and the common constraints, where $\hat{H}$ is the Pass-A incumbent headway at the same cost level.

Common constraints. Both passes share the following constraints.

$$\bar{H} \geq H^t \qquad \forall\, t \in \mathcal{T} \tag{M1}$$

$$H^t \geq \sum_{k\in\mathcal{K}^+} \pi_k \big[\ell F_k + (1-\ell) G_k^{r(t)}\big] - M_{t,\ell}^H (1 - b_\ell^t) \qquad \forall\, t \in \mathcal{T},\ \ell \in L_t \tag{M2}$$

$$G_k^r \geq h_{\mathrm{TTD},k,ij}\, x_{ij}^r \qquad \forall\, k \in \mathcal{K}^+,\ r \in \mathcal{R},\ (i,j) \in A_q,\ q \in Q \tag{M3}$$

$$G_k^r \geq L_k \qquad \forall\, k \in \mathcal{K}^+,\ r \in \mathcal{R} \tag{M4}$$

$$\sum_{j:\,(m_q,j)\in A_q} x_{m_q j}^r = 1 \qquad \forall\, q \in Q,\ r \in \mathcal{R} \tag{M5}$$

$$\sum_{i:\,(i,n_q)\in A_q} x_{i\, n_q}^r = 1 \qquad \forall\, q \in Q,\ r \in \mathcal{R} \tag{M6}$$

$$\sum_{i:\,(i,p)\in A_q} x_{ip}^r = \sum_{j:\,(p,j)\in A_q} x_{pj}^r \qquad \forall\, p \in V_q,\ q \in Q,\ r \in \mathcal{R} \tag{M7}$$

$$y_p^r = \sum_{i:\,(i,p)\in A_q} x_{ip}^r \qquad \forall\, p \in V_q,\ q \in Q,\ r \in \mathcal{R} \tag{M8}$$

$$y_p^r = 1 \qquad \forall\, p \in \mathcal{M},\ r \in \mathcal{R} \tag{M9}$$

$$\sum_{\ell\in L_t} b_\ell^t = 1 \qquad \forall\, t \in \mathcal{T} \tag{M10}$$

$$\rho_t \geq \rho_{t-1} \qquad \forall\, t \in \mathcal{T} \setminus \{0\} \tag{M11}$$

The infrastructure transition cost uses the boundary-change indicators, linearised for all $p \in \mathcal{B}^{\mathrm{feas}}$ and $r \in \mathcal{R}$ with $y_p^{-1} = \hat{y}_p$ by

$$a_p^r \geq y_p^r - y_p^{r-1}, \qquad a_p^r \leq y_p^r, \qquad a_p^r \leq 1 - y_p^{r-1}, \tag{M12a}$$

$$d_p^r \geq y_p^{r-1} - y_p^r, \qquad d_p^r \leq y_p^{r-1}, \qquad d_p^r \leq 1 - y_p^r. \tag{M12b}$$

Because the first-stage predecessor y_p^{-1} is the fixed brownfield seed $\hat{y}_p$, the same six inequalities cover the initial and all subsequent transitions.

The variable domains are

$$
\begin{aligned}
& y_p^r,\ a_p^r,\ d_p^r \in \{0,1\} && \forall\, p \in \mathcal{B}^{\mathrm{feas}},\ r \in \mathcal{R},\\
& x_{ij}^r \in \{0,1\} && \forall\, (i,j) \in A_q,\ q \in Q,\ r \in \mathcal{R},\\
& b_\ell^t \in \{0,1\} && \forall\, \ell \in L_t,\ t \in \mathcal{T},\\
& G_k^r \geq 0 && \forall\, k \in \mathcal{K}^+,\ r \in \mathcal{R},\\
& H^t \geq 0 && \forall\, t \in \mathcal{T},\\
& \bar{H} \geq 0, \quad C \geq 0.
\end{aligned}
$$

Relation to the MIP formulation. The minimax constraint M1, the section-headway and layout constraints M3–M9, and the deployment constraints M10–M11 are identical to those introduced in Section 5.5.4. The MILP differs in only two respects. First, the binary indicators a_p^r and d_p^r of (B.3) replace the positive-part transition terms in the cost objective. Second, the bilinear expected-headway constraint is replaced by one linear inequality (M2) for each admissible deployment level. Because exactly one b_ℓ^t is active in each year, the inequality for the selected level remains active while the others are relaxed by $M_{t,\ell}^H$. Forbidden positions are represented structurally rather than by a fixed-zero constraint, because they are excluded from $\mathcal{B}^{\mathrm{feas}}$ and from the interval vertex sets.

The resulting formulation is objective-equivalent to the MIP. The epigraph variables G_k^r and H^t may retain slack in a stage or year that does not determine $\bar{H}$, but they can be tightened to their defining values without changing the discrete decisions or the objective value.

B.2. Big-M calibration

The single big-M family in the formulation is the per-level constant $M^H_{t,\ell}$ that deactivates the expected headway inequality (M2). It is built from a conservative scenario-specific bound $M_{G,k}$ on the slow track-releasing line headway.

For each scenario k and each mandatory interval $q \in Q$ with endpoints m_q and n_q, no active TTD section can span a mandatory boundary, so every feasible section lies within one interval and is bounded by

$$B_{k,q} = t_{\mathrm{MA}} + \max_{p=m_q,\dots,n_q-1} (w_{k,p} + \Delta_{k,p}) + \sum_{p=m_q}^{n_q-1} \delta_{k,p} + \max_{p=m_q+1,\dots,n_q} c_{k,p} + t_{\mathrm{release}}. \tag{B.4}$$

The production implementation applies a 5 % safety factor and takes

$$M_{G,k} = \max\left\{F_k,\ 1.05 \max_{q\in Q} B_{k,q}\right\}. \tag{B.5}$$

This bound is conservative because its entry, running time, and clearing components need not attain their respective maxima for the same section. The per-level deactivation constant follows as

$$M^H_{t,\ell} = \sum_{k\in\mathcal{K}^+} \pi_k \big[\ell\, F_k + (1-\ell)\, M_{G,k}\big]. \tag{B.6}$$

Because the weight on $M_{G,k}$ is $(1-\ell)$, the bound is tighter in years with high TIM deployment, and $M_{G,k}$ is used only to deactivate non-selected inequalities rather than as an explicit upper bound on G^r_k.

B.3. Certified lower bound

The production formulation includes the certified lower bound L_k to strengthen the linear relaxation. The bound is defined in Equation (5.11) and computed once per scenario by a bottleneck-path dynamic program over the interval graph. Let $D_k(p)$ denote the smallest maximum arc headway of any path from m_q to position p. Starting from $D_k(m_q) = 0$,

$$D_k(j) = \min_{i:\,(i,j)\in A_q} \max\big\{D_k(i),\ h_{\mathrm{TTD},k,ij}\big\}, \tag{B.7}$$

and the interval bound is $D_k(n_q)$, so that

$$L_k = \max_{q\in Q} D_k(n_q).$$

At integer solutions this bound is already implied by the selected section path, but it strengthens the linear relaxation when the arc variables are fractional. The recurrence was verified against exhaustive enumeration on reduced instances.

C Optimisation verification and preprocessing diagnostics

This appendix supports the model of Chapter 5 in two parts. The first part documents the analytical headway model detail used in preprocessing: the computation of the precomputed headway components, a sensitivity analysis of these components, and the station-stop motion correction. The second part reports the numerical verification of the implemented optimisation, comprising the big-M and Linear Programming (LP) relaxation diagnostics and the validation suite that confirm the reformulation of Appendix B.

C.1. Computation of the headway components

This section documents how the five headway components that are precomputed independently of the optimisation solve are obtained. The setting, sight and reaction, and release times are treated as fixed constants, while the approach and clearing times are precomputed per candidate TTD boundary position and enter the optimisation depending on which candidate positions the layout selects as section entries or exits. The sixth component, the running time through the release section t_{run}, depends on the release regime under ETCS HTD and is developed in Section 5.2 of the main report rather than here.

C.1.1. Approach time

The approach time is the time associated with the braking distance of the following train before it reaches the release section entry. The full ETCS supervised braking model depends on complex train-specific deceleration profiles, gradient and adhesion corrections, brake build-up times, and national safety parameters (European Union Agency for Railways, 2023; European Union Agency for Railways et al., 2023b). This study therefore uses a representative speed-dependent braking deceleration table $a(v)$ that the Dutch infrastructure manager provides.

The approach time at each candidate TTD boundary position is a precomputed position-dependent term: its value at every candidate entry position is obtained before the optimisation solve (Section 5.5.2), and the value that enters the headway calculation for a given layout is the one at the boundary position that the layout selects as the section entry.

Because the braking distance depends on the following train's maximum permissible speed profile on approach to the release section entry, it differs from one location to another, so the framework computes the approach time for every candidate release section entry along the corridor. A braking curve intersection method obtains it: the method constructs a braking curve backwards from the entry position, starting from zero speed. The full speed-dependent braking distance from speed v to standstill is

$$D(v) = \int_0^v \frac{u}{a(u)}\, du \tag{C.1}$$

and follows from numerical evaluation using the speed-dependent deceleration table $a(v)$. The intersection between this braking curve and the running speed profile of the following train determines the speed v^* at which the follower enters the braking zone. The approach time in the blocking time model is the time the follower requires to traverse the braking distance at running speed (Pachl, 2021). Under the approximation of constant deceleration $a(v^*)$ at the intersection speed, the braking distance simplifies to $D(v^*) \approx v^{*2}/(2a(v^*))$ and the approach time becomes

$$w_p = \frac{D(v^*)}{v^*} \approx \frac{v^*}{2a(v^*)}. \tag{C.2}$$

The first equality is the definition: braking distance divided by running speed. The second equality holds exactly under constant deceleration and is first-order accurate when $a(v)$ varies slowly over the speed range $[0, v^*]$. The formulation accepts any speed-dependent deceleration table $a(v)$ as input and is therefore not restricted to a specific train type or braking system. This constant-deceleration assumption evaluates the approach time at the single intersection speed v^* rather than by integrating the full ETCS braking curve over the approach distance. The sensitivity analysis in Section C.2 shows why this is acceptable for layout optimisation.

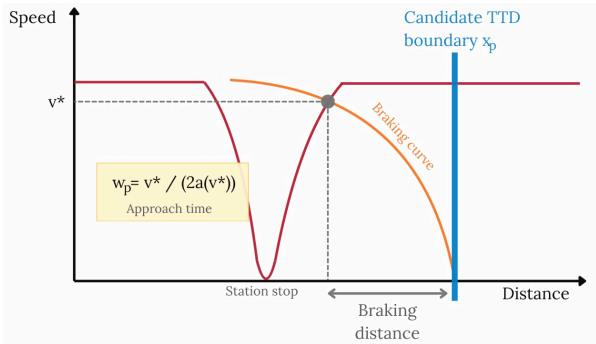

Figure C.1: Approach time computation by braking curve intersection.

C.1.2. Clearing time

The clearing time is the time required for the rear of the leading train to pass the clearing point at the release section exit, that is, the TTD boundary at which the trackside train detection system registers the release section as vacated. For a possible TTD boundary exit position x_q, it follows from the cumulative running time profile of the leader over its train length:

$$c_{q,\text{leader}} = t_{\text{leader}}(x_q + l_{\text{leader}}) - t_{\text{leader}}(x_q). \tag{C.3}$$

Like the approach time, the framework precomputes the clearing time for each candidate TTD boundary exit position and leader type. It is not a free decision variable, but the value entering a release section headway depends on which exit boundary x_q the layout selects. In contrast to t_{run}, it depends only on the exit position and the leading train, not on the release section entry.

C.1.3. Constant components

This study treats three of the six components, the movement authority setting time t_{setting}, the sight and reaction time t_{sr}, and the release time t_{release}, as constants. Chapter 6 specifies the numerical values adopted for them. The remaining components depend on train-specific parameters such as speed, acceleration, deceleration, stopping pattern, and train length (Pachl, 2021). For compact notation, the setting time and the sight and reaction time are combined as

$$t_{\text{MA,follower}} = t_{\text{setting}} + t_{\text{sr}}. \tag{C.4}$$

C.2. Sensitivity of the analytical headway components

To assess the relative influence of each headway component on the slow track-releasing pair headway $h_{\text{TTD},pq}$, a one at a time sensitivity analysis is performed. Each component of the analytical headway expression is varied by $\pm 20\,\%$ from its nominal value while keeping all other components at their base values. The result is presented as a tornado diagram in Figure C.2.

The figure contains three panels that span the speed and section length range of a mixed-traffic corridor: $v = 40\,\text{km/h}$ with $d = 1\,\text{km}$ (station approach and departure zones), $v = 80\,\text{km/h}$ with $d = 2\,\text{km}$ (speed restrictions at switches or curves), and $v = 130\,\text{km/h}$ with $d = 3\,\text{km}$ (near cruise speed on open track). In each panel the five headway components are ordered by influence, and the horizontal axis shows the change in slow track-releasing pair headway $\Delta h_{\text{TTD},pq}$ under a $\pm 20\,\%$ variation of each component. The spread Δh between the two cases measures the sensitivity to that component.

Three observations emerge from the tornado diagram:

1. The running time through the release section is the dominant component in all three panels, with a spread of $\Delta h \approx 33\text{-}36\,\text{s}$. This component depends on the release section length together with the train's running profile through that section, and is therefore the most relevant term for TTD boundary placement.
2. The approach time sensitivity increases with line speed (from $\Delta h = 3.2\,\text{s}$ at $v = 40\,\text{km/h}$ to $\Delta h = 12.6\,\text{s}$ at $v = 130\,\text{km/h}$), while the clearing time sensitivity decreases (from $\Delta h = 7.2\,\text{s}$ to $\Delta h = 2.2\,\text{s}$). Both remain smaller than the running time contribution. The approach time is a precomputed position-dependent term that varies smoothly along the corridor, so approximations affect neighbouring sections similarly and have limited influence on the relative ranking of candidate TTD boundaries.
3. The system constants t_{MA} and t_{release} produce a constant, small spread ($\Delta h = 2.8\,\text{s}$ and $\Delta h = 0.8\,\text{s}$, respectively) regardless of operating conditions. Because they are uniform across the corridor, they shift all headways equally and do not affect the optimal layout.

These observations justify the modelling simplifications adopted in the analytical headway model. Because the running time through the release section dominates, the following approximations have limited influence on the relative ranking of TTD layouts:

- The approach time w_p is evaluated at a single speed v^* at the braking curve intersection rather than by integrating the full ETCS braking curve over the approach distance. This is adequate because w_p varies smoothly along the corridor and enters the headway expression as a precomputed position-dependent term whose value is selected by the layout.
- The clearing time c_p is computed by integrating the leader's cumulative running time profile over its own length (Equation (C.3)). This avoids the overestimation that occurs with a single-point speed evaluation at low-speed positions such as station stops (see Section 5.5.2).
- The system constants t_{MA} and t_{release} are treated as uniform values. In practice, t_{MA} depends on the communication latency between the onboard unit and the RBC, but the variation is small relative to the running time component and does not affect the relative ranking of possible TTD boundaries.

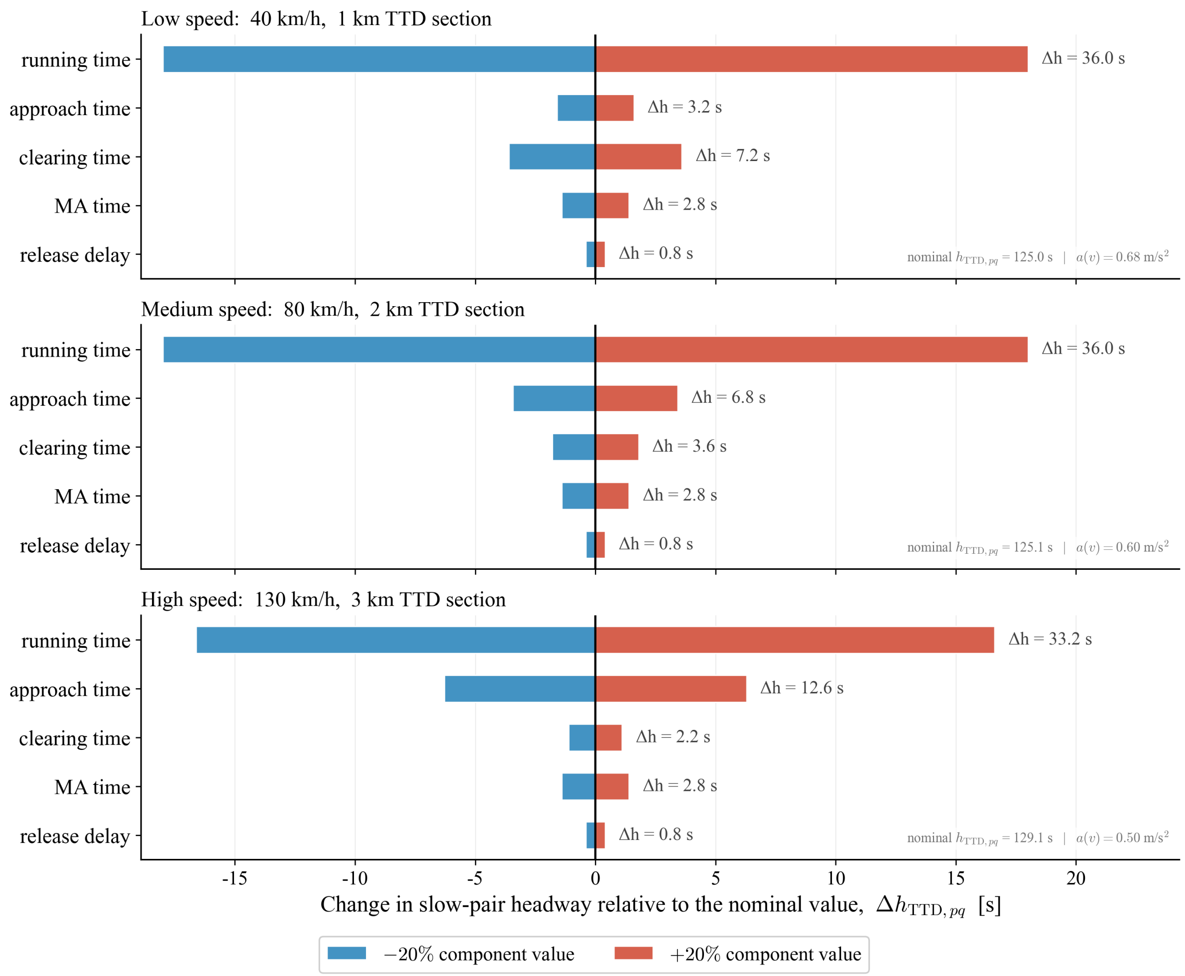

Figure C.2: Tornado diagram showing the sensitivity of the slow track-releasing pair headway $h_{\mathrm{TTD},pq}$ to ±20 % variation of each headway component.

C.3. Station-stop motion correction

This section documents the station-stop motion correction that compensates the speed floor used near platforms in the preprocessing that builds the cumulative running time profile $t_{\mathrm{train}}(x)$ of Section 5.5.2. It reuses the same deceleration table $a(v)$ as the braking curve approximation of Section C.1.1, so the two remain mutually consistent: a change to the table propagates to both.

For numerical stability the preprocessing clips the speed profile at a floor $v_{\min} = 20\,\mathrm{km/h}$ before computing per-segment running times as $\Delta t = \ell_{\mathrm{seg}}/v$, so that the increment does not diverge as $v \to 0$ at a scheduled stop. As a consequence, the last stretch of the brake-in curve (where the physical speed would fall from $v_{\min}$ to 0) and the first stretch of the accel-out curve (from 0 to $v_{\min}$) are integrated at $v = v_{\min}$ rather than along the true profile, systematically under-counting the running time in those two zones. The omitted time is added back analytically as a scalar correction per stop.

Brake side. The true time to decelerate from $v_{\min}$ to standstill is $\int_0^{v_{\min}} \mathrm{d}v/a(v)$, whereas the preprocessing traverses the same distance $D(v_{\min})$ of Equation (C.1) at the constant clipped speed $v_{\min}$, i.e. in $D(v_{\min})/v_{\min}$. The under-counted time is the difference,

$$\tau_{\mathrm{brake}} = \int_0^{v_{\min}} \frac{1}{a(v)} \left(1 - \frac{v}{v_{\min}}\right) \mathrm{d}v, \tag{C.5}$$

evaluated numerically from the same table $a(v)$.

Acceleration side. With a constant departure acceleration a_{acc}, the true time from 0 to $v_{\min}$ is $v_{\min}/a_{\text{acc}}$, while the preprocessing covers the corresponding distance $v_{\min}^2/(2a_{\text{acc}})$ at speed $v_{\min}$, i.e. in $v_{\min}/(2a_{\text{acc}})$. The difference has the closed form

$$\tau_{\text{accel}} = \frac{v_{\min}}{2\,a_{\text{acc}}}. \tag{C.6}$$

Total correction. The per-stop motion correction is

$$\tau_{\text{corr}} = \tau_{\text{brake}} + \tau_{\text{accel}}, \tag{C.7}$$

added once, together with the scheduled dwell, at the first grid node strictly downstream of each station (the snap-after index). Because every running time increment is obtained by finite differencing of $t_{\text{train}}(x)$, this guarantees that the station effect is counted exactly once: a section that crosses the station inherits the full dwell and correction, while a section that does not cross it inherits only movement time (the property verified in Chapter 6).

Magnitude. For the case study train types ($v_{\min} = 20\,\text{km/h} = 5.556\,\text{m/s}$, $a_{\text{acc}} = 0.5\,\text{m/s}^2$, $a(v_{\min}) = 0.70\,\text{m/s}^2$) the two contributions evaluate to $\tau_{\text{brake}} = 3.97\,\text{s}$ and $\tau_{\text{accel}} = 5.56\,\text{s}$, giving a total correction of $\tau_{\text{corr}} \approx 9.52\,\text{s}$ per ST stop. The correction is not hard-coded: it is recomputed from $a(v)$, $v_{\min}$, and a_{acc} at every run, so changes to the deceleration table, the speed floor, or the departure acceleration propagate automatically and stay consistent with the braking curve approximation of Section C.1.1.

C.4. Big-M and LP diagnostics

This section reports the numerical diagnostics used to assess the optimisation formulation. The big-M diagnostic exposes implementation errors in the calibrated constants, and big-M validity itself is established analytically in Section B.2. The LP relaxation check is used as a diagnostic for the strength of the formulation, not as a proof of mathematical correctness. All reported diagnostics use the Type 1 case study instance at $c_{\text{train}} = €\,84{,}420$ at a 3 % relative MIP gap tolerance chosen to keep the diagnostic solves fast, whereas the production optimisation runs in Chapter 7 use a 1 % tolerance.

C.4.1. LP strengthening cut effectiveness

Table C.1 reports the effect of the two built-in LP-strengthening components, the tight scenario-specific big-M calibration (Section B.2) and the certified lower bound L_k (Section B.3), on the root-node LP bound. The baseline is the formulation with both components disabled. Each row shows the LP bound when only the indicated component is active. The final row shows the combined effect.

Table C.1: Effect of the two built-in LP-strengthening components on the root-node LP bound.

Configuration	Root LP bound on $\bar{H}$ [s]	Ratio to F_{weighted}
Baseline (both components off)	142.3778	0.9678
Tight $M^H_{t,\ell}$ only	142.3778	0.9678
Certified $G^r_k \geq L_k$ only	147.1138	1.0000
Both components active	147.1138	1.0000

With both components active, the root-node LP bound equals $F_{\text{weighted}} = \sum_{k\in\mathcal{K}^+} \pi_k F_k = 147.11\,\text{s}$ exactly. The certified lower bound alone attains this floor: raising every G^r_k to at least L_k combined with the convex-combination structure of the expected headway constraint lifts $\bar{H}$ to F_{weighted} in the LP relaxation. The baseline and the configuration with only the tightened big-M both give a root LP bound of 142.38 s, about 4.74 s below F_{weighted}, because the LP relaxation exploits fractional level selection variables b^t_ℓ to deactivate the disjunctive C2 constraints partially, allowing $\bar{H}$ to fall slightly below the layout-independent floor. The tightened per-level $M^H_{t,\ell}$ does not improve the root LP bound in isolation because the size of the big-M affects the LP polytope only where the disjunction is not partially deactivated. It nevertheless remains valuable for branch and bound performance through better numerical conditioning. The resulting LP-to-MIP gap is reported in Section C.4.4.

C.4.2. Scenario-specific big-M constants

The big-M constants are computed from the preprocessed speed profiles and differ between leader-follower scenarios. Table C.2 reports the fast-release headway floor F_k and the conservative slow-release bound $M_{G,k}$ used in the reported arc-flow optimisation runs. The bound $M_{G,k}$ enters the per-level expected headway deactivation constant $M^H_{t,\ell}$ (Section B.2), which is the only big-M family actually applied in the arc-flow formulation.

Table C.2: Scenario-specific fast-release headway floor F_k and conservative slow-release bound $M_{G,k}$ used in the reported arc-flow optimisation runs.

Scenario	F_k [s]	$M_{G,k}$ [s]
ST-IC	255.60	483.76
IC-ST	64.39	197.84
IC-IC	66.74	197.84
ST-ST	134.24	350.89

The largest $M_{G,k}$ occurs for the ST-IC case, because the stopping leader accumulates the largest running time difference relative to the non-stopping follower across the longest inter-mandatory segment. The homogeneous IC-IC case has a small $M_{G,k}$, because there is no leader-follower running time difference and no dwell time.

C.4.3. Big-M implementation check

The validity of the per-level big-M constants $M^H_{t,\ell}$ follows analytically from Section B.2: they are constructed from the conservative bound $M_{G,k}$, which is at least the worst-case section headway over any feasible layout of scenario k, and therefore bound the expected headway disjunction on the entire feasible set. The numerical comparison in Table C.3 is retained as an *implementation* smoke-check on the calibrated constants: solving the model twice, once with the tightened bounds and once with a deliberately loose slow-release bound of 5000 s, should return incumbents that agree within the numerical tolerance implied by the reported MIP gap. Both solves use a 3 % relative gap here, and the resulting numerical tolerance is

$$\epsilon_M = 3 \cdot 0.03 \cdot \max(\bar{H}_{\text{tight}}, \bar{H}_{\text{loose}}).$$

Table C.3: Big-M implementation check on the Type 1 instance at four cost levels with a 3 % MIP gap, comparing the calibrated tight formulation against a deliberately loose slow headway bound ($M_G = 5000\,\text{s}$), with the numerical tolerance $\epsilon_M = 3 \cdot \text{gap} \cdot \max(\bar{H}_{\text{tight}}, \bar{H}_{\text{loose}})$ below which observed differences are attributable to the reported MIP gap rather than to a bound implementation error.

$C_{\max}$ [€]	$\bar{H}_{\text{tight}}$ [s]	$\bar{H}_{\text{loose}}$ [s]	Difference [s]	ϵ_M [s]
1,950,000	168.7486	168.7486	0.0000	15.1874
2,094,736	159.3223	159.3223	0.0000	14.3390
2,300,000	156.3778	155.1503	1.2275	14.0740
2,500,000	154.1981	154.4552	0.2571	13.9010

All differences remain below the numerical tolerance ϵ_M, so the two solves return incumbents consistent with the reported MIP gap. The check therefore reveals no implementation error in the tightened bounds. At the two cheapest cost levels the tight and loose formulations return identical $\bar{H}$ to four decimals. The small residual differences at the two higher cost levels (1.23 s at $C_{\max} =$ € 2.30M, 0.26 s at € 2.50M) fall within the alternative-optima ambiguity implied by the 3 % MIP gap. Big-M validity itself does not depend on this numerical comparison. It is established analytically in Section B.2.

C.4.4. LP relaxation diagnostics

The LP relaxation was solved for selected cost levels of the arc-flow formulation with both built-in LP-strengthening components active. In this relaxation, the binary boundary placement variables and level selection variables are replaced by continuous variables, such that $y^r_p \in [0, 1]$ and $b^t_\ell \in [0, 1]$. The resulting objective value provides a lower bound on the optimal MIP objective. With the two built-in components (Sections B.2 and B.3), the LP relaxation value equals $F_{\text{weighted}} = \sum_{k\in\mathcal{K}} \pi_k F_k$ at the root node. The LP relaxation is used as a diagnostic for the strength of the formulation rather than as a validation of the TTD layout itself, since fractional boundary placements do not correspond to implementable TTD layouts.

Table C.4: LP relaxation diagnostics for selected cost levels (Type 1, $c_{\text{train}} = €\,84{,}420$, both built-in LP-strengthening components active).

$C_{\max}$ [€]	$\bar{H}_{\text{LP}}$ [s]	$\bar{H}_{\text{MIP}}$ [s]	LP gap [%]
1,950,000	147.1138	168.7486	12.82
2,065,790	147.1138	161.0353	8.64
2,210,526	147.1138	158.3414	7.09
2,500,000	147.1138	153.9298	4.43

The LP relaxation lower bound is constant at $F_{\text{weighted}} = 147.11\,\text{s}$ across all cost levels, for the reason given in Section C.4.1. The MIP $\bar{H}$ decreases monotonically with the budget, from 168.75 s at $C_{\max} =$ €1.95M to 153.93 s at $C_{\max} =$ €2.50M, narrowing the LP-to-MIP gap from 12.82 % to 4.43 %.

C.4.5. Handcase validation output

The handcase validation (test T7 in the validation suite, Section C.5) compares the MIP implementation with brute-force enumeration on a small synthetic corridor. Because the number of possible TTD boundary positions is small, all feasible layouts can be enumerated explicitly. The brute-force solution therefore provides an independent global optimum against which the MIP solution can be checked.

The synthetic test cases are designed to isolate different parts of the formulation: endpoint-only layouts, the selection of additional boundaries among a 5-position free subset, the all-boundaries case at an unrestricted cost level, full integrity operation with $\rho = 1$, mixed integrity operation with $\rho = 0.4$, and infrastructure restrictions through a forbidden boundary position. The instance is the Type 1 corridor restricted to the 12 mandatory positions plus a 5-position evenly spaced free subset, giving $2^5 = 32$ enumerable layouts per case. The IC-IC scenario is selected with unit weight to keep the brute-force evaluation tractable. In all six cases, the MIP and brute-force enumeration returned the same maximum line headway within a tolerance of 10^{-4} s.

Table C.5: Handcase validation results for Type 1 restricted to a 5-position free subset (IC-IC single scenario), where MIP and brute-force headways agree exactly in all cases.

Case	Test condition	$\bar{H}_{\text{brute}}$ [s]	$\bar{H}_{\text{MIP}}$ [s]	Result
1	$\rho = 0$, endpoints only	179.8188	179.8188	PASS
2	$\rho = 0$, one extra boundary	142.5178	142.5178	PASS
3	$\rho = 0$, all boundaries affordable	142.5178	142.5178	PASS
4	$\rho = 1$, $\bar{H}$ must equal F	66.7378	66.7378	PASS
5	$\rho = 0.4$, mixed F/G formulation	112.2058	112.2058	PASS
6	Forbidden position excluded	163.4493	163.4493	PASS

C.5. Validation suite

This section documents the unified validation suite for the optimisation model. The automated implementation is archived with the model and is not reproduced here. On the Type 1 case study instance ($c_{\text{train}} = €\,84{,}420$, 3 % MIP gap), the suite verifies preprocessing consistency and arc-set completeness (T1), the implemented optimisation and linearisation constraints (T3), post-hoc headway reconstruction (T4), the F_{weighted} lower bound $\bar{H} \geq \sum_k \pi_k F_k$ on every Pareto point (T6), the big-M implementation check (T5, Section C.4.3), the brute-force handcase (T7, Section C.4.5), homogeneous-Δ behaviour (T10), the endogenous deployment trajectory and cost decomposition (D1-D3), and the initial-state shape (N5). All executed checks pass, as summarised in Table C.6. The weight-extreme and monotonicity checks (T2, T8, T9) require longer-gap solves and are not included in the reported suite.

Table C.6: Validation suite results on the Type 1 case study instance.

Test	Description	Result
T1	Preprocessing consistency and complete arc set	PASS
T3	MIP constraint verification	PASS
T4	Post-hoc headway match	PASS
T5	Big-M implementation check	PASS
T6	F_{weighted} lower bound on $\bar{H}$ across the Pareto sweep	PASS
T7	Handcase brute-force	PASS
T10	Homogeneous Δ	PASS
D1	Endogenous ρ_t selection	PASS
D2	Deployment level expected headway disjunction	PASS
D3	TIM cost decomposition	PASS
N5	Brownfield init shape	PASS

D RailSys simulation data

This appendix provides additional RailSys output supporting the microscopic evaluation in Chapter 7. Section D.1 reports the analytical versus microscopic line headway comparison underlying the consistency checks, and Section D.2 adds the time-distance diagrams used to inspect how the 15 min standstill of the first ST service at AMPO propagates through the timetable.

D.1. Analytical versus microscopic line headway comparison

This section reports the numerical basis for the analytical-microscopic consistency checks summarised in Chapter 7. Table D.1 compares the analytical and microscopic RailSys line headway at the current ETCS L2 reference condition, and Table D.2 reports the stage-specific comparison for the selected heterogeneous T1 and T2 migration layouts under partial TIM deployment. In both cases the RailSys line headway is derived from the UIC 406 compression without timetable allowances, so that it is comparable with the analytical value. The homogeneous T3 layout is part of the capacity consumption evaluation but is not included in this headway comparison.

Table D.1: Reference condition comparison of the analytical and microscopic RailSys expected line headway for the current brownfield 44-boundary layout at $\rho = 0$, with all relevant movements dependent on TTD level release.

Condition	Analytical [s]	RailSys [s]	Difference [s]
ETCS L2, 44 boundaries, $\rho = 0$	155.15	162.00	+6.85

Table D.2: Stage-specific comparison of the analytical and microscopic RailSys expected line headway for the selected T1 and T2 layouts under partial TIM deployment, with the difference given as RailSys minus analytical headway.

Configuration	Metric [s]	S1	S2	S3	S4
T1	Analytical	162.2	161.1	162.1	162.2
T1	RailSys	167.2	165.8	165.1	160.9
T1	Difference	+5.0	+4.7	+3.0	−1.3
T2	Analytical	162.6	162.4	162.3	162.2
T2	RailSys	167.8	166.5	166.0	167.2
T2	Difference	+5.2	+4.1	+3.7	+5.0

D.2. Deterministic disturbance time-distance diagrams

Figures D.1 and D.2 show the time-distance diagrams for the selected T1 ETCS HTD migration layout across all four migration stages, under both the best-case and worst-case TIM assignments. In the discrete KPI 3 trainset representation, stage 4 has all eight trainsets TIM-equipped, so its best-case and worst-case assignments coincide. The diagrams are diagnostic: they illustrate how the disturbance propagates to downstream services, while the quantitative recovery comparison is reported in Chapter 7.

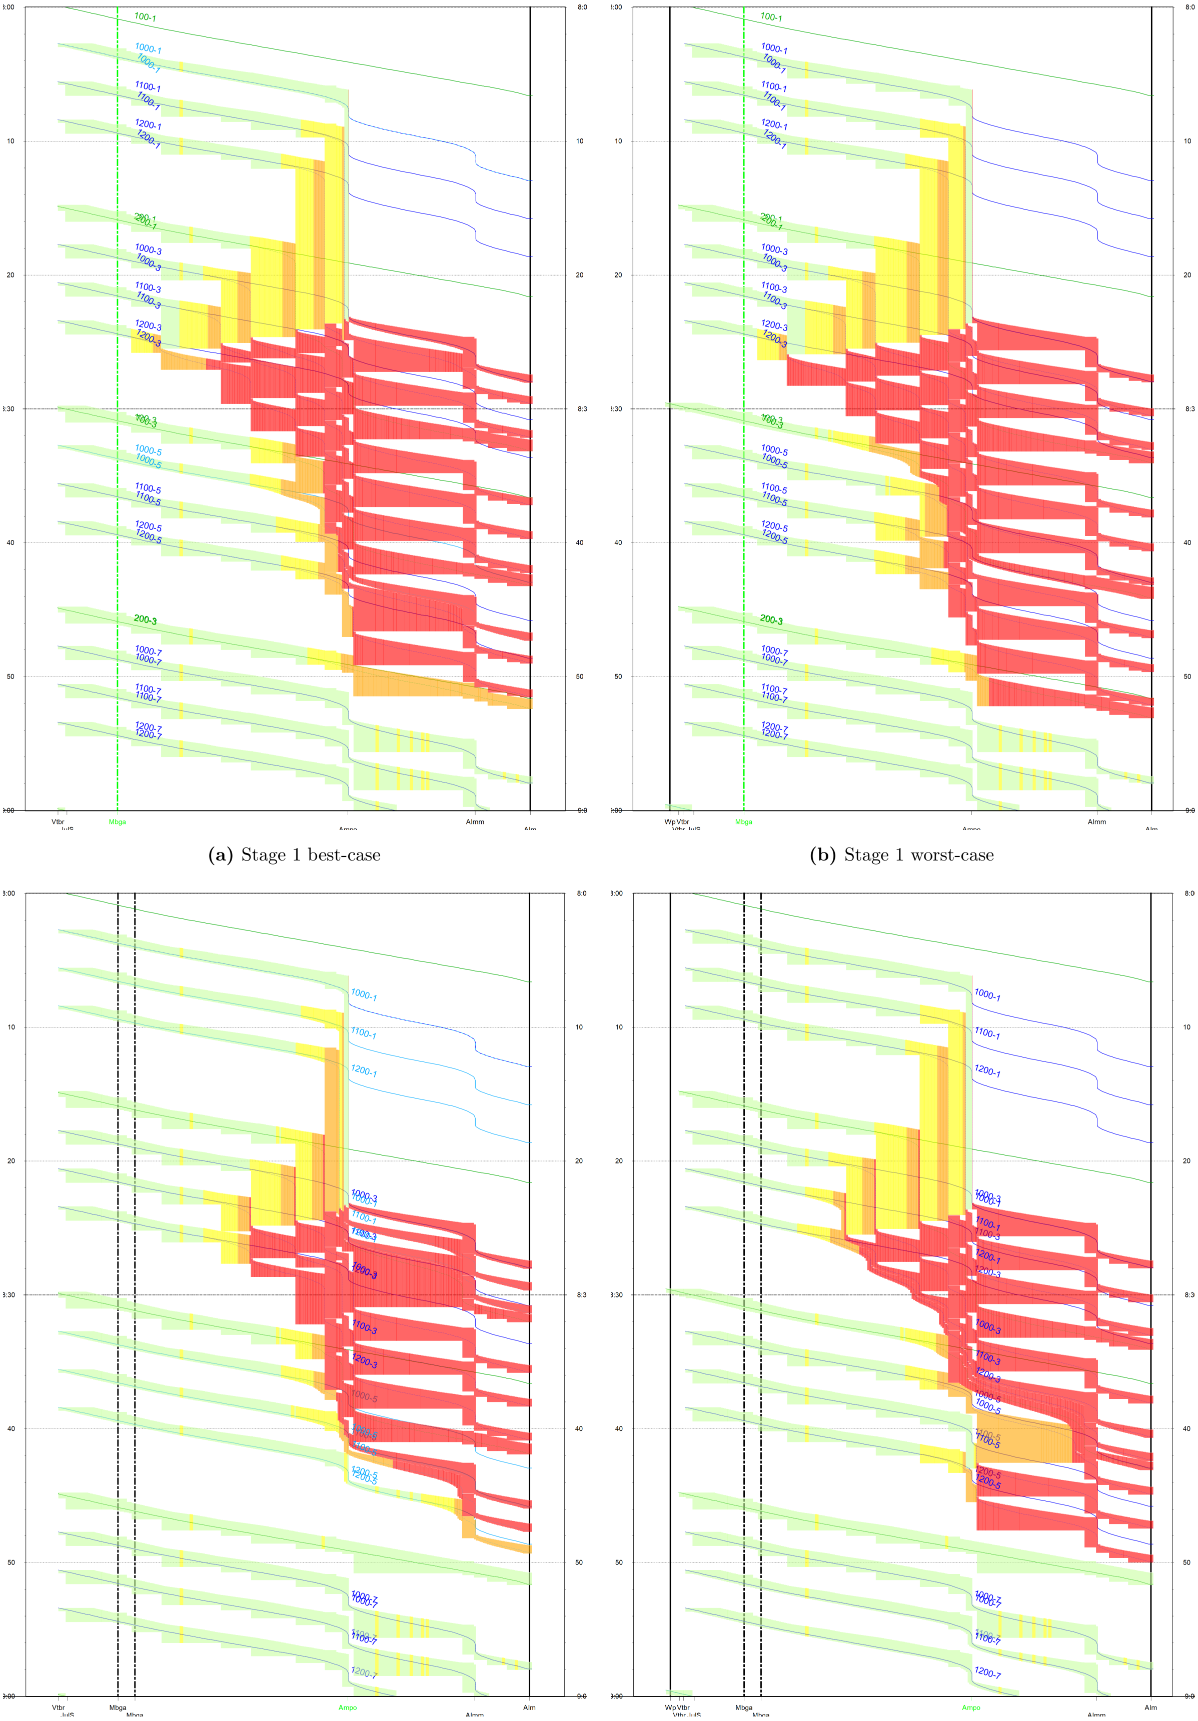

(a) Stage 1 best-case

(b) Stage 1 worst-case

(c) Stage 2 best-case

(d) Stage 2 worst-case

Figure D.1: Case study time-distance diagrams for the deterministic T1 disturbance at AMPO, stages 1 and 2, under the best-case and worst-case TIM assignments.

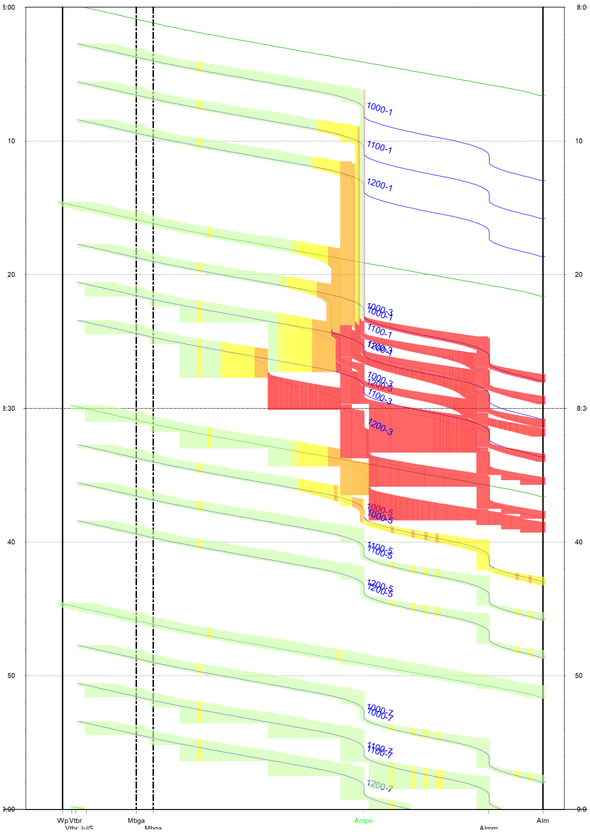

(a) Stage 3 best-case

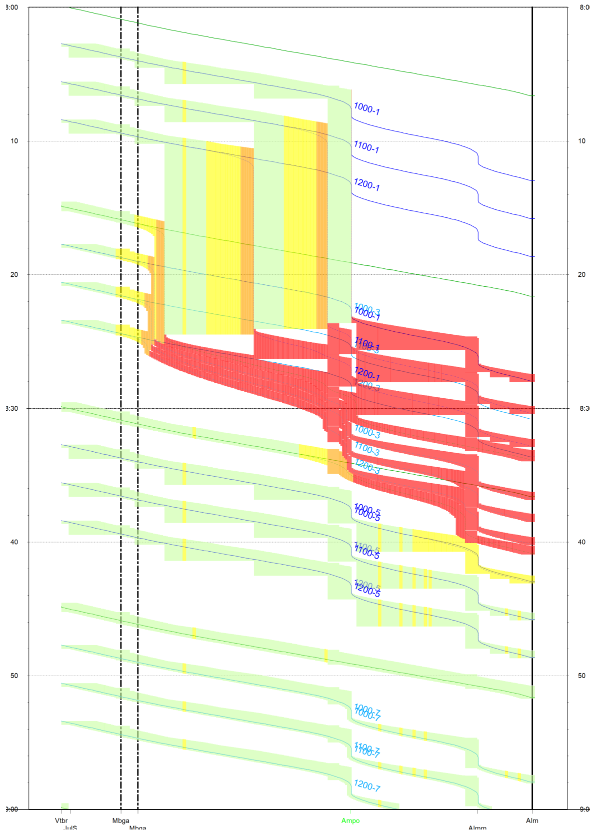

(b) Stage 3 worst-case

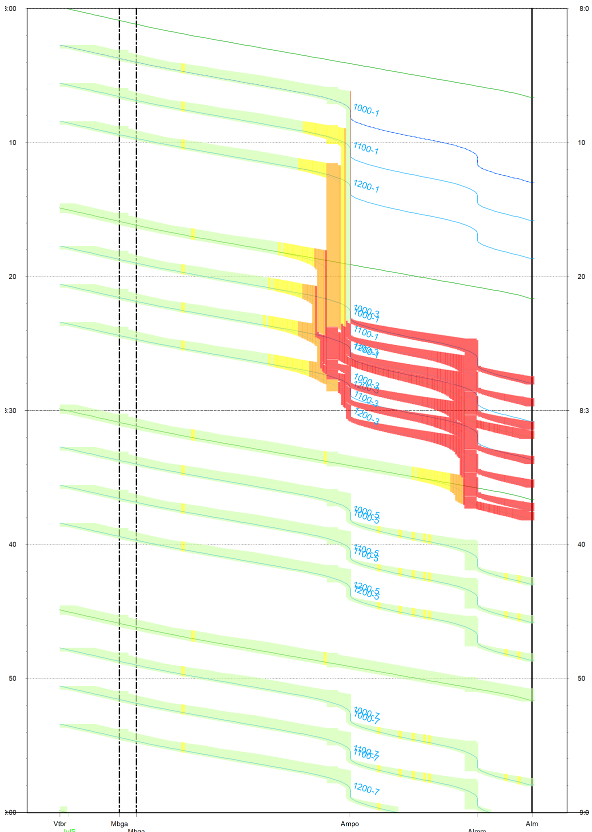

(c) Stage 4 (all trainsets TIM-equipped)

Figure D.2: Case study time-distance diagrams for the deterministic T1 disturbance at AMPO, stages 3 and 4, under the best-case and worst-case TIM assignments, with the two assignments coinciding in stage 4 where all eight trainsets are TIM-equipped.

E AI Statement

During the preparation of this thesis, several artificial intelligence (AI) tools were used as supporting tools for writing, translation, coding, consistency checking, post-processing support, and research-related brainstorming. ChatGPT (OpenAI, 2026), Claude (Anthropic, 2026), and DeepL (DeepL SE, 2026) were used to improve the efficiency and quality of the research process. The final responsibility for the content, modelling choices, interpretation of results, and written text remains entirely my own.

ChatGPT was used throughout the thesis process as a writing and reasoning assistant. I used it to rewrite and improve text that I had written myself, to help formulate academic paragraphs based on my own input, and to check whether chapters had a clear structure, logical flow, and sufficient scientific reasoning. ChatGPT was also used for brainstorming sessions, for discussing the formulation of the optimisation model, and for support with LaTeX, Python, and Excel formulas. This included debugging Python code, understanding programming errors and LaTeX errors, drafting LaTeX tables and figure captions, and formulating Excel formulas for the post-processing of RailSys and KPI data. These interactions helped me to better understand the modelling and programming steps and contributed to a steep learning curve during the development of the optimisation framework.

DeepL was used to translate text that I had written in Dutch into English. The translated text was always reviewed and adapted by me before being included in the thesis, to ensure that the terminology, technical meaning, and academic style were appropriate.

Claude AI Pro and Max, including Claude Code, were used as supporting tools while working with the thesis `.tex` files. These tools helped with checking labels, cross-references, terminology consistency, acronyms, and the structure of LaTeX files. They were also used to rewrite or refine text based on text fragments and instructions provided by me. In addition, they supported practical editing tasks in the report files, such as checking whether references and labels were used consistently and whether terminology remained consistent across chapters.

AI tools were also used to support the structuring of chapters and the development of the research argumentation. This included brainstorming about the interpretation of results, limitations, recommendations, and design guidelines. In these cases, AI was used as a discussion and reflection tool to test whether the reasoning was clear and complete, not as a replacement for my own judgement.

AI was not used to independently generate final research results without control. It did not make final modelling decisions, define the research scope, select the final case-study assumptions, or interpret results without my assessment. Numerical outputs from the optimisation model, Excel post-processing, and RailSys simulations were produced using the project files and were checked by me before being included in the thesis.

The AI tools were therefore used as assistants in the writing, coding, editing, translation, and reflection process, not as autonomous authors of the thesis. All AI-generated suggestions, text edits, code edits, LaTeX edits, and structural recommendations were personally reviewed, assessed, and modified where necessary. Final decisions on the research scope, model formulation, assumptions, case-study setup, results interpretation, conclusions, design guidelines, and recommendations were made by me. The use of AI accelerated the research process and supported my learning, but the final thesis reflects my own understanding, judgement, and responsibility.