Load Hosting Capacity of Urban Traction Networks

Case study: Calandlijn metro network in Rotterdam

Yash Dwarkasing

Technische Universiteit Delft
Rotterdamse Elektrische Tram N.V.

Load Hosting Capacity of Urban Traction Networks

Case study: Calandlijn metro network in Rotterdam

by

Yash Dwarkasing

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on July 15, 2026 at 11:00.

Student number:	5458242	
Project duration:	November 10, 2025 – July 3, 2026	
Thesis committee:	Prof. dr. ir. P.T.M. Vaessen,	TU Delft
	Dr. ir. D. van der Born,	TU Delft
	Ing. J. Verschoor,	Rotterdamse Elektrische Tram N.V.
	Dr. ir. I. Diab,	TU Delft
	Dr. Dipl.-Ing. A. Lekić,	TU Delft

An electronic version of this thesis is available at http://repository.tudelft.nl/.

Preface

This thesis marks the provisional end of an educational journey that I have followed for most of my life. From my earliest introduction to electricity and electrical energy in high-school physics, I was both fascinated and somewhat intimidated by the subject. Nevertheless, I felt drawn to pursue it further, particularly in light of the transition towards alternative methods of generating and transporting electrical energy.

This project gave me the opportunity to explore traction power grids from both theoretical and practical perspectives. My interactions with the engineers involved in the project allowed me to develop a broader understanding of the many facets of traction power systems, including aspects beyond the scope of this research.

First, I would like to thank my company supervisor, Jaap Verschoor, for his support throughout this project and for his patience while discussing (occasionally) overly convoluted approaches. I am equally grateful to my daily supervisor, Dennis van der Born, for his academic guidance and for the discussions that helped shape the direction of this research.

I would also like to thank Kenneth Bruninx and Peter Vaessen for their involvement in the project and their valuable feedback on this report. My gratitude also extends to the other members of the thesis committee for their time and contributions.

Beyond the academic and professional support I received, I would like to thank my friends and family for their encouragement, understanding, and patience throughout this period. Finally, and most importantly, I am deeply grateful to my parents. Without their support, reaching this point would not have been possible.

Yash Dwarkasing
July 2026

Abstract

Dutch urban distribution grids face growing congestion, with many areas unable to accommodate new connections until reinforcement is completed years from now. Light-rail traction power systems, dimensioned to power trams and metros, represent an underexplored asset on the same grid: their substations are sized for peak vehicle loads, leaving capacity that could in principle host additional external loads such as fast chargers or local industry.

This thesis develops a method to quantify and allocate that hosting capacity. Train traffic makes the demand on each substation vary unpredictably from second to second; the network's thermal limits depend not on instantaneous load but on how long an overload persists; and a load added at one substation shifts power flows across the entire line. A worst-case assessment returns conservatively small numbers; a single-snapshot average ignores the events that actually constrain operation. Treating the substations collectively, as a single aggregate capacity, hides the spatial trade-offs that determine where loads can actually be sited.

The method addresses this in three stages. A simulation of the network generates a large pool of realistic operating scenarios. The pool is then reduced to a representative subset using a new clustering approach that preserves the scenarios most relevant to feasibility. The reduced set feeds a multi-objective optimisation, formulated with one capacity objective per substation, that exposes the spatial trade-offs between total hosting capacity and how evenly the load can be distributed. Probabilistic operating limits replace deterministic worst-case bounds where the underlying engineering standards admit them, so that brief admissible transients do not subordinate the result. Of the two metaheuristics evaluated, swarm optimisation handles the search space most effectively at moderate confidence levels, though its advantage narrows as the reliability requirement tightens.

Applied to the Calandlijn of the Rotterdam metro, operated by Rotterdamse Elektrische Tram N.V., the method identifies a hosting capacity of 2 to 5 MW across the line's twelve substations, with the achievable value depending on how evenly the load is distributed. A more important finding emerges from the analysis: the binding limit is not uniform across the network. On some feeders it is the contractual cap on utility intake rather than the traction infrastructure itself, while on others the smaller substations run close to their thermal limit. The distinction matters for the operator, since the two cases call for different remedies: contract renegotiation in one, physical reinforcement in the other.

Nomenclature

Abbreviations

AC, DC	Alternating current, direct current
BESS	Battery energy storage system
BFS	Backward-Forward sweep
CCP	Chance-constrained programming
CDF	Cumulative distribution function
CPL	Constant power load
CVaR	Conditional value-at-risk
DDSR	Distribution-driven scenario reduction
EV	Electric vehicle
FACTS	Flexible AC transmission system
GA	Genetic algorithm
GS	Gauss–Seidel (solver)
HVAC	Heating, ventilation, and air conditioning
IEC	International Electrotechnical Commission
LVDC	Low-voltage direct current
MOOPF	Multi-objective optimal power flow
NR	Newton–Raphson (solver)
NSGA-II/III	Non-dominated sorting genetic algorithm, versions II and III
OASR	Optimality-aware scenario reduction
OPF	Optimal power flow
P&G	Pulse-and-glide (driver policy)
PCA	Principal component analysis
PDSR	Problem-driven scenario reduction
PF	Power flow
PSO	Particle swarm optimisation
PV	Photovoltaic
RBE	Regenerative braking energy
RES	Renewable energy source
RET	Rotterdamse Elektrische Tram

SCADA	Supervisory control and data acquisition
SMPSO	Speed-constrained multi-objective PSO
SO	Scenario optimisation
SR	Scenario reduction
SS	Substation
TPSS	Traction power supply system

Sets and indices

Ω	Full simulator-generated scenario set
$\mathcal{S}_{\mathrm{R}}$	Reduced scenario set
$\mathcal{X}$	Decision space
$\mathcal{T}$	Set of admissible rectifier tap positions
$\mathcal{N}_{\mathrm{DC}}$	Set of DC nodes
$\mathcal{S}_{\mathrm{SS}}$	Set of substations
$\mathcal{V}$	Set of vehicles
$\mathcal{E}$	Set of conductor branches
$\mathcal{U}$	Set of utility connections
$\mathcal{G}$	Set of regenerative vehicle nodes
$\mathcal{M}$	Set of motoring vehicle nodes
$\mathcal{P}$	Set of passive nodes
N_{SS}	Number of substations ($= 12$ for the Calandlijn)
N_{U}	Number of utility connections ($= 4$ for the Calandlijn)
k	Substation index (also: NR iteration index)
s	Scenario index; also vehicle position
i, j	DC node indices; (ij) denotes a branch
u	Utility-connection index
m	Gauss–Seidel outer-iteration index

Roman symbols

A, B, C	Davis resistance coefficients
e_{ι}	One-hot encoding of binding-constraint label ι
E_u	Energy drawn at utility connection u
$\overline{E}_u$	Contractual energy cap at utility connection u
f_k^{CPL}	Per-substation CPL allocation objective
f^{imb}	Loading-imbalance objective

f^{diss}	Network-dissipation objective
f_u^{util}	Per-utility intake objective
f	Combined multi-objective vector
$F, F_{\mathrm{res}}, F_{\mathrm{grade}}, F_{\mathrm{net}}$	Applied, resistive, grade, and net longitudinal forces on a vehicle
$\mathcal{F}^{\mathrm{hard}}$	Hard-constraint feasibility predicate
g_j	Constraint function indexed by j
g_{chop}	Chopper conductance
$\mathrm{G}_{\mathrm{base}}, \mathrm{G}_{\mathrm{eff}}$	Base and effective DC conductance matrices
I_{ij}	Current in branch (ij)
$I_{\mathrm{SS},k}$	Substation rectifier current
I_{max}	Maximum admissible branch current
I_{thr}	Rectifier conduction threshold
J	Jacobian of the residual vector
K	Number of medoids
k_{max}	NR iteration limit
m	Vehicle mass (kg)
$\mathrm{P}_{\mathrm{CPL}}$	External CPL allocation vector
$P_{\mathrm{CPL},k}$	External CPL allocation at substation k
$\overline{P}_{\mathrm{CPL}}$	Upper bound on per-substation CPL allocation
$P_{\mathrm{sub},k}$	Substation DC power demand
P_{chop}	Chopper dissipation
P_{loss}	DC transmission loss
$P_{\mathrm{util},u}$	Active power at utility connection u
P_{aux}	Vehicle auxiliary load
P_{demand}	Signed vehicle power demand
$\mathrm{P}_{\mathrm{ref}}$	Power-demand vector in the DC power-flow problem
r	Residual vector in the NR formulation
$S_{\mathrm{rated},k}$	Rated apparent power of substation k
t	Rectifier tap-position vector
U_{mean}	Mean useful pantograph voltage
U_{min}	Minimum mean useful pantograph voltage (EN 50388-1)
v	Vehicle velocity
V_i	Voltage at DC node i
$V_{\mathrm{min}}, V_{\mathrm{max}}$	DC node voltage bounds
$V_{\mathrm{SS,min}}, V_{\mathrm{SS,max}}$	Substation-busbar voltage bounds
V_{nom}	Nominal DC voltage (750 V)

$V_{\min}^{\mathrm{veh}}$	Minimum vehicle-node voltage in a scenario
w_s	Empirical weight of scenario s in a reduced set
x	Decision vector $[\mathrm{P_{CPL}}; \mathrm{t}]$

Greek symbols

α	NR damping factor; chance-constraint significance level
δ	Uniform-scaling CPL multiplier (in p.u. of S_{rated})
$\delta_{\max}(s)$	Per-scenario feasibility margin under uniform scaling
δ_{hi}	Upper bracket of the feasibility line search
ε	NR convergence tolerance
η_{regen}	Regenerative braking efficiency
θ	Track gradient angle
$\iota(s)$	Categorical label of the constraint binding at $\delta_{\max}(s)$
λ_k	Substation loading ratio $P_{\mathrm{sub},k}/S_{\mathrm{rated},k}$
$\overline{\lambda}_k(\tau)$	Moving average of λ_k over window of length τ
λ_{trip}	Instantaneous trip-threshold loading ratio
$\overline{\lambda}_1, \overline{\lambda}_2$	Duration-dependent loading bounds at windows τ_1, τ_2
τ_1, τ_2	Moving-average windows (1 min, 4 min)
$\phi 0(s)$	Optimality fingerprint of scenario s
Ω	Full scenario set (see Sets)

Contents

Nomenclature **iv**

1 Introduction **1**
1.1 Problem statement . . . 1
1.2 Technical challenges . . . 1
1.3 Research questions . . . 2
1.4 Research methodology . . . 2
1.5 Thesis outline . . . 3

2 Traction power systems and the RET Calandlijn case study **4**
2.1 Introduction . . . 4
2.2 DC light-rail traction architecture . . . 4
2.2.1 Topology and supply . . . 4
2.2.2 Voltage-driven grid interactions . . . 6
2.3 The RET Calandlijn . . . 6
2.3.1 Substations and feeding clusters . . . 6
2.3.2 Feeder geometry . . . 8
2.3.3 Rolling stock and the traction load . . . 8
2.4 Operating envelope and limits . . . 9
2.4.1 Voltage limits . . . 9
2.4.2 Current limits . . . 10
2.4.3 Duration-dependent thermal limits . . . 10
2.4.4 Utility intake limits . . . 10
2.4.5 Voltage and current profiles . . . 11
2.5 Summary . . . 12

3 Literature review **13**
3.1 Introduction . . . 13
3.2 Capacity assessment . . . 14
3.2.1 (Traction) power systems: current practice in load allocation . . . 14
3.2.2 Capacity assessment under uncertainty in distribution grids . . . 14
3.2.3 Capacity assessment through optimisation . . . 15
3.3 Scenario analysis . . . 18
3.3.1 Scenario generation . . . 19
3.3.2 Scenario reduction . . . 19
3.4 Research gap and positioning . . . 22
3.4.1 Relevance metrics . . . 22
3.4.2 Gap identification . . . 22

4 Design methodology **24**
4.1 The analysis pipeline . . . 25
4.2 Scenario generation methodology . . . 26
4.2.1 Modelling requirements . . . 27
4.2.2 Power flow solver . . . 27
4.2.3 Definition and span of a scenario . . . 27
4.3 Scenario reduction methodology . . . 28
4.4 Grid margin estimation methodology . . . 28
4.5 Validation methodology . . . 32
4.6 Data flow pipeline . . . 33
4.7 Conclusion . . . 33

5 Traction grid modelling and scenario generation **35**
5.1 Introduction . . . 35
5.2 Co-simulation framework . . . 35
5.3 Vehicle model . . . 37
5.3.1 Vehicle dynamics . . . 37
5.3.2 Train modelling . . . 38
5.3.3 Co-simulation interface . . . 38
5.4 LVDC traction feeder simulator . . . 39
5.4.1 Operational pipeline . . . 39
5.4.2 Power flow problem definition . . . 41
5.4.3 Newton-Raphson solver with active-set switching . . . 42
5.4.4 Gauss-Seidel baseline . . . 44
5.4.5 Comparative validation . . . 45
5.5 Distribution grid model . . . 46
5.6 Scenario generation procedure . . . 47
5.6.1 Recorded data products per scenario . . . 47
5.6.2 Generation procedure . . . 49
5.7 Simulation results . . . 49
5.8 Conclusion . . . 49

6 Traction power system scenario reduction **52**
6.1 Introduction . . . 52
6.1.1 Selected methods for implementation . . . 53
6.1.2 Feature selection - Distribution driven (PCA + k-medoids) . . . 53
6.1.3 Optimality-aware scenario reduction . . . 54
6.2 Reduced set evaluation . . . 56
6.3 Conclusion . . . 57

7 Traction network optimisation **58**
7.1 The optimisation problem . . . 58
7.2 Objective formulations evaluated . . . 58
7.3 Empirical constraint handling under uncertainty . . . 59
7.4 NSGA-III implementation . . . 59

7.5 SMPSO implementation 60
7.6 Conclusion 60

8 Results 62
8.1 Introduction 62
8.2 Scenario reduction. 62
8.2.1 Power allocation - quantitative comparison 62
8.2.2 Validation results 64
8.2.3 Utility feeder 64
8.2.4 Verdict: scenario reduction 64
8.3 Optimisers 66
8.3.1 Power allocation - quantitative comparison 66
8.3.2 Pareto front tap setting comparison 66
8.3.3 Validation metrics per optimiser. 67
8.3.4 Verdict: optimisers 68
8.4 Risk analysis 69
8.4.1 Power allocation - confidence levels 69
8.4.2 Utility limits 69
8.4.3 Grid voltage stability analysis 70
8.4.4 Violation assessment. 71
8.4.5 Verdict: risk level 71
8.5 Spatial distribution of load allocation 73
8.6 Technical feasibility 73
8.7 Conclusion 74
8.7.1 Objective findings 74
8.7.2 Interpretation and implications 77
8.8 Limitations 78

9 Conclusion 80
9.1 Problem and objectives 80
9.2 Answering the research questions 80
9.3 Recommendations. 82

A RET Calandlijn technical specifications 91
A.1 Rolling stock 91
A.2 Traction grid specifications 91
A.3 Normative references 93
A.4 Constraint violation results: Scenario reduction, optimiser choice, confidence level 95
A.4.1 Scenario reduction 95
A.4.2 Optimiser choice 99
A.4.3 Confidence level 101

1 Introduction

1.1. Problem statement

Urban areas concentrate the demand for electricity in modern energy systems, accounting for roughly 75 % of global energy consumption. [1]. This figure is expected to rise, partly due to the ongoing electrification of road transport and space heating as the two largest non-electric energy sinks in most developed economies, now shifting onto the electrical network [2]. In the Netherlands, the combination has produced a structural congestion crisis: national peak demand is projected to rise from approximately 19 GW today to 27 GW by 2030, requiring an estimated EUR 200 billion in grid investment through 2040 [3], [4]. Reinforcement is already lagging local need; the distribution grid covering Rotterdam-Zuid and surrounding municipalities reached maximum capacity in November 2024, with no reinforcement expected before 2029 [5]. As of March 2025, more than 12 000 large consumers and 7 500 generation projects nationwide were waitlisted for grid connections, unable to expand or initiate operations until physical reinforcement of the network is completed [6], [7].

Against this backdrop, the urban rail traction power system operated by Rotterdamse Elektrische Tram (RET) represents an underexplored asset: dimensioned for peak vehicle loads that occur only a small fraction of operational time, the system possesses latent capacity that could in principle support external loads without requiring reinforcement of the surrounding distribution grid.

1.2. Technical challenges

Traction power supply systems (TPSS), powering light-rail (trams and subway) in dense urban areas are already being employed for EV-charging [8] [9] as these allow for flexible charging schemes. However, several key challenges exist in identifying the uninterrupted hosting capacity of the traction power system for external loads, denoted as a constant power load (CPL or P_{CPL}). In contrast to finding scenario-based hosting margins based on load forecasting, this thesis aims to determine the guaranteed connection capacity available to third-party users. Contrary to distribution power networks and LVDC microgrids, the power flow in traction grids is very dynamic and nonlinear as it is subjected to various control actions and state variables such as braking resistors and voltage-dependent vehicle load shedding [10]. This makes it hard to find (quasi) steady-state solutions for the tolerable margins at which the system can potentially be loaded.

The power demand of trains subjects the traction power system to fluctuating, spatially-dependent loading, where the magnitude and location of demand vary continuously as vehicles accelerate, coast, and brake along the network. It is therefore important that the operational conditions are represented and characterised according to a realistic distribution generated by exercising the network model over an extended period of operation. Lastly, the loading limits imposed on traction grid components are duration-dependent [11] [12], which reduces the effectiveness of snapshot-based evaluation of loading limits.

The combination of these properties rules out the simplifications typically applied in capacity assessment. A worst-case static bound that pairs minimum voltage with maximum simultaneous train load underestimates available capacity by a wide margin because such an operating point is improbable and short-lived; allocating an additional load against it produces conservatively small values that defeat the purpose of the exercise.

1.3. Research questions

The technical challenges identified in the previous section motivate the central question of this thesis:

> *Under realistic stochastic operating conditions, how much capacity can the RET DC light-rail traction power system make available to external loads, and how should that capacity be allocated across substations?*

Based on the key technical challenges mentioned in section 1.2 the main question is decomposed into three sub-questions pertaining to characterising traction power systems (scenario generation and reduction) and determining the P_{CPL} limits.

> ***RQ1.*** *How can the state of a traction grid be quantified and represented in terms of the energetic quantities that reflect on the operability of the traction power supply system?*

> ***RQ2.*** *Given that it is infeasible to determine load hosting capacity for every possible traction network state, which finite set of representative operating scenarios can be constructed to approximate the system's stochastic behaviour and then be used as input to downstream hosting capacity computations?*

> ***RQ3.*** *What is the load hosting capacity of the traction power system under uncertainty, and what does does the resulting allocation reveal about the governing limit(s)?*

1.4. Research methodology

First, an overview of traction power suppy systems as well of their characteristics is presented in chapter 2. This also covers the configuration of the system which is the subject of this case study. The operating envelopes, traction power demand, and component limits are introduced, which form the scope of the analysis. Chapter 3 investigates the current state of stochastic power forecasting, power hosting capacity within distribution and traction grids. This forms the basis of the design of a model, which is described in chapter 4 and implemented in the subsequent chapters. The design objectives formulated in chapter 4 are based on the research questions and the corresponding gaps identified in the literature survey.

1.5. Thesis outline

Following the research questions introduced in Section 1.3, relevant system information is presented in chapter 2 as this sets up the scope for the project. Then, the literature review and the research gap pertaining to the research questions will be established in chapter 3, after which the methodology, defined in chapter 4 discusses the actual research process of this thesis. The design and implementation of the developed models are presented across three implementation chapters: chapter 5 addresses RQ1 through the generation of representative operational scenarios; chapter 6 addresses RQ2 by reducing the resulting scenario space to a tractable subset for downstream use; and chapter 7 addresses RQ3 by formulating and solving the capacity allocation problem on the reduced scenario set. Lastly, chapter 9 synthesises the findings against the research questions, discusses the limitations of the work, and proposes directions for future research.

2

Traction power systems and the RET Calandlijn case study

2.1. Introduction

This chapter establishes the technical object of study. It first describes the *system background*, explaining how a DC light-rail traction power system is constructed and how it differs from a conventional distribution network. It then introduces the *loads*, which are determined by moving vehicles rather than by static demand.

The chapter also defines the *limits* that bound admissible operation. These limits form a voltage, current, thermal, and contractual envelope. Finally, it presents the *voltage and current profiles* that arise from the interaction between these loads and operating limits.

The chapter is organised accordingly. Section 2.2 describes the generic architecture of a DC light-rail supply and the voltage-driven interactions that make its loading behaviour distinctive. section 2.3 introduces the RET Calandlijn as the case study, including the substation register, the feeder layout, and the rolling stock that generates the load. Section 2.4 states the operating envelope, the voltage, current, duration-dependent thermal, and contractual limits that bound the network state, and shows the measured and simulated profiles against which those limits act. This envelope is the constraint set carried into the modelling and optimisation of the subsequent chapters.

2.2. DC light-rail traction architecture

2.2.1. Topology and supply

Traction power supply systems are built in a range of configurations, distinguished by voltage level, by AC or DC excitation, and by hybrid combinations of the two. Urban light rail in Europe is most commonly a low voltage DC system at a nominal $600\,\mathrm{V}$ or $750\,\mathrm{V}$ on the contact line, supplied from a medium-voltage AC distribution network (typically $10\,\mathrm{kV}$ to $30\,\mathrm{kV}$) through rectifier substations that interface the two grids. The configuration considered in this thesis is a nominal $750\,\mathrm{V}$ DC supply fed from a $10\,\mathrm{kV}$ AC distribution network.

The contact line, an overhead catenary or a third rail, is fed by a series of rectifier substations connected bilaterally, so that each feeder section is fed from substations on both sides. Bilateral feeding limits conductor voltage drop and provides redundancy during N$-$1 conditions. On the AC side the substations are grouped into feeding clusters, each supplied from its own utility connection and electrically isolated from neighbouring clusters by open tie-breakers in normal operation, so an AC-side fault stays contained within a single cluster. On the DC side the picture is different: the bilaterally fed contact line couples adjacent substations into a lightly meshed network, so power can flow between substations and, where clusters meet, between clusters. The circuit is closed through the running rails. Because this return current can leak into the surrounding soil and corrode buried infrastructure, the return path is bonded and, with rare exceptions, not sectionalised, as a stray-current mitigation measure required by EN 50122-1 [13]. Figure 2.1 shows the electrical representation for a two-substation section.

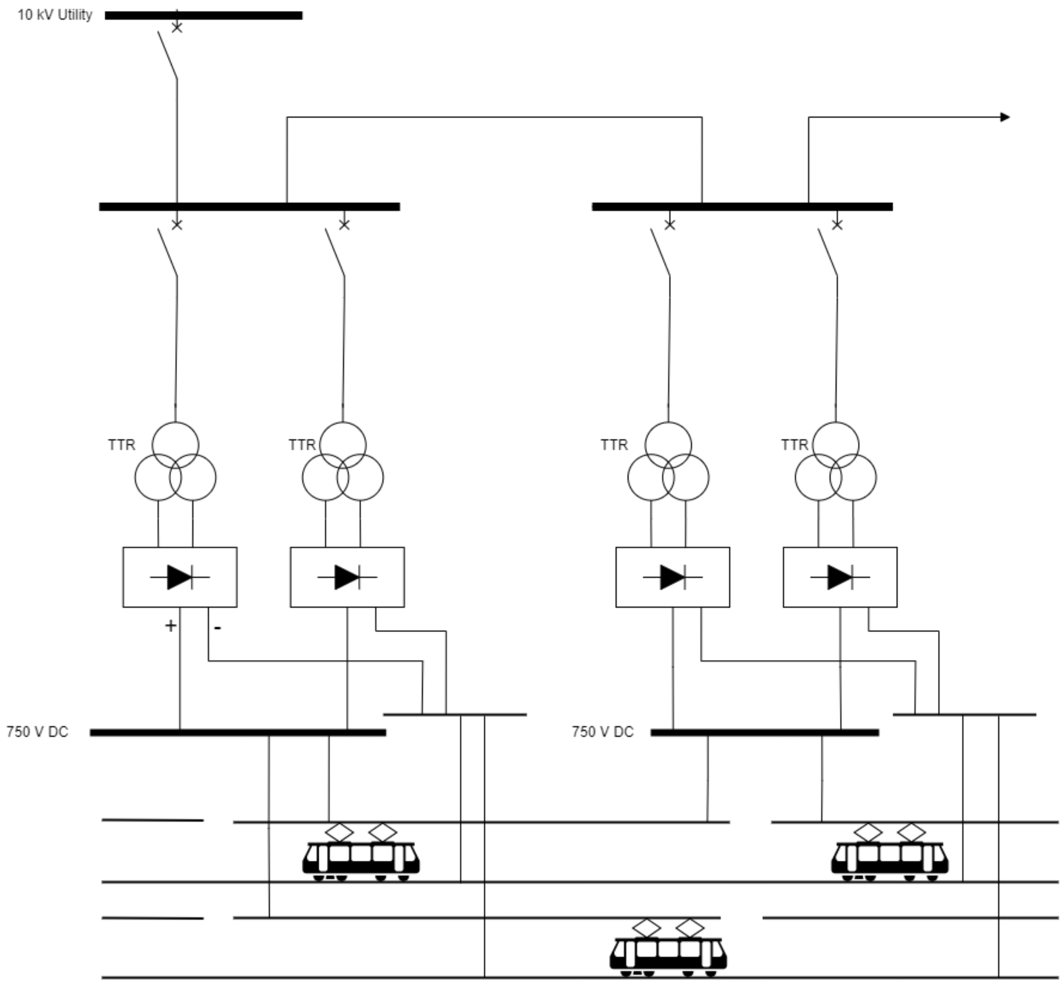

Figure 2.1: Simplified schematic of a two-substation $750\,\mathrm{V}$ DC traction supply section. Each substation interfaces the medium-voltage AC distribution ($10\,\mathrm{kV}$) to the DC contact line through a rectifier transformer and a unidirectional diode rectifier. Because the rectifier blocks reverse current, regenerative braking energy cannot return to the AC side and must be absorbed by other vehicles or dissipated on board. The contact line is fed bilaterally from both substations, with the running rails forming the common return path.

The rectifier sets the substation output voltage. Its no-load DC output is above the nominal contact-line voltage, approximately $850\,\mathrm{V}$ in the model used here, and is adjustable through the rectifier-transformer tap changer over five discrete positions spanning $\pm 5\,\%$ in $2.5\,\%$ steps. The tap position is one of the decision variables in the later optimisation. It suffices to note that the substation operating voltage is set by the rectifier and is not a free network quantity.

2.2.2. Voltage-driven grid interactions

The DC traction grid has no controllable generation. The instantaneous power flow is set entirely by the positions and demands of the vehicles, and the grid mediates the balance through its node voltages. Because the rectifier substations are unidirectional, they can only source power to the contact line and cannot absorb it. Regenerative braking energy produced by a decelerating vehicle must therefore be taken up by another vehicle motoring on the same connected section, or dissipated on board.

When no motoring vehicle is available to absorb it, the braking energy raises the local contact-line voltage. As the voltage rises above the rectifier output level the nearest substation's diodes become reverse-biased and stop conducting, so the substation is effectively disconnected from the network. Whether a substation conducts is therefore not fixed but set at run time by the prevailing voltage, and the active network topology shifts with the load distribution. Should the voltage continue past the upper operating limit, the on-board chopper switches a resistor across the DC link, dissipating the surplus as heat and clamping the voltage at its ceiling; the recovered energy is then lost. At the other extreme, heavy simultaneous demand, particularly from vehicles electrically distant from a substation, depresses the contact-line voltage, and below the lower limit the rolling stock curtails its own traction demand to prevent stalling. The grid voltage thus acts as the control signal for substation blocking, braking dissipation, and load shedding alike.

The demand seen at a substation is correspondingly variable: brief peaks during coincident acceleration, near-zero intervals when braking energy is reused locally, and a time-averaged loading well below the instantaneous peak. This high peak-to-average ratio is visible in the substation load trace of Figure 2.2. Together with the voltage-dependent switching of choppers and the blocking of substations, it is what separates the traction grid from a static distribution feeder and is the reason the loading must be treated statistically rather than by a single design value.

2.3. The RET Calandlijn

The Calandlijn is one of the principal heavy-metro lines of the Rotterdamse Elektrische Tram (RET) network, running through the Rotterdam region. The line is electrified at $750\,\mathrm{V}$ DC throughout and is supplied by twelve rectifier substations grouped into four feeding clusters, each cluster supplied from an independent utility connection on the AC distribution side. The electrical structure, substation ordering, cluster membership, feeder links, and utility intakes, is shown schematically in Figure 2.3.

2.3.1. Substations and feeding clusters

The twelve substations and their assignment to feeding clusters are listed in Table A.2. The installed transformer capacity totals $26.22\,\mathrm{MV\,A}$ and is markedly heterogeneous: ratings range from $1.44\,\mathrm{MV\,A}$ at the smaller single-transformer stations to $2.88\,\mathrm{MV\,A}$ at the stations with two parallel transformers. This heterogeneity matters for the allocation problem, because a unit of external load placed at a small $1.44\,\mathrm{MV\,A}$ station consumes a much larger fraction of local headroom than the same load at a $2.88\,\mathrm{MV\,A}$ station.

The four clusters define the islanding boundaries on the AC side. MSMCP supplies the

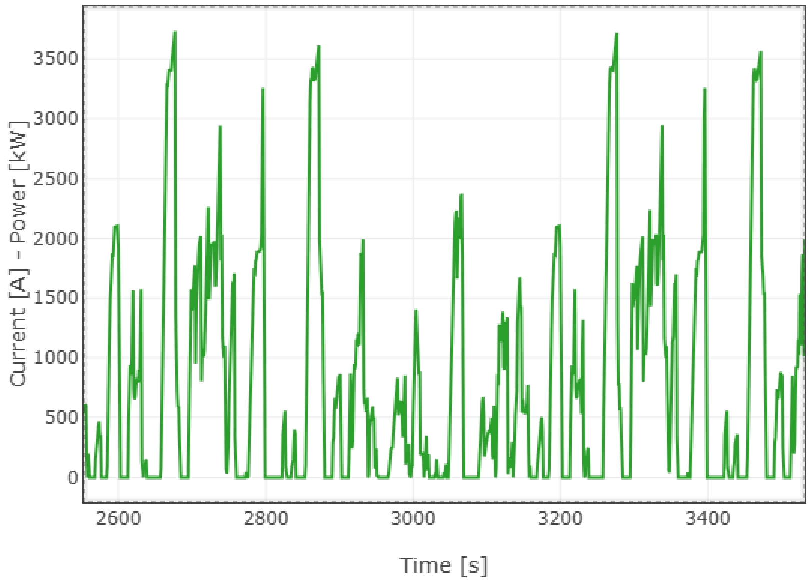

Figure 2.2: Simulated power demand at substation GDW over a representative window. The load is dominated by short motoring peaks separated by near-idle intervals, giving traction load its characteristic high peak-to-average ratio. Simulation results are provided by KRUCH Energy Flow Simulations.

western section, GSGDW the central section, GSTRP the single branch substation, and GSNSL the eastern section and terminus. The cluster boundaries are open in normal operation, so AC-side faults remain contained within a cluster; on the DC side, however, the clusters are coupled through the bilaterally fed contact line, so a regenerative event in one cluster can support motoring demand in an adjacent cluster, subject to the voltage-drop limit on the intervening feeder.[1]

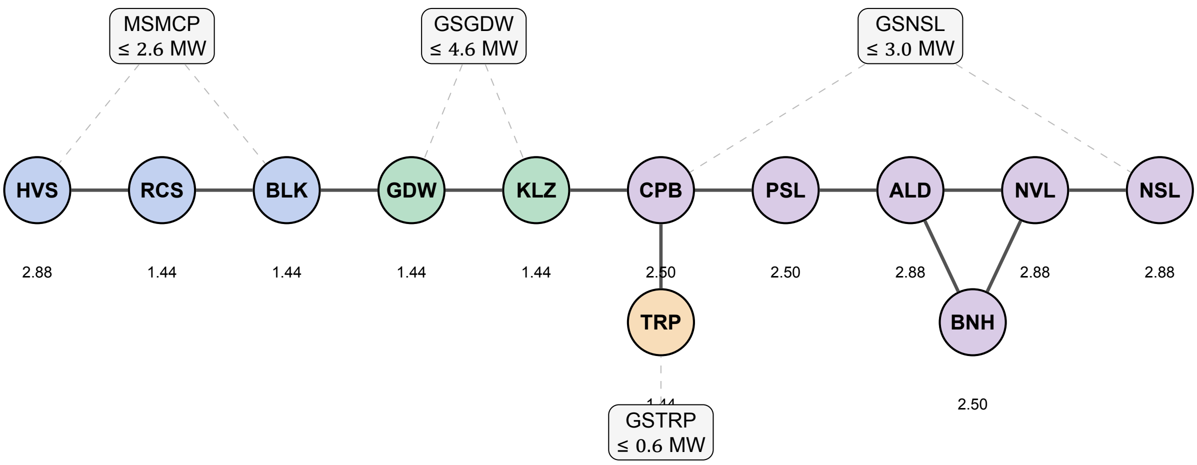

Figure 2.3: Electrical structure of the Calandlijn. Circles are rectifier substations (label inside, rated capacity in MVA below), coloured by feeding cluster; solid lines are the $750\,\mathrm{V}$ DC contact line; dashed lines are the AC utility intakes, annotated with the $15\,\mathrm{min}$ rolling contractual cap of each cluster (each intake additionally has a $10\,\mathrm{MW}$ instantaneous cap). TRP is the single substation on the GSTRP branch; BNH sits on the ALD–NVL loop.

[1]The substation-to-cluster assignment used throughout this thesis is the AC intake mapping; it is approximate at the two backbone connections (MSMCP and GSGDW), where the AC solve resolves small redistributions and resistive losses of a few percent that the snapshot aggregation ignores.

Table 2.1: Calandlijn substation register: anchor station, rated transformer capacity, and feeding cluster. Ratings are apparent-power transformer ratings; stations marked $2\times$ have two parallel transformers.

Substation	**Anchor station**	**Rated capacity [MVA]**	**Feeding cluster**
HVS	DHV	$2 \times 1.44 = 2.88$	MSMCP
RCS	EDP	1.44	MSMCP
BLK	BLK	1.44	MSMCP
GDW	GDW	1.44	GSGDW
KLZ	KLZ	1.44	GSGDW
TRP	TRP	1.44	GSTRP
CPB	CPB	2.50	GSNSL
PSL	PSL	2.50	GSNSL
ALD	ALD	$2 \times 1.44 = 2.88$	GSNSL
BNH	BNH	2.50	GSNSL
NVL	NVL	$2 \times 1.44 = 2.88$	GSNSL
NSL	NSL	$2 \times 1.44 = 2.88$	GSNSL
Total installed		**26.22**	

2.3.2. Feeder geometry

The substations are connected by bilateral DC feeders that allow power to flow in either direction along the contact line depending on the instantaneous load distribution. Two conductor types are present: a low-resistance third-rail conductor on the western and central trunk, and a higher-resistance overhead conductor (denoted *bvl*) on the eastern section. Feeder lengths range from a few hundred metres between closely spaced trunk substations to roughly $3\,\mathrm{km}$ on the CPB–TRP branch, and these lengths, with the per-metre conductor resistances, set the inter-substation coupling that governs how load and regenerative energy redistribute across the network. The resulting conductances enter the node-admittance matrix of the DC power-flow solver developed in chapter 4.

2.3.3. Rolling stock and the traction load

The Calandlijn is mainly operated with Bombardier SG3 vehicles, bidirectional light-rail units of approximately $45\,\mathrm{m}$ length operating singly or in coupled pairs. The datasheet for these vehicles can be found in A.1 in the appendix.

The load character follows directly from these parameters. During accelerationa vehicle draws a short, high power peak; while coasting it draws only the auxiliary load. While braking above $5\,\mathrm{m\,s^{-1}}$, the vehicle returns power to the contact line at the efficiency given in Table A.1, with the on-board chopper absorbing whatever cannot be taken up by another vehicle. The available power of the rolling stock is limited by the pantograph voltage; a sag below 625 V reduces the available traction power of the vehicle as a measure of grid stability security. A measured power-and-voltage trace at the pantograph is shown in Figure 2.4, and illustrates both the peaky motoring demand and the voltage excursions that accompany it.

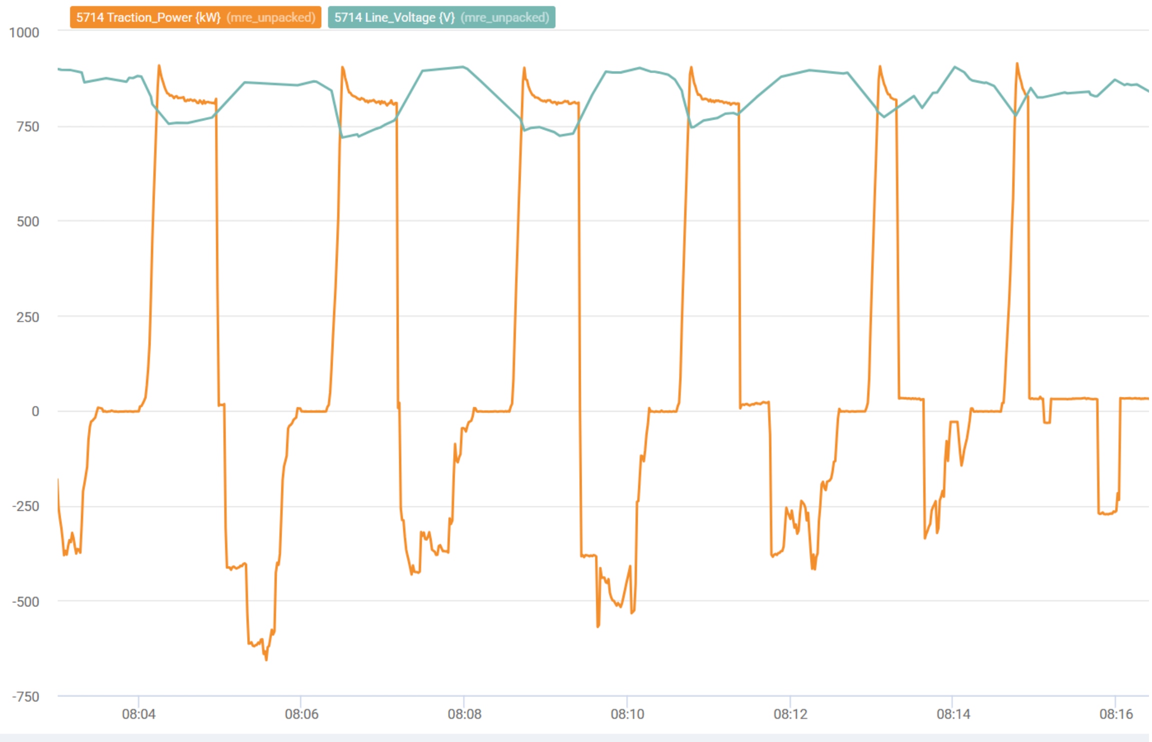

Figure 2.4: Measured power demand and pantograph voltage at vehicle level. Positive power is motoring (drawn from the contact line); negative power is regenerative braking returned to it. The voltage rises during braking and dips during coincident motoring. Courtesy of Railfleet One.

2.4. Operating envelope and limits

The admissible operating envelope is bounded by voltage, current, thermal, and contractual limits, set by international standards and by the utility supply contract. Table 2.2 collects the limits used in the power-flow and optimisation models of the subsequent chapters, with the source of each. The remainder of this section discusses each group and the constraint that the optimisation enforces from it.

2.4.1. Voltage limits

The contact-line voltage band is defined by EN 50163 [14] (whose IEC counterpart is IEC 60850). For a nominal $750\,\mathrm{V}$ DC system the standard sets a permanent maximum of $900\,\mathrm{V}$, a non-permanent maximum of $1000\,\mathrm{V}$ admissible only for short transients, and a permanent lower limit of $500\,\mathrm{V}$. The $500\,\mathrm{V}$ floor at the vehicle is the value below which sustained operation is no longer assured and the rolling stock curtails traction to avoid stalling; it is the lower voltage limit enforced in this thesis.

The rolling-stock interface adds a stricter criterion. NEN-EN 50388 [15] requires a minimum *mean useful voltage* at the pantograph, defined as a moving mean over a specified window rather than an instantaneous floor; the value used here is $625\,\mathrm{V}$.[2] This criterion guards against sustained under-supply that brief excursions below $500\,\mathrm{V}$ would not trigger.

[2]The mean-useful-voltage figure is the project value and should be confirmed against the applicable edition of NEN-EN 50388 for the $750\,\mathrm{V}$ DC class.

Table 2.2: Operating envelope of the Calandlijn $750\,\mathrm{V}$ traction power system as enforced in the model.

Quantity	Limit used in the model	Source
Contact-line voltage	$U_{\mathrm{nom}} = 750\,\mathrm{V}$; $U_{\mathrm{max,perm}} = 900\,\mathrm{V}$; $U_{\mathrm{max,non\text{-}perm}} = 1000\,\mathrm{V}$	EN 50163 [14]
Vehicle voltage (enforced floor)	$U_{\mathrm{veh}} \geq 500\,\mathrm{V}$	EN 50163 [14]
Pantograph mean useful voltage	$U_{\mathrm{mean,useful}} \geq 625\,\mathrm{V}$	NEN-EN 50388 [15]
Substation current	$\lvert I\rvert \leq 10\,\mathrm{kA}$; reverse current blocked	EN 50123 [16]; project data
Transformer loading	1 pu continuous; ≤ 2.5 pu on a $1\,\mathrm{min}$ rolling average; full duration curve per Figure 2.5	IEC 60076-12 [17]
Utility intake (per cluster)	$\leq 10\,\mathrm{MW}$ instantaneous; $15\,\mathrm{min}$ rolling cap per cluster (Table 2.3)	RET / Stedin supply contract

2.4.2. Current limits

On the current side the binding limit is the substation current, capped at $10\,\mathrm{kA}$; the unidirectional rectifier additionally blocks any reverse current toward the AC side. The DC switchgear and busbar ratings, together with their short-circuit withstand and insulation levels, follow the EN 50123 series [16].

2.4.3. Duration-dependent thermal limits

The transformer and rectifier components are not characterised by a single power limit but by a load-duration curve. Figure 2.5 shows the admissible loading of a Class VI dry-type transformer under IEC 60076-12 [17], expressed as a loading ratio $\mathrm{LR} = S/S_{\mathrm{rated}}$ against overload duration. Nominal loading (1 pu) is sustainable indefinitely, while short overloads are admissible for limited times: up to 2.5 pu for one minute and roughly 1.23 pu for sixty minutes.

The relevant capacity statement for such an asset is therefore joint with a time duration: a single deterministic power figure does not capture the envelope the standard defines. This thesis enforces the limit through a $1\,\mathrm{min}$ rolling loading ratio held to 2.5 pu. Because the limit is a probabilistic statement about a duration-averaged quantity, it motivates the chance-constrained treatment developed in chapter 4, where the duration-dependent limit is enforced at a chosen confidence level rather than as a hard worst-case bound.

2.4.4. Utility intake limits

Each feeding cluster draws from the AC distribution network through a single utility connection, and the supply contract bounds that intake in two ways: an instantaneous hard cap of $10\,\mathrm{MW}$ per connection, and a sustained limit expressed as a $15\,\mathrm{min}$ rolling-average cap that differs by cluster. The per-cluster rolling caps are listed in Table 2.3. The clusters differ sharply: the single-substation GSTRP branch is held to $0.6\,\mathrm{MW}$ while the central GSGDW connec-

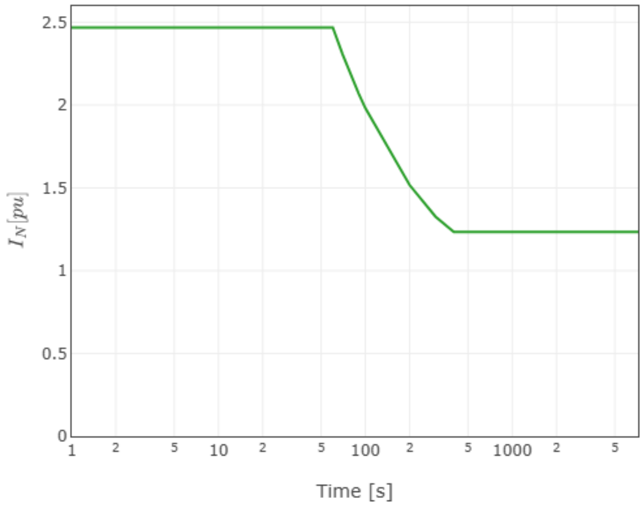

Figure 2.5: Class VI dry-type traction-transformer duration-dependent overload limit under IEC 60076-12: admissible loading ratio against overload duration, flat at 2.5 pu up to about one minute and decaying to roughly 1.23 pu beyond one hour.

tion allows $4.6\,\mathrm{MW}$. As the results chapters show, this $15\,\mathrm{min}$ rolling cap, rather than the instantaneous limit or the transformer thermal limit, is the constraint that most often binds the achievable external-load allocation on the GSNSL and GSTRP clusters.

Table 2.3: Per-cluster utility intake limits. The instantaneous cap applies to every cluster; the $15\,\mathrm{min}$ rolling cap is cluster-specific.

Cluster	15-min rolling cap [MW]	Instantaneous cap [MW]	Substations
MSMCP	2.6	10	HVS, RCS, BLK
GSGDW	4.6	10	GDW, KLZ
GSTRP	0.6	10	TRP
GSNSL	3.0	10	CPB, PSL, ALD, BNH, NVL, NSL

2.4.5. Voltage and current profiles

The envelope above acts on profiles that are anything but flat. At the vehicle level (Figure 2.4) the pantograph voltage swings with the local power balance: it dips during coincident motoring, when several vehicles draw from the same section, and rises during braking, occasionally to the chopper clamp. At the substation level (Figure 2.2) the power trace is dominated by short motoring peaks separated by near-idle intervals, so the instantaneous peak sits well above the time-averaged loading. The substation current follows the same shape, since it is essentially the power divided by the slowly varying busbar voltage, and it is the brief coincident-motoring peaks that approach the $10\,\mathrm{kA}$ current cap and the short-duration end of the transformer loading curve (Figure 2.5). These profiles are why the limits cannot be checked against a single design point and why a population of operating snapshots, rather than one worst case, is needed to certify available capacity.

The undervoltage function does not trip the supply directly, since brief excursions below

$500\,\mathrm{V}$ are tolerable in practice; instead, sustained undervoltage raises an alarm and triggers load shedding at the rolling-stock level.

2.5. Summary

This chapter has set out the system background, the loads, the limits, and the profiles that define the case study. The Calandlijn is a $750\,\mathrm{V}$ DC light-rail network of twelve heterogeneous rectifier substations ($26.22\,\mathrm{MVA}$ installed) in four AC-islanded but DC-coupled clusters, carrying a load that is set by moving vehicles and is therefore peaky, spatially distributed, and bidirectional through regenerative braking. Its admissible states are bounded by an envelope that is partly standard-derived (the EN 50163 voltage band, the NEN-EN 50388 mean useful voltage, the EN 50123 current ratings, the IEC 60076-12 duration-dependent transformer loading) and partly contractual (the per-cluster $15\,\mathrm{min}$ utility intake caps). Two features of this envelope shape the rest of the thesis: the thermal limit is duration-dependent rather than a single power value, and the binding limit on external load is most often the contractual utility cap rather than a hardware rating. Both make a static worst-case assessment inadequate and call for the stochastic, scenario-based modelling and the chance-constrained optimisation developed in the following chapters.

3

Literature review

3.1. Introduction

The technical challenges identified in section 1.2 and the operating envelope established in chapter 2 together define a specific class of problem: the allocation of continuous load on the existing capacity of a DC light-rail traction power system, under duration-dependent thermal limits and stochastic vehicle traffic. This chapter surveys the literature relevant to that problem. The aim is not an exhaustive treatment of railway electrification or power system optimisation, but a characterisation of the state of practice in the areas that bear directly on the research questions of section 1.3.

The review follows the chain that the problem itself imposes. The starting point is *capacity assessment*. Capacity assessment is done under a finite representation of the operating-condition distribution. Section 3.3 therefore reviews *scenario analysis*: how a large scenario set is reduced to a smaller, representative subset without distorting the underlying probablistic distribution of worst-case scenarios, how scenarios are generated in the first place, and what power flow solver the generation step rests on.

> ***A scenario in this thesis is defined as a snapshot of the traction power system with corresponding state variables. These state variables include power flow quantities among which instantaneous vehicle demands, node voltages, branch currents, and substation loading (instantaneous and window-averaged).***

Each section opens by motivating its content from the one before and closes with a short statement of what the literature leaves open. Section 3.4 establishes the research gap, identified from a matrix of relevance metrics derived from the research questions, and positions the thesis at the intersection the matrix leaves empty.

Related topics in railway electrification such as dynamic stability, harmonic analysis, electromagnetic compatibility, and signalling integration are real areas of research but are not covered. The broader literature on distribution-grid optimisation is surveyed only where its methods transfer to the traction-network context.

3.2. Capacity assessment

Assessing how much external load a traction network can host is a capacity-allocation question, and capacity allocation can be posed as a constrained optimisation: maximise the load placed on the network subject to its technical and contractual limits under uncertain operation. This section builds toward that formulation. It first examines how traction practice has so far treated capacity on existing infrastructure, then introduces the capacity-assessment framework that the distribution-grid community has matured under the name of hosting capacity, and finally reviews the optimisation schemes (multi-objective formulation, chance constraints, and stochastic load flow).

3.2.1. (Traction) power systems: current practice in load allocation

Utilising traction grids for connecting smart loads is a well-established research direction, with dedicated work on EV charging integration into traction infrastructure [11], [18]–[20]. These studies focus on controlled power flow, typically through EV charging schemes or FACTS devices coupled to the traction substations, in order to achieve secondary objectives such as improved power quality, energy recovery, or loss minimisation. Closer to the capacity question, Diab et al. [21] propose a prediction scheme for surplus capacity estimation based on a probabilistic simulation of traffic on a single line between neighbouring substations, and Smith et al. [18] characterise power and thirty-minute energy margins of a traction grid from actual measurements, constrained by both technical and contractual limits.

The framework of Diab et al. [21] produces margin estimates but its case study is limited to the segment between two neighbouring substations; how surplus capacity should be distributed when multiple substations and clusters interact is not posed. The measurement-based work of Smith et al. [18] is retrospective, characterising the margin that existed under observed conditions rather than the margin that could be allocated under stochastic future conditions. The chance-constrained sizing work of Chen et al. [22], methodologically adjacent, treats surplus capacity as an input parameter to an energy-storage sizing problem rather than as the unknown to be computed. Pierrou and Hug [23] embed EV charging into a railway-station energy management system but assume station-level capacity is given, and the flexible-substation work of [24] treats the substation as a controllable asset for upstream grid services without addressing the underlying capacity question. Where capacity-related questions have been studied, then, the framing has been narrow in geographical scope, retrospective in time, or oriented toward sizing new equipment rather than allocating capacity on existing infrastructure. The systematic allocation of continuous external load across the substations of a real network, under stochastic traffic and duration-dependent thermal limits, is not addressed in the surveyed traction-specific work.

3.2.2. Capacity assessment under uncertainty in distribution grids

The probabilistic framing that traction practice lacks is, in the distribution-grid literature, mature. The most analogous question to traction-network capacity allocation is renewable / distributed generation (DG) hosting capacity: the maximum level of distributed generation or load that can be connected without causing operational limit violations. The concept originated in the context of high penetration of distributed photovoltaic generation [25] and has developed into an analytical framework that addresses some of problems this thesis confronts.

Hosting capacity methods can be classified along two axes: the treatment of uncertainty (deterministic worst-case, stochastic time-series, or probabilistic) and the computational strategy (Monte Carlo simulation, analytic probabilistic load flow, or optimisation-based) [26]–[28]. The worst-case approach evaluates the most adversarial combination of load and generation conditions; it is computationally cheap but notoriously conservative, systematically underestimating available capacity in much the same way as the static worst-case bound described in section 1.2. Stochastic time-series methods sample profiles from empirical distributions and process them sequentially; they preserve temporal correlation but require many samples to characterise worst case scenario probability [26]. Probabilistic methods derive the distributions of the limiting variables analytically, often via linearised sensitivity analysis, allowing the hosting capacity to be reported as a quantile of an admissible-injection (load connection, in the case of this project) distribution [27], [28]. Its application to the traction power system would involve gathering substation-specific measurements, as well as a

Recent work has extended these formulations to multi-dimensional risk assessment. The framework of [29] aggregates security, reliability, and economic risk indicators into a composite risk index and computes both an optimal hosting capacity and a maximum acceptable hosting capacity under a defined risk threshold. The gap between the two exceeds $30\,\%$ on the IEEE 33-node test system, illustrating the practical magnitude of the conservatism introduced by single-criterion deterministic methods. The spatiotemporal voltage sensitivity analysis of [27] computes the probability of voltage limit violations at any bus given random generation location and magnitude, providing a probabilistic foundation for capacity assessment that does not rely on exhaustive Monte Carlo simulation.

This literature supplies a strong methodological template, but its assumptions do not transfer directly to traction networks. Hosting-capacity work is typically AC, balanced, and oriented toward distributed generation rather than distributed load. The limiting constraint is most often voltage rise from reverse power flow rather than the duration-dependent thermal loading that governs traction substations, the topology is radial rather than dense and meshed, and the primary uncertainty source is aggregated renewable generation rather than spatially coincident vehicle motoring. The quantile-based capacity statement and the risk-weighted trade-off are nonetheless the right principles for the traction case, and the route to operationalising them is an explicit optimisation, treated next.

3.2.3. Capacity assessment through optimisation

The load allocation problem is to maximise the external load placed on the network subject to its technical, contractual, and probabilistic constraints, evaluated through a power flow over the operating envelope.

Optimal power flow formulations A defining feature of the traction capacity allocation problem is that the substations of a real network are not interchangeable: they differ in transformer rating, cluster membership on the AC side, feeder geometry, and the local distribution of vehicle traffic. These differences manifest as spatial heterogeneity in the available capacity and in the binding constraints. The spatiotemporal correlation analysis of Gao et al. [30] shows explicitly that spatial dependence between substations must be preserved in any reduced representation, since aggregating it away produces over-representation of certain operating regimes. Operationally, a CPL increment at one substation shifts load not only locally but also onto its DC-coupled neighbours through the bilateral feeders. The effect is asymmetric across the network because the cluster topology and the substation ratings are themselves asymmetric.

The methodological response to this heterogeneity is a multi-objective formulation. A single aggregate objective, for example the sum of CPL values across all substations, is mathematically simpler but implicitly assigns equal weight to every kilowatt regardless of where it is allocated. In practice, this concentrates the allocation on the substations with the largest local headroom, irrespective of the stress imposed on neighbours through DC coupling. Multi-objective formulations make the spatial trade-offs explicit by treating per-substation (or per-cluster) loading as distinct objectives, allowing the nondominated (Pareto) front to expose the underlying conflicts, including the trade-off between aggregate capacity and network-level loading balance.

The most commonly considered objectives in traction-network and distribution-grid OPF are summarised in Table 3.1. Despite the methodological maturity of multi-objective OPF in the

Table 3.1: Objectives considered in optimal power flow formulations, with representative traction-network instances where available.

Objective	Description	Ref.
Energy cost minimisation	Classical OPF objective; in traction, minimising the cost of returned braking energy.	[31]
Transmission loss minimisation	Primary or secondary objective in distribution-grid OPF.	[32]
Voltage quality	Profile flatness, unbalance, or sag mitigation.	[22]
Peak shaving / load balancing	Reduce substation loading variance; limits transformer thermal ageing.	[32]
Renewable utilisation	Maximise absorption of local PV or storage output.	[33]
Reliability / resilience	Quantify contingency impact on supplied load.	[33]

broader literature, its application to traction networks remains uncommon. The dominant pattern in traction OPF is single-objective: minimisation of recuperated energy cost [31], minimisation of utility intake, or maximisation of regenerative absorption. Where multiple objectives have been considered, the formulations have typically aggregated them via fixed weightings rather than navigating the Pareto front explicitly. A fully multi-objective formulation that treats per-substation CPL allocations as distinct objectives, and exposes their spatial trade-offs through the Pareto front, has not been documented for traction-network capacity allocation in the surveyed literature.

Multi-objective evolutionary algorithms The traction capacity allocation problem combines non-convexity[1] from the underlying power flow with mixed-integer decision variables (tap settings, integer power allocation quantities in the kW scale). Typically, mathematical gradient search methods and relaxations are used in optimal power flow, in which the decision variables are mapped towards the objective functions, of which the optimal points are found through iteration [32]. As the optimisation problem considered in this thesis is constraint-dominated[2], these methods are less suited to the determination of the load hosting limits. The natural alternative is metaheuristic search, which treats the computation of the objective functions and its constraints as a black box evaluator, and samples the decision space under selection pressure: Babiker et al. [32] report metaheuristics as the dominant approach for nonconvex OPF, and Martinez Fernández et al. [34] find evolutionary algorithms the most widely used family in railway optimisation specifically.

Two sub-families are relevant: non-dominated-sorting evolutionary algorithms and swarm-based particle methods. Their defining mechanisms are compared in Table 3.2.

Table 3.2: Comparison of the two multi-objective metaheuristic families considered for the capacity allocation problem.

	NSGA-III [35]	SMPSO [36]
Search paradigm	Population with generational offspring	Particle swarm with external archive
Offspring / update	Binary tournament, simulated binary crossover, polynomial mutation	Velocity update with archive-drawn leader
Diversity mechanism	Reference directions (Das–Dennis [37]), replacing the crowding distance of NSGA-II [38]	Leader sampled inversely to local archive density
Constraint handling	Constraint domination [39]: feasible over infeasible, infeasible ranked by aggregate violation	Velocity clipping to curb divergent search
Many-objective behaviour	Remains discriminative as objective dimension grows	No explicit niching; prone to premature convergence on dominant objectives
Implementation	More involved	Simpler, trivially vectorisable

Performance in the many-objective literature is typically assessed with the hypervolume metric [40], which captures both convergence to the Pareto front and diversity along it.

Chance-constrained programming Constraint formulations under uncertainty fall into three principal categories. Deterministic worst-case formulations require every constraint to hold

[1] A not smoothly converging objective function, in this case, the power balance residuals in the traction power flow, which is subject to control actions (substation blocking, regenerative braking resistance, vehicle load curtailment)

[2] The aim in this study is to find the limits as to how much capacity can be connected to the grid under some degree of violation prevalence

under every realisation of the uncertainty and are inherently conservative. Robust optimisation relaxes the worst-case requirement to an uncertainty set chosen by the modeller, trading conservatism for tractability. Chance-constrained programming, introduced by Charnes and Cooper [41], replaces the deterministic statement $g(x, \xi) \leq 0$ ($g > 0$ represents a constraint violation) with the probabilistic statement $\mathbb{P}[g(x, \xi) \leq 0] \geq 1 - \alpha$, admitting a small fraction α of operating conditions in which the constraint may be violated. The Conditional Value-at-Risk (CVaR) formulation of Rockafellar and Uryasev [42] generalises this to the case in which the engineering cost of a violation depends on its magnitude.

Chance-constrained formulations have been applied to power-system problems with stochastic loads or generation, including in traction networks. Chen et al. [22] use a chance-constrained formulation to size storage and power-flow controllers for railway smart grids under uncertain traction load, treating voltage unbalance probabilistically and showing that the probabilistic envelope yields less conservative sizing than the deterministic worst case. The application is limited to the sizing problem with a single primary objective; the use of chance constraints inside a multi-objective formulation for capacity allocation has not been ocumented in the surveyed work.

Stochastic load flow A chance constraint is only meaningful relative to a distribution of operating states, which is what stochastic load flow provides. Stochastic power flow methods extend the deterministic power flow problem to admit random inputs and produce as output a probability distribution over network state variables rather than a single operating point. Ju et al. [43] present a stochastic dynamic framework for power system analysis under uncertain variability, integrating analytical and Monte Carlo techniques with fast convergence on medium-sized networks. Analytical methods such as the point-estimate method, cumulant methods, and polynomial chaos expansions trade exactness for computational efficiency, while Monte Carlo approaches retain distributional fidelity at the cost of many repeated solves. The vast majority of this work targets distribution grids with stochastic renewable injections rather than traction networks, where the dynamic topology induced by vehicle motion and the discrete state of chopper and rectifier elements present additional challenges that the standard formulations do not accommodate.

Casting capacity assessment as a multi-objective, chance-constrained optimal power flow therefore presupposes a finite set of operating scenarios over which the probabilistic constraints are evaluated. How that scenario set is constructed is the subject of the next section.

3.3. Scenario analysis

The demand variability as well as its dependence on travel patterns represent the operating-condition distribution. Two questions follow.

1. How are the scenarios generated?

2. Given a large set of snapshots, how is it reduced to a subset small enough to optimise over while retaining the probability characteristics of worst-case scenarios?

3.3.1. Scenario generation

TPSS power flow methods As noted above, simulation-based scenario generation warrants closer analysis, since the simulator can represent additions to the power system that historical data cannot. It comprises two components:

1. travel-demand traction profiles, which simulate vehicle power demand and location;
2. the power flow solver, which evaluates the resulting traction demand on the infrastructure.

The DC traction power flow problem itself differs from the conventional distribution power flow problem in three respects: the system is nonlinear due to chopper dynamics and unidirectional rectifier behaviour; the network topology is dense and meshed rather than radial; and the operating point depends on the discrete state of multiple control elements (chopper activation, substation conduction) that must be resolved consistently with the continuous power balance. Power flow methods relevant to this problem can be classified into several families [31], [44]: Newton-Raphson based methods, Gauss-Seidel iterative techniques, forward-backward sweep methods, physics-informed and other artificial neural network approaches [45], and data-driven surrogate models. Table 3.3 summarises the principal advantages and limitations of each family in the DC traction context, with the suitability of each method for scenario-generation use foregrounded.

Cai et al. [47] introduced iterative formulations for DC traction power flow with regenerative braking, comparing conductance, current, voltage, and power balanced formulations. The voltage iterative method was identified as the most robust for convergence, but the comparison was conducted from a static-simulation perspective without addressing the scenario-throughput requirements of an optimisation pipeline. Li et al. [31] develop a fast quasi-optimal power flow specifically for flexible DC traction systems, achieving rapid convergence by superimposing the natural traction load with the controlled response of voltage-source converters at neighbouring substations; the formulation, however, omits several nonlinearities relevant to the present case study and targets minimisation of recuperated energy cost as a single objective. Neither solver has been benchmarked under repeated execution over thousands of scenarios, which is precisely the regime statistical validation of an allocation result requires.

3.3.2. Scenario reduction

Scenario reduction is a practice originating from stochastic programming (e.g. optimisation under uncertainty) in which the probability distribution of variables produces too many combinations of possible state configurations for numeric analysis [48]. Scenario reduction is increasingly applied in the stochastic penetration of renewables in the grid, as this method can effectively deal with multivariate uncertainty [49]–[52], which is somewhat analogous to the stochastic load patterns in traction power systems.

The goal of this method is to pick some collection of scenarios from the existing set or distribution, while maintaining the representativeness with respect to the original. This representativeness is based on a similarity metric, per the rubric of [50]. The measure of representativeness and the selection method characterises the scenario reduction method, 4 methods are shortly compared in Table 3.4.

Table 3.3: Comparison of solution methods for the DC traction power-flow problem, as reported in the literature.

Method	**Advantages**	**Limitations**
Gauss-Seidel	Robust per node evaluation, easy to encode chopper logic, substation unidirectionality and voltage limits.	Slow linear convergence
Sequential Newton-type solver	Better local nonlinear correction than Gauss-Seidel; integrates naturally with traction-specific switching rules.	Not fully coupled; may require damping and validity checks to avoid non-physical branches.
Full Newton-Raphson	Fully coupled formulation; rigorous Jacobian-based treatment of nonlinear interactions.	Higher implementation and per-iteration cost; active-set switching is still required.
Forward-backward sweep	Usually robust for radial feeders; avoids full Jacobian assembly.	Per node evaluation / sensitive to oscillations.
Artificial neural networks [45]	Fast online evaluation after training; favourable scaling for large networks.	Reliable supervised training data are not generally available; physical consistency is hard to guarantee.
Other data-driven surrogates [46]	Potentially fast approximate prediction once trained.	Reliable supervised training data are not available; extrapolation across operating modes is unreliable.
Probabilistic / Monte Carlo methods	Useful for stochastic fleet-level analysis and energy-absorption studies.	Not a stand-alone solver; requires many repeated inner solves.

Probability-distance methods

The most established works in scenario reduction [48], [53] minimise the Kantorovich distance between the full and reduced distributions. The Kantorovich distance, equivalent to the Wasserstein metric, measures the minimum cost of transporting one probability distribution onto another, so minimising it selects a reduced set whose distribution most closely matches the original This can be formulated as:

$$W_p(P,Q) = \left(\min_{\pi_{\omega,\omega_s} \geq 0} \sum_{\omega \in \Omega} \sum_{\omega_s \in \Omega_s} \pi_{\omega,\omega_s}\, d(\xi_\omega, \xi_{\omega_s})^p \right)^{1/p}$$

$$\sum_{\omega_s \in \Omega_s} \pi_{\omega,\omega_s} = p_\omega, \quad \omega \in \Omega,$$

$$\sum_{\omega\in\Omega} \pi_{\omega,\omega_s} = q_{\omega_s}, \quad \omega_s \in \Omega_s.$$

Here, Ω is the full scenario set, while Ω_s is the reduced scenario set. The original distribution P assigns probability p_ω to each scenario $\omega \in \Omega$, and the reduced distribution Q assigns probability q_{ω_s} to each retained scenario $\omega_s \in \Omega_s$. The variable π_{ω,ω_s} represents the amount of probability mass transported from full scenario ω to reduced scenario ω_s. Finally, $d(\xi_\omega, \xi_{\omega_s})$ measures the distance between the corresponding scenario values, and p is the order of the Wasserstein metric, commonly chosen as $p = 1$ or $p = 2$.

The probability distributions refer to the specific state variables of a scenario.

Clustering methods represent scenarios by medoids. Medoids preserve feasibility by returning actual scenarios, whereas centroids yield biased objective estimates [54], [55]. Dynamic time warping [56] and PCA (principal component analysis) preprocessing [57] adapt clustering to the phase-shifted, low-rank structure of train profiles. The reduction remains blind to the downstream objective. **Moment matching** [58] produces a compact summary but under-represents worst-case scenarios, which is disqualifying for a quantile-driven margin objective [59]. **Problem-driven methods** [50], [60] couple the reduction to the optimiser and target optimality directly. The criterion extends to the multi-objective case through a hypervolume-ratio gap,

$$\mathrm{gap}_{HV}(\Omega_s) = 1 - \frac{HV(\mathcal{P}^*(\Omega_s))}{HV(\mathcal{P}^*(\Omega))}, \tag{3.1}$$

where $\mathcal{P}^*(\cdot)$ is the Pareto front on the respective scenario set. Each outer iteration requires a full multi-objective solve, which is the main cost of this method.

Table 3.4: Comparison of scenario reduction families for the multi-objective traction capacity allocation problem.

Family	**Strengths**	**Weaknesses**
Kantorovich FFS/SBR [48], [53]	Strong theoretical guarantees; Lipschitz bound on optimality gap; preserves actual scenarios	Objective-agnostic; tail scenarios under-represented; scales poorly with N and dimension
Clustering [56], [57], [61]	Returns feasible scenarios; robust to outliers; metric-flexible; exploits temporal structure	Sensitive to random initialisation; objective-blind; choice of K non-trivial
Moment matching [58]	Compact and interpretable; preserves covariance structure	Matched moments do not imply matched distributions; poor tail fidelity
Problem-driven / decision-based [50], [60]	Targets optimality directly; best in-sample to out-of-sample fidelity; extensible to hypervolume-gap for MOPF	Expensive outer loop (repeated MOPF solves); multi-objective extension non-standard

The literature presents a structural trade-off. Distribution-driven methods preserve the empirical distribution but are blind to the optimisation they support. Problem-driven methods address optimality directly but require an outer loop that is incompatible with a decoupled reduction-and-optimisation pipeline. No surveyed method occupies the middle ground, and none has been documented for a traction-network capacity allocation problem.

3.4. Research gap and positioning

The preceding sections each closed on something the literature leaves open. Rather than restating those points as separate threads, this section consolidates them through a small set of relevance metrics derived from the research questions, scores the surveyed literature against those metrics in a single matrix, reads the gap off the matrix, and positions the thesis at the intersection it leaves empty.

3.4.1. Relevance metrics

The main research question of section 1.3 asks how the traction network can supply external loads while adhering to operational and contractual limits, decomposed into the representation of the operating state (RQ1), the construction of a reduced scenario subset (RQ2), and the achievable hosting capacity together with the nature of its dominating constraint (RQ3). Four metrics capture the themes a work must address to bear directly on these questions.

Spatial-aware hosting capacity. Whether the work quantifies allocatable capacity with spatial resolution across multiple nodes, rather than as a single aggregate figure or a single-site margin. This is the object of the main question and of RQ3.

Traction power-system application. Whether the method is developed for, or applied to, a traction or railway power system, as opposed to a generic AC distribution grid. This grounds the main question in the case-study network.

Uncertainty / stochastic scenario representation. Whether operating uncertainty is represented through stochastic scenarios or probabilistic load flow, the prerequisite for RQ1 and RQ2.

Power-system reliability. Whether the work explicitly evaluates adherence to operational or contractual limits, the reliability dimension of the main question and of the binding-constraint part of RQ3.

Table 3.5 scores the surveyed literature against the four metrics. Broad reviews are excluded, since they report the field rather than contribute a position within it.

3.4.2. Gap identification

The matrix shows a clear separation. The traction works occupy the upper block: they reliably satisfy the application metric and, in the more recent cases, the uncertainty metric, but they treat capacity either as a single-site margin or as a given input, so spatial-aware hosting capacity is at most partial and reliability is addressed only incidentally. The hosting-capacity works occupy the middle block: they satisfy the spatial-aware capacity, uncertainty, and (often) reliability metrics, but none is set in a traction context, and their AC, generation-oriented, voltage-limited assumptions do not transfer to a meshed DC network bound by duration-dependent thermal limits.

Table 3.5: Positioning of the surveyed literature against the four themes most relevant to this thesis. ✓ marks a theme addressed by the work, and ∗ indicates partial or implicit treatment.

Reference	Spatial-aware hosting capacity	Traction power system application	Uncertainty / stochastic scenarios	Power-system reliability
Diab et al. [21]	∗	✓	∗	
Smith et al. [18]	∗	✓		
Chen et al. [22]	∗	✓	✓	∗
Pierrou and Hug [23]		✓		
Chen et al. [24]		✓	✓	✓
Li et al. [31]		✓		
Cheng et al. [62]		✓		✓
Gao et al. [30]		✓	✓	
Bollen and Hassan [25]	✓		∗	∗
Qammar et al. [26]	✓		✓	∗
Munikoti et al. [27]	✓		✓	✓
Chihota et al. [28]	✓		✓	
Liu et al. [29]	✓		✓	✓
Ju et al. [43]			✓	✓
Zhuang et al. [50]			✓	
Wafaa and Dessaint [63]	∗		✓	✓
This work	✓	✓	✓	✓

No surveyed work occupies the intersection of all four metrics. The gap is therefore the probabilistic, spatially resolved assessment of hosting capacity for a real traction power system, in which operating uncertainty is represented through scenarios and the result is held to explicit operational and contractual limits. Reaching that intersection means more than transplanting hosting-capacity methods: it requires a simulation-based scenario generator and an optimality-aware yet decoupled scenario reduction (neither of which the scenario literature supplies for traction networks), feeding a per-substation multi-objective formulation under chance-constrained, duration-dependent thermal limits (a combination undocumented in the traction OPF literature).

The specific methodological commitments responsive to the gap, and the chapters that implement them, are stated in chapter 4.

4

Design methodology

The preceding chapters have done three things: posed the research questions in section 1.3 that the thesis sets out to answer, established in chapter 2 the technical context within which those questions are situated, and identified in section 3.4 the research gap in the literature that the thesis must close in order to provide credible answers. The present chapter translates that combination of questions and gaps into a set of *design objectives* that the methodology must satisfy, and introduces the analysis pipeline as the structural response.

A design objective, as used here, is a requirement on the methodology expressed in declarative form. Each design objective traces back to a specific research question (which it operationalises) and to one or more gap threads (which it addresses). The four design objectives of this thesis are stated below; the rest of the chapter develops the methodological commitments that meet them.

DO1: Produce a representative scenario set. The methodology shall generate a set of operational scenarios of sufficient fidelity to capture the nonlinear and stochastic phenomena identified in section 1.2, and at a computational cost that admits its use as input to a multi-scenario optimisation. The scenario set must span the operating envelope defined in section 2.4 rather than only the historically observed conditions. This objective operationalises RQ1 and responds to the capacity-allocation and probabilistic-assessment aspects of the research gap (section 3.4).

DO2: Reduce the scenario set in an optimality-aware manner. The methodology shall produce a reduced subset of the full scenario set on which the downstream optimisation can be conducted, and shall do so in a way that preserves both the distributional structure of the operating-condition set and the optimisation-relevant structure of the scenarios that determine the binding constraints. This objective operationalises RQ2 and responds to the optimality-aware scenario-reduction aspect of the research gap (section 3.4). The combination of distributional fidelity and optimality-awareness is itself a methodological contribution and is supported by comparison against an established distribution-driven baseline.

DO3: Formulate capacity allocation as a multi-objective chance-constrained optimisation. The methodology shall pose the P_{CPL} allocation question as a constrained optimisation programme with per-substation decision variables, multiple objectives that preserve the spatial heterogeneity of the network, and a probabilistic treatment of the duration-dependent thermal constraints implicit in the operating envelope. This objective operationalises RQ3 and responds to the probabilistic-assessment and multi-objective chance-constrained aspects of the research gap (section 3.4).

DO4: Validate the optimised allocations against dynamic feasibility. The methodology shall provide an internal validation step in which selected allocations from the Pareto front are tested for chronological feasibility under the dynamic co-simulation, thereby distinguishing solutions that are statically feasible but dynamically inadmissible. This objective operationalises the second half of RQ2 (generalisation under dynamic evaluation) and contributes to the binding-constraint identification required by RQ3.

The four design objectives together define the requirements the methodology must satisfy. The remainder of this chapter introduces the analysis pipeline that responds to these requirements and develops the commitments at each pipeline stage.

4.1. The analysis pipeline

The analysis pipeline is the structural response to the four design objectives. It comprises four stages, each of which produces a specific data product consumed by the next stage: scenario generation, scenario reduction, optimisation, and validation. Figure 4.1 shows the pipeline schematically.

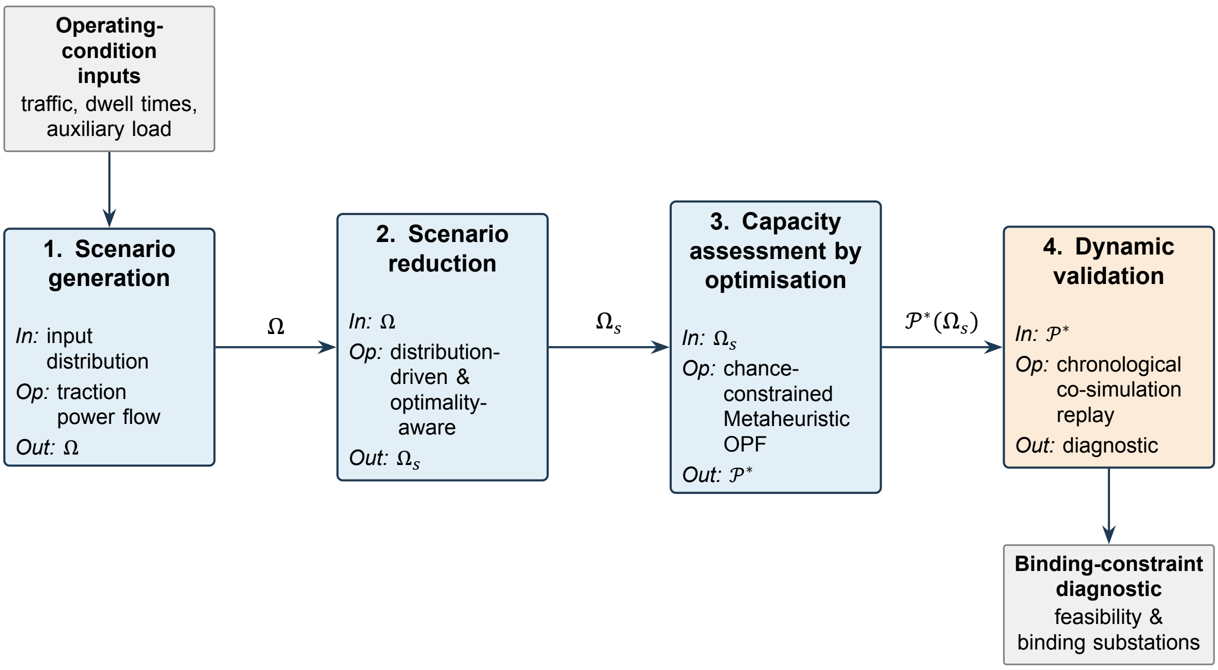

Figure 4.1: The analysis pipeline of this thesis. Each stage produces a data product consumed by the subsequent stage: scenario generation produces the full operating-condition set Ω; scenario reduction produces a tractable subset $\Omega_s \subset \Omega$; optimisation produces the Pareto front $\mathcal{P}^*(\Omega_s)$. The validation stage (in amber) operates on the optimised allocations and closes the loop against the dynamic co-simulation, producing a feasibility and binding-constraint diagnostic that informs the discussion rather than feeding the next pipeline stage.

The first stage, *scenario generation* (chapter 5), produces the full operating-condition scenario set Ω by exercising a co-simulation model of the network under a chosen distribution of inputs. The set is large by construction: each scenario is a per-second snapshot of network state, generated across a representative operating window. Modelling and scenario generation are treated jointly at this stage, since the model exists in service of the scenario-generation task; the model is not an end in itself.

The second stage, *scenario reduction* (chapter 6), compresses Ω into a tractable subset $\Omega_s \subset \Omega$ that retains both the distributional structure of the full set and the optimisation-relevant features of its margin-critical scenarios. The reduced set is the input to the downstream optimisation. Two reduction methods are implemented in parallel: a distribution-driven baseline grounded in the probability-distance and clustering literature surveyed in subsection 3.3.2, and an optimality-aware variant that introduces the contribution of this thesis to the reduction step.

The third stage, *optimisation* (chapter 7), operates on Ω_s and produces a Pareto front $\mathcal{P}^*(\Omega_s)$ of P_{CPL} allocations. The optimisation is multi-objective and chance-constrained, with the specific formulations developed in chapter 7. Several formulations are evaluated, since the literature in subsection 3.2.3 provides no canonical multi-objective treatment of traction-network capacity allocation and several formulation choices are open.

The fourth stage, *validation* (chapter 8), selects representative points from the Pareto front and replays them through the dynamic co-simulation as introduced by DO1. This stage tests whether allocations that are statically optimal on Ω_s remain feasible under chronological operation, and identifies the binding constraints that govern the extent of the Pareto front.

The data products flow strictly downstream: $\Omega \rightarrow \Omega_s \rightarrow \mathcal{P}^*(\Omega_s) \rightarrow$ validation results. The pipeline does not iterate between the optimisation and reduction stages, which distinguishes it from the fully problem-driven scenario reduction schemes of, for example, Zhuang et al. [50], whose method requires repeated MOOPF solves inside an outer reduction loop and is incompatible with a decoupled pipeline. The trade-off this choice implies, and the optimality-aware reduction method developed to mitigate it, are discussed in section 4.3.

4.2. Scenario generation methodology

The scenario generation stage is governed by DO1: produce a representative scenario set on a timescale that meets the peak-load and energy-management requirements of NEN50388 (Figure A.1) and captures the nonlinear and stochastic phenomena identified in section 1.2. Meeting these requirements means resolving each scenario at per-second resolution, for overcurrent evaluation, over a span of several hours, for energy management. Evaluating one complete scenario set on this basis requires on the order of 10^5 power flow solves, which the proprietary software used by RET completes in 10 to 20 hours.

The simulator is not only a scenario generator. It is also the evaluator applied to every candidate solution during optimisation, so its throughput determines whether the optimisation is practical at all. The evaluation count of the optimisation loop is the product of its factors: 12 combinations of design variables, a reduced scenario set on the order of 10^3 (extracted from a full set of order 10^4), a population of 100 candidates, tens of generations, and 2 power flow solves per candidate per scenario. This yields a power flow solve count in the order of $10^7 - 10^8$. The theoretical computational effort required by RET's energy flow simulation software can take upwards of years based on the normal simulation computational effort.

Such a runtime is disqualifying on two grounds. It defeats the purpose of the project, and it outlives the assumption that the network infrastructure remains unchanged over the course of the solve. A faster power flow evaluation is therefore a precondition for the rest of the methodology, not a refinement of it.

4.2.1. Modelling requirements

The use of simulation-based generation places specific requirements on the model itself. First, the model must represent the subsystems of the traction network at the fidelity required to capture the technical phenomena of section 1.2: the chopper-rectifier nonlinearity of the DC side, the spatially correlated vehicle traffic, and the electrical coupling between the DC traction grid and the upstream AC distribution. Second, the model must be modular, so that the vehicle, DC feeder, and AC distribution subsystems can be developed, verified, and parameterised independently. Third, the model must be fast enough to produce the required scenario throughput within the total computational budget of the analysis.

These requirements lead to the choice of a co-simulation paradigm rather than a monolithic simulation. A co-simulation orchestrator couples the three subsystems through a defined data-exchange interface at a chosen synchronisation timescale, preserving their independent solver structure while enforcing consistency at each exchange step. Modular co-simulation is established practice for multi-physics power-system problems and meets all three of the requirements above; alternatives, such as integrated solvers implemented in proprietary environments [64], were considered but rejected on tractability grounds.

4.2.2. Power flow solver

The principal computational cost of scenario generation lies in the DC power flow solve at each scenario timestep. The power-flow families surveyed in subsection 3.3.1 are evaluated against the dual requirements of robustness (convergence across the operating envelope, including the discrete state changes induced by chopper activation and rectifier blocking) and throughput (per-solve wall-clock cost). A Newton-Raphson formulation with active-set switching for the discrete elements is adopted: it provides the fully coupled treatment of the nonlinear interactions documented by Cai et al. [47] and Li et al. [31], while admitting the explicit active-set handling required by the chopper and rectifier topology. A Gauss-Seidel implementation is retained as a verification baseline, not as a competing approach. Detailed solver development is deferred to chapter 5.

4.2.3. Definition and span of a scenario

A scenario in this thesis is a snapshot of the network's state under a specific instantaneous configuration of vehicle positions, vehicle power demands, and substation operating points, together with the resulting solution of the DC power flow. Scenarios are generated at a temporal resolution of 1 second across an operating window chosen to be representative of weekday service.

4.3. Scenario reduction methodology

The scenario reduction stage is governed by DO2: compress the full scenario set Ω into a tractable subset Ω_s that preserves both distributional fidelity and optimisation-relevant structure. The literature in subsection 3.3.2 establishes a trade-off between these two principles: distribution-driven methods (probability-distance, clustering, moment matching) preserve the empirical distribution but are blind to the optimisation; problem-driven methods address optimality directly but require an iterative coupling between reduction and optimisation that is incompatible with the decoupled pipeline architecture introduced in section 4.1.

The methodology of this thesis responds to the trade-off by implementing *two* reduction methods in parallel, comparing their behaviour on the same downstream optimisation, and reporting which structural properties of the reduced set affect the optimisation outcome.

The first method is a *distribution-driven baseline* grounded in the clustering family surveyed in subsection 3.3.2. A principal-component preprocessing step reduces the dimensionality of the scenario feature space, followed by a k-medoids clustering that returns physically realisable scenarios (rather than synthetic centroids) as the representatives of the reduced set. The choice of k-medoids over k-means follows the feasibility-preservation argument established in the literature: the optimiser evaluates the power flow on each scenario, so each scenario must be a realisable operating point. The selection of K, the number of representatives, is governed by a Wasserstein-based criterion that quantifies the distributional fidelity of the reduced set against the full set.

The second method is an *optimality-aware variant* introduced by this thesis as a relaxation of full problem-driven reduction. Each scenario in Ω is characterised not by its raw feature vector but by its per-scenario P_{CPL} headroom, computed through a fast static evaluation that does not require iteration with the optimiser. Scenarios are then clustered in this optimality space rather than in the raw feature space, producing a reduced set in which margin-critical scenarios are represented in proportion to their optimisation relevance. The method preserves the decoupled pipeline architecture without calling the expensive multi-objective optimiser, while introducing optimisation-awareness into the reduction step. The relationship between the optimality vector and the full problem-driven framework of, for example, Zhuang et al. [50], is developed in chapter 6.

The comparison between the two methods is itself a contribution. Both methods produce reduced sets that are subsequently fed through the same optimisation pipeline; the difference in the resulting Pareto fronts, and in the binding constraints those fronts expose, is reported in chapter 8. The implementation of both methods, the choice of K, the clustering procedures, and the feature engineering decisions are developed in chapter 6.

4.4. Grid margin estimation methodology

The determination of the power allocation limit is governed by DO3: formulate P_{CPL} allocation as a constrained optimisation programme with per-substation decision variables, multiple objectives that preserve the spatial heterogeneity of the network, and a probabilistic treatment of duration-dependent thermal constraints.

The literature surveyed in subsection 3.2.3 establishes that fully multi-objective formulations are uncommon in traction-network optimisation, and that the small body of existing work has used fixed-weighting aggregations rather than navigating the Pareto front directly. There is, accordingly, no single canonical formulation that the thesis can adopt: the formulation itself is part of the methodological territory the thesis must explore.

Decision variables. The principal decision variables are the per-substation P_{CPL} values, each a continuous quantity bounded below by zero and above by an upper bound consistent with the operating envelope of Table 2.2. Furthermore, substation voltage level tuning by means of tap settings is another decision variable which can potentially benefit the power capacity margin as well as traction substation loading distribution in the system, as per the findings of the study by Diab et al. [9]. Therefore, the tap settings of 12 rectifier stations are added to the decision space.

Objectives The primary objective of this thesis is finding the sizing of external power loads at which the system can reliably transmit power to the vehicles. Hence, the primary objective is defined as that which maximises the power load when connected to the DC bus at the substation level. The argument for a DC-only connection instead of tapping into the 10 kV distribution grid is from an engineering standpoint simpler to implement, as there are no costly high-voltage components to integrate, and is more suitable for potential use cases such as charging electric vehicles.

The main objective is thus formulated in the following manner:

$$f_k^{\mathrm{CPL}}(\mathrm{x}) = P_{\mathrm{CPL},k}, \qquad k = 1, \dots, N_{\mathrm{SS}}, \qquad \text{(maximise)}, \tag{4.1}$$

i.e. one objective per substation rather than a single network-wide total. The decision vector $\mathrm{x} = [\mathrm{P_{CPL}};\ \mathrm{t}]$ collects the N_{SS} continuous load allocations $\mathrm{P_{CPL}} \in [0, \overline{P}_{\mathrm{CPL}}]^{N_{\mathrm{SS}}}$ and the discrete rectifier tap positions $\mathrm{t} \in \mathcal{T}^{N_{\mathrm{SS}}}$; for the Calandlijn $N_{\mathrm{SS}} = 12$. Taps appear in the decision space because they adjust substation voltage and can trade against load allocation, as shown by Diab et al. [9]. The per-substation form preserves the spatial heterogeneity of the network: a scalar aggregate $\sum_k P_{\mathrm{CPL},k}$ would hide which substations absorb the marginal kilowatt, and is retained only as a baseline in chapter 7.

Secondary objectives The per-substation hosting-capacity objectives drive the search towards the boundary of the feasible region. The technical and contractual constraints determine which combinations of $\mathrm{P_{CPL}}$ are admissible; the secondary objectives are therefore not required to prevent the decision variables from reaching their upper bounds. Instead, they distinguish feasible allocations that provide comparable hosting capacity but differ in their operational characteristics.

The implemented formulation retains two secondary objectives: the distribution of loading among substations and the aggregate intake from the upstream utility connections. Network dissipation and chopper losses are recorded separately during the dynamic validation and are therefore treated as evaluation metrics rather than optimisation objectives.

Loading imbalance. An allocation may satisfy all component limits while concentrating the highest loading at only a small number of substations. To quantify this concentration, define the maximum normalised loading of substation k over the reduced scenario set $\mathcal{S}_{\mathrm{R}}$ as

$$\lambda_k^{\max}(\mathrm{x}) = \max_{s\in\mathcal{S}_{\mathrm{R}}} \frac{P_{\mathrm{sub},k}^{(s)}(\mathrm{x})}{S_{\mathrm{rated},k}}. \tag{4.2}$$

The loading-imbalance objective is then

$$f^{\mathrm{imb}}(\mathrm{x}) = \mathrm{Var}_k\left(\lambda_k^{\max}(\mathrm{x})\right), \qquad \text{(minimise)}. \tag{4.3}$$

Because the loading is normalised by the corresponding substation rating, this objective penalises allocations for which the peak utilisation is concentrated at a subset of substations. The metric is interpreted jointly with the capacity objectives: its magnitude also increases with the overall loading level and should therefore not be interpreted as a stand-alone measure of relative fairness.

Utility intake. The Calandlijn is supplied through N_{U} independent AC-side utility connections. Their contractual limits are enforced separately in the constraint set, so a violation at one connection cannot be compensated by unused capacity at another. As a secondary efficiency criterion, the optimisation minimises the total scenario-weighted mean utility intake. First, define the mean intake at connection u as

$$\overline{P}_{\mathrm{util},u}(\mathrm{x}) = \sum_{s\in\mathcal{S}_{\mathrm{R}}} w_s P_{\mathrm{util},u}^{(s)}(\mathrm{x}), \qquad \sum_{s\in\mathcal{S}_{\mathrm{R}}} w_s = 1, \tag{4.4}$$

where w_s is the empirical probability weight assigned to scenario s. The aggregate utility-intake objective is

$$f^{\Sigma\mathrm{util}}(\mathrm{x}) = \sum_{u\in\mathcal{U}} \overline{P}_{\mathrm{util},u}(\mathrm{x}), \qquad \text{(minimise)}. \tag{4.5}$$

This objective exposes the additional upstream power required to support a given hosting allocation, while the connection-specific rolling and instantaneous limits remain enforced as constraints.

Dissipation as a validation metric. Although dissipation is not included in the optimisation objective vector, it is evaluated for the selected Pareto designs during the chronological validation. The reported metric is

$$m^{\mathrm{diss}}(\mathrm{x}) = \sum_{s\in\mathcal{S}_{\mathrm{R}}} w_s \left[\sum_{(ij)\in\mathcal{E}} P_{\mathrm{loss},ij}^{(s)}(\mathrm{x}) + \sum_{i\in\mathcal{G}} P_{\mathrm{chop},i}^{(s)}(\mathrm{x})\right], \tag{4.6}$$

where $\mathcal{E}$ is the set of conductor branches and $\mathcal{G}$ is the set of regenerating vehicle nodes. Reporting this quantity separately permits an assessment of whether an allocation increases conductor losses or improves the utilisation of regenerative braking energy without increasing the dimensionality of the optimisation problem.

Because both optimisation algorithms are implemented in minimisation form, the per-substation capacity objectives are represented as

$$f_k^{\mathrm{CPL}}(\mathrm{x}) = -P_{\mathrm{CPL},k}, \qquad k = 1, \ldots, N_{\mathrm{SS}}. \tag{4.7}$$

The complete objective vector is therefore

$$\mathrm{f}(\mathrm{x}) = \left(f_1^{\mathrm{CPL}}, \ldots, f_{N_{\mathrm{SS}}}^{\mathrm{CPL}}, f^{\mathrm{imb}}, f^{\Sigma\mathrm{util}}\right). \tag{4.8}$$

For the Calandlijn, $N_{\mathrm{SS}} = 12$, so the implemented optimisation contains $N_{\mathrm{SS}}+2 = 14$ objective components.

Constraints. Hard and soft constraints are incorporated as evaluation metrics in the margin assessment. Hard constraints cover the non-permanent voltage limits and feeder current, as listed in Table 4.1, whereas soft constraints permit brief exceedances of the rated thresholds. *Chance-constrained limits* are being used with growing frequency in optimal power flow analyses.The relevance to this traction-specific case is supported by the fact that strictly maintaining predefined limits is not feasible; even in the absence of load disturbances on the network, traction loads can still cause voltage violations, exceed substation loading thresholds, and trigger overcurrent protection relays.

The decision vector x is bounded by a combination of equipment ratings, safety thresholds, and contractual limits. Table 4.1 lists the constraints retained in the optimisation. Each is expressed in its physical, deterministic form; the duration-dependent subset admits a probabilistic relaxation developed in chapter 7.

Table 4.1: Operating constraints on the multi-objective optimisation. State variables on the left-hand side of each inequality are returned by the power-flow solver (chapter 5) under the decision x and operating scenario $s \in \mathcal{S}_{\mathrm{R}}$.

ID	Constraint	Type	Source
g_1	$V_{\min} \leq V_i^{(s)}(\mathrm{x}) \leq V_{\max},\quad \forall\, i \in \mathcal{N}_{\mathrm{DC}},\, s \in \mathcal{S}_{\mathrm{R}}$	hard	EN 50163
g_2	$U_{\mathrm{mean},j}^{(s)}(\mathrm{x}) \geq U_{\min},\quad \forall\, j \in \mathcal{V},\, s$	hard	EN 50388-1
g_3	$\left\lvert I_{ij}^{(s)}(\mathrm{x})\right\rvert \leq I_{\max,ij},\quad \forall\, (ij) \in \mathcal{E},\, s$	hard	EN 50123
g_4	$\overline{\lambda}_k^{(s)}(\mathrm{x};\, \tau_1) \leq \overline{\lambda}_1,\quad \forall\, k \in \mathcal{S}_{\mathrm{SS}},\, s$	duration-dep.	IEC 60076-12
g_5	$\overline{\lambda}_k^{(s)}(\mathrm{x};\, \tau_2) \leq \overline{\lambda}_2,\quad \forall\, k \in \mathcal{S}_{\mathrm{SS}},\, s$	duration-dep.	IEC 60076-12
g_6	$\overline{P}_u^{(s)}(\mathrm{x};\, \Delta t) \leq \overline{P}_u^{\mathrm{cap}},\quad \forall\, u \in \mathcal{U},\, s$	contractual	utility contract

Symbol values for the Calandlijn: $V_{\min} = 500\,\mathrm{V}$, $V_{\max} = 900\,\mathrm{V}$; $U_{\min} = 625\,\mathrm{V}$; $I_{\max,ij}$ branch-specific; $\overline{\lambda}_1 = 2.5$ for $\tau_1 = 1$ min; $\overline{\lambda}_2 = 1.2$ for $\tau_2 = 5$ min; $\overline{P}_u^{\mathrm{cap}}$ per the per-cluster contract over $\Delta t = 15$ min. The sets $\mathcal{N}_{\mathrm{DC}}, \mathcal{S}_{\mathrm{SS}}, \mathcal{V}, \mathcal{E}, \mathcal{U}$ denote DC nodes, substations, vehicles, conductor branches, and utility connections; $\lambda_k = P_{\mathrm{sub},k}/S_{\mathrm{rated},k}$ is the substation loading ratio and $\overline{\lambda}_k(\tau)$ is its moving average over a window of length τ.

g_1 — *DC node voltage band.* The catenary voltage is bounded by the equipment ratings of the rolling stock and by the insulation of the feeder cables. Operation below $V_{\min}$ triggers load curtailment at the vehicles and risks stalling; operation above $V_{\max}$ threatens insulation. The non-permanent ceiling of 1000 V allowed by EN 50163 for short regenerative transients is enforced separately by the on-board chopper inside the model (Eqs. 5.13–5.14) and does not require a constraint here.

g_2 — *Pantograph mean useful voltage.* EN 50388-1 specifies a minimum mean useful voltage at the pantograph, averaged over a defined moving window, of 625 V for 750 V DC systems. The constraint protects the rolling stock from sustained undervoltage that compromises traction performance even when the instantaneous voltage briefly recovers above $V_{\min}$; it is the time-averaged complement to g_1.

g_3 — *Branch ampacity.* Each conductor branch in the DC feeder has a current-carrying capacity determined by its cross-section and thermal environment. Exceeding the ampacity damages the conductor or its insulation through accelerated thermal ageing; unlike the transformer thermal limits below, no overload curve is associated with conductors at the scale of this network, so the constraint is enforced strictly at every scenario. Values are taken from Table 2.1 (feeder cables) and Table 2.4 (busbars and switchgear).

g_4, g_5 — *Transformer / rectifier thermal overload envelope.* Between the continuous rating and the trip threshold, the transformer admits temporary overload according to the duration-dependent profile of Figure 2.4 (IEC 60076-12). The continuous curve is discretised here into two representative windows: a 1-minute moving average bounded at 2.5 p.u. (g_4) and a 5-minute moving average bounded at 1.2 p.u. (g_5). Each average must remain below its corresponding bound. The underlying curve is itself a statistical statement about admissible operation, obtained by integrating an equivalent-ageing criterion over an assumed load-duration profile; consequently, a deterministic enforcement on every operating window is more conservative than the standard requires, and the probabilistic enforcement consistent with the loading guide's intent is developed in chapter 7.

g_6 — *Utility connection cap.* Each of the four AC-side feeding clusters draws from an independent utility connection under a contracted active-power cap evaluated over a 15-minute rolling window. Exceeding the cap incurs a per-kWh penalty under the current tariff.

Solvers and exploration strategy. Two solvers are used. NSGA-III, justified in Table 3.2.3 by its reference-direction-based selection which remains discriminative in high-dimensional objective spaces, is the principal solver. SMPSO is the swarm-based comparator, used to provide independent verification of the Pareto front rather than as a competing solution method. The combination of multiple formulations, two solvers, and two reduced scenario sets produces a structured comparison whose results are reported in chapter 8 and synthesised in chapter 9; the synthesis identifies which formulation choices have measurable consequences for the resulting allocation and which do not.

The implementation of both solvers, the chance-constraint integration into the NSGA-III selection operator, and the specific formulations investigated are developed in chapter 7.

4.5. Validation methodology

The validation stage is governed by DO4: distinguish allocations that are statically optimal on Ω_s from those that remain feasible under chronological operation. The need for this step arises from the distinction between static and dynamic feasibility flagged in RQ2: a multi-objective optimisation conducted over a set of independent scenarios may produce allocations that satisfy every constraint on every individual scenario without satisfying the sequential-state requirements of dynamic operation, in which the network's state at one timestep is the initial condition for the next.

The validation step replays selected Pareto-optimal allocations through the dynamic co-simulation under representative chronological operating windows. A Pareto-optimal allocation is declared *dynamically feasible* if no operational constraint is violated during the replay, and *dynamically infeasible* otherwise. The distinction is itself diagnostic: an allocation that passes the static test but fails the dynamic one reveals a binding constraint that the static formulation has not captured, typically a duration-dependent thermal envelope whose exceedance accumulates across consecutive scenarios in a way the per-scenario chance constraint cannot detect.

External validation against operational SCADA measurements is out of scope for this thesis. The consequences of this exclusion are discussed in chapter 9. The dynamic replay procedure, the selection of allocations from the Pareto front for validation, and the binding-constraint analysis are developed in chapter 8.

4.6. Data flow pipeline

The pipeline can be read in two complementary ways. Figure 4.1 presents it as a sequence of *operations*; Figure 4.2 closes the chapter by presenting it as a sequence of *data products*, making explicit what each stage holds and how each transformation reshapes it. The input parameters are turned by the simulator into a per-snapshot feature vector of vehicle and power-flow quantities; scenario reduction compresses this to a small set of representative vehicle locations and power draws; and optimisation maps that set to the hosting allocation and transformer tap position of each substation.

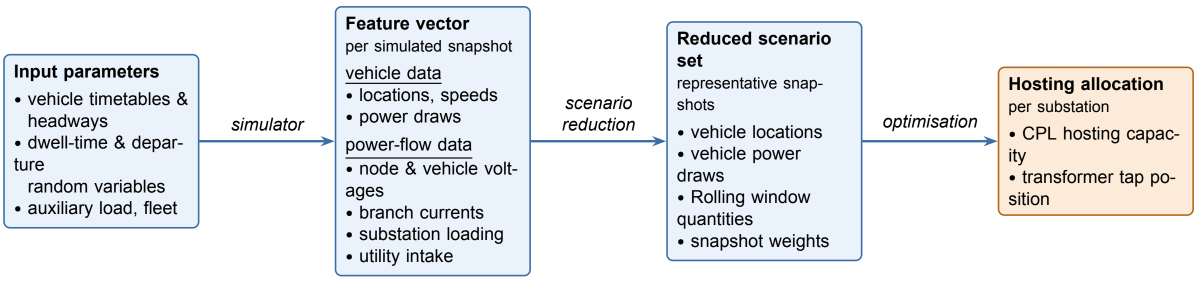

Figure 4.2: Data flow of the analysis pipeline, the complement to the operational view of Figure 4.1. Each container lists the data it holds; the arrows name the transformation that maps one container to the next. The simulator turns the input parameters into a per-snapshot feature vector of vehicle and power-flow quantities; scenario reduction retains only the representative vehicle locations and power draws (with their weights); and optimisation yields the constant-power-load hosting allocation and transformer tap position for each substation.

4.7. Conclusion

This chapter has translated the research questions and the literature gap into a concrete methodology. The capacity-allocation problem of the Calandlijn is framed through four design objectives and answered by a four-stage pipeline: simulation-based scenario generation over the operating envelope, scenario reduction in both a distribution-driven and an optimality-aware form, a per-substation multi-objective optimisation in which the duration-dependent thermal and contractual limits enter as chance constraints, and a dynamic co-simulation replay that separates statically optimal allocations from those that remain feasible under chronological operation. The constraint set is tied directly to the operating envelope of chapter 2, and the data exchanged between stages is made explicit in Figure 4.2.

What this chapter fixes is the structure of the analysis, not its implementation. The solver integration, the embedding of the chance constraints in the NSGA-III selection operator, and the specific objective formulations investigated are developed in chapter 7; the structured comparison they enable, across two solvers, two reduced scenario sets, and several formulations, is reported in chapter 8; and the methodological assumptions, in particular the reliance on simulator-derived scenarios and internal validation, are revisited in chapter 9.

5

Traction grid modelling and scenario generation

5.1. Introduction

This chapter implements Design Objective DO1 of chapter 4 and the scenario-generation methodology committed to in section 4.2: it develops the co-simulation traction-grid model and the procedure by which that model produces the operating-condition scenario set Ω consumed by the downstream scenario reduction stage of chapter 6.

Section 5.2 introduces the co-simulation framework and the four components that compose it. Section 5.3 develops the vehicle model and its associated driver policy. Section 5.4 develops the LVDC traction feeder simulator and the Newton-Raphson power flow solver, with a brief comparative validation against a Gauss-Seidel baseline. Section 5.5 describes the AC distribution grid model. Section 5.6 formalises the scenario generation procedure and the operating-condition distribution Ω that the model produces. Section 5.7 presents representative simulation outputs.

5.2. Co-simulation framework

The traction-grid model is implemented as a modular co-simulation in which the principal subsystems – the vehicle fleet, the DC traction feeder, and the AC distribution grid – run as independent simulation components coordinated by a master scheduler. Each component is responsible for the solution of its own subsystem at each timestep and exchanges only the boundary quantities required by the other components. Figure 5.1 shows the modular structure.

This process is iteratively represented in figure Figure 5.5.

The four components correspond to the four subsystems of the traction power system:

Vehicle simulator Each vehicle is represented as an independent agent whose state variables (position along its assigned track section, instantaneous velocity, instantaneous power demand) are integrated locally. The simulator receives mission assignments from the master scheduler – start signal, current and next station – and emits power demand and position to the LVDC feeder. Variable per-vehicle parameters (auxiliary load, passenger mass, number of coupled cars) are applied at instantiation and remain fixed for the duration of a single scenario. The component is detailed in section 5.3.

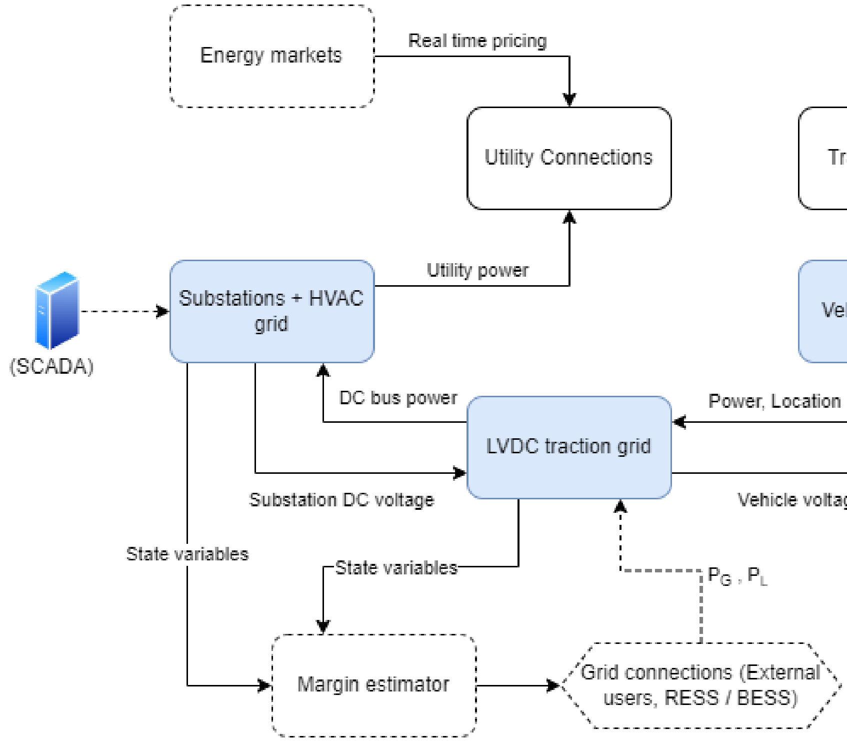

Figure 5.1: Modular structure of the traction-grid co-simulation. Four simulation components (vehicle fleet, LVDC traction feeder, AC distribution grid, master scheduler) run independently and exchange boundary quantities at each one-second synchronisation step. The vehicle simulator reports per-vehicle power demand and position to the LVDC feeder; the LVDC feeder returns the per-vehicle pantograph voltage and reports per-substation power demand to the AC distribution grid; the AC grid returns the per-substation busbar voltage used by the LVDC solver as its fixed-voltage boundary condition. The master scheduler coordinates the components and dispatches mission start signals to the vehicle fleet.

LVDC feeder simulator. The DC feeder simulator assembles the instantaneous network topology, solves the DC power flow under the prevailing vehicle positions and power demands, and emits the resulting per-substation power demands to the AC distribution model. The active-set logic for chopper activation on regenerative nodes and unidirectional rectifier handling at substations is encapsulated in this component. The component is detailed in section 5.4.

AC distribution grid simulator. The upstream AC network is modelled as a fixed-topology balanced load flow with the substation rectifier buses appearing as time-varying loads. The utility connections are treated as slack buses with fixed voltage. The component returns the per-substation busbar voltage to the LVDC simulator and reports the utility feeder loading for use in capacity assessment. The component is detailed in section 5.5.

Master scheduler. The master orchestrates the four components, enforces the one-second synchronisation step, manages the dispatch of vehicle missions per the operating timetable, and collects per-step output for downstream processing. Vehicle headways and dwell times are handled within the master rather than within the vehicle simulator, so that traffic-policy variations across scenarios modify the master's dispatch table while leaving the vehicle dynamics unchanged.

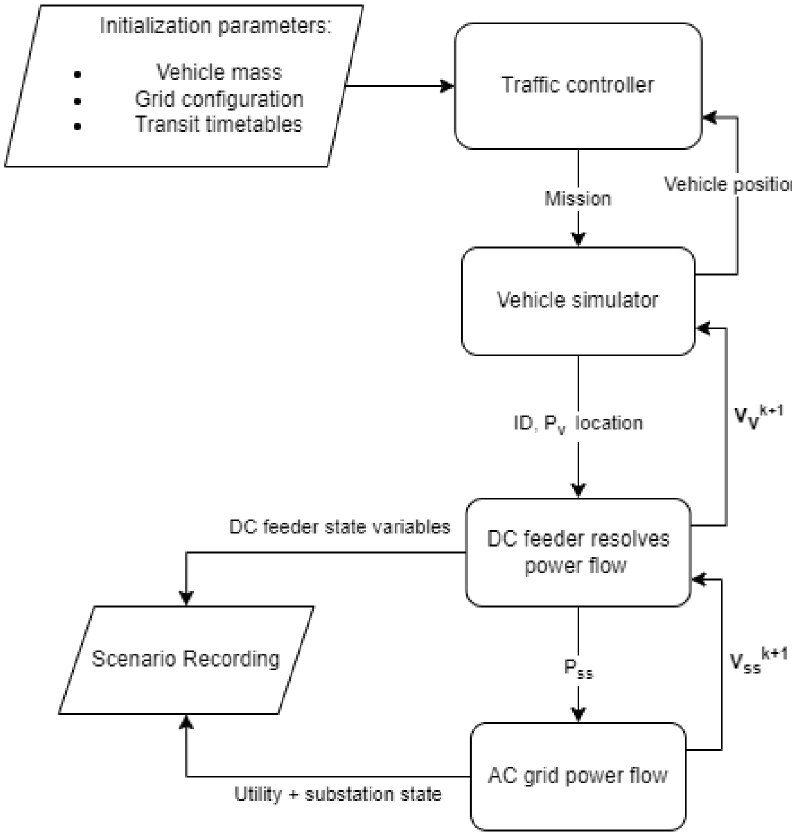

Figure 5.2: Operational procedure of the scenario generation during one full step

5.3. Vehicle model

5.3.1. Vehicle dynamics

The traction vehicle is modelled as a point mass subject to longitudinal forces, integrated at a fixed timestep of $\Delta t = 1\,\mathrm{s}$. The net tractive force at each timestep is determined by the difference between the applied force F and the running resistance F_{res}:

$$F_{\mathrm{net}} = F - F_{\mathrm{res}}. \tag{5.1}$$

The running resistance is computed using the Davis equation [34], which accounts for constant mechanical friction, speed-proportional bearing and flange losses, and speed-squared aerodynamic drag:

$$F_{\mathrm{res}} = A + Bv + Cv^2, \tag{5.2}$$

with coefficients $A = 1000\,\mathrm{N}$, $B = 30\,\mathrm{N\,s\,m^{-1}}$, and $C = 2\,\mathrm{N\,s^2\,m^{-2}}$ calibrated for the RET rolling stock. On graded track an additional resistance component is superimposed:

$$F_{\mathrm{grade}} = mg\sin\theta, \tag{5.3}$$

where m is the vehicle mass, $g = 9.81\,\mathrm{m\,s^{-2}}$, and θ is the track gradient, corrected for the direction of travel. The vehicle velocity and position are updated using a trapezoidal integration scheme:

$$v_{k+1} = v_k + \frac{F_{\mathrm{net}}}{m}\,\Delta t, \qquad s_{k+1} = s_k + \tfrac{1}{2}(v_k + v_{k+1})\,\Delta t. \tag{5.4}$$

The traction power demand reported to the DC feeder simulator is computed from the wheel-rail force and the instantaneous velocity, with a constant auxiliary load $P_{\mathrm{aux}} = 50\,\mathrm{kW}$ added to represent onboard hotel and HVAC systems:

$$P_{\mathrm{demand}} = \begin{cases} F\,v_{k+1} + P_{\mathrm{aux}} & \text{(motoring)}, \\ \eta_{\mathrm{regen}}\,F\,v_{k+1} + P_{\mathrm{aux}} & \text{(braking)}, \end{cases} \tag{5.5}$$

where $\eta_{\mathrm{regen}} = 0.8$ is a fixed regenerative braking efficiency representing electromechanical conversion losses in the drive system, consistent with values reported by [65]. Below a speed threshold of $5\,\mathrm{m\,s^{-1}}$, regenerative braking is suppressed and mechanical brakes are assumed to absorb the remaining kinetic energy. The signed-demand convention adopted throughout the model is $P_{\mathrm{demand}} > 0$ when the vehicle draws power from the DC network (motoring) and $P_{\mathrm{demand}} < 0$ when it injects power into the network (regenerating). The other parameters relating to the modelled rolling stock are available in A.1.

5.3.2. Train modelling

The driver policy governs the applied force F at each timestep and therefore determines the shape of the power demand profile delivered to the DC feeder. Deterministic methods built on the Davis resistance framework are the dominant approach in the traction literature [34]. The aim of the policy in this thesis is to produce a statistically representative distribution of substation power demand rather than to replicate the behaviour of an individual driver with high temporal precision, so the policy is required to be physically consistent and energy-realistic but not to be calibrated against operational driver data.

A pulse-and-glide (P&G) policy is adopted. The applied force is piecewise: full traction until the section speed limit is reached, coasting until a lower hysteresis bound, and a braking phase governed by a kinematic terminal constraint at the next station [66]. The empirical fit of the P&G policy against measured vehicle power profiles is shown alongside two alternative formulations in Figure 5.3: an inverse-reinforcement-learning policy following the framework of [67], and a regression-based machine-learning policy following [68]. P&G is selected on three grounds: it fits the measured profile most closely without requiring operational training data, it is the industry-standard driving instruction for energy-efficient operation, and the energy gap between an optimal P&G policy and a competent human driver is on the order of $2.7\,\%$ [69], which is within the uncertainty of the network parameters and passenger loading and therefore does not constitute a binding source of error in the scenario set produced by the model.

5.3.3. Co-simulation interface

The vehicle simulator is integrated into the Mosaik co-simulation framework and exchanges the following signals with the traffic controller and the LVDC feeder at each synchronisation step. The signal interface is illustrated in Figure 5.4.

- **Start signal** (input): issued by the traffic controller to indicate that the vehicle may depart.

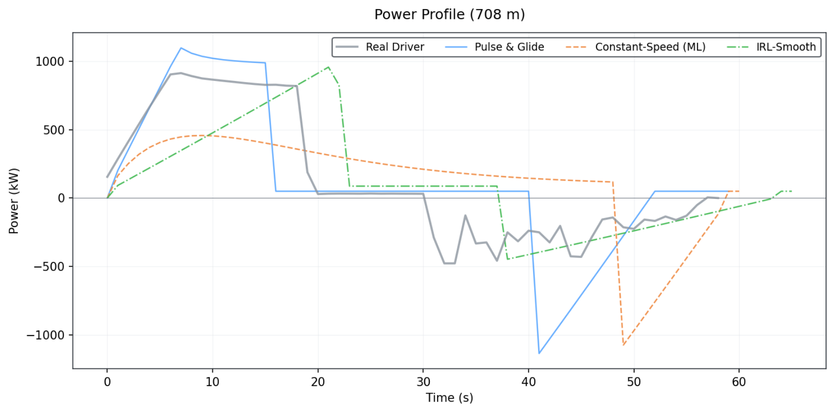

Figure 5.3: Vehicle power profiles produced by three candidate driver policies, evaluated against measured RET rolling-stock data on the same track section. The pulse-and-glide (P&G) policy reproduces the observed traction-coast-brake structure of the measurement profile most closely. The inverse-reinforcement-learning (IRL) and machine-learning (ML) policies, by contrast, exhibit smoother force profiles that under-represent the characteristic coasting phases. The P&G policy is adopted as the driver model for all subsequent scenarios.

Dwell times are handled implicitly within the traffic controller.

- **Current and next station** (input): the topological reference used by the LVDC feeder to locate the vehicle electrically.
- **Arrived** (output): a pulse signal sent to the traffic controller upon reaching the target station, enabling headway and safety-constraint enforcement.
- **Power demand** (output): the signed traction power P_{demand} forwarded to the LVDC solver.
- **Distance** (output): the accumulated distance travelled along the current section, used by the LVDC solver to compute the vehicle's electrical distance to the surrounding substations.
- **Speed** (state): the instantaneous velocity, retained internally for the physics integration and emitted for monitoring.

5.4. LVDC traction feeder simulator

The LVDC feeder simulator solves the DC power flow over the instantaneous network at each timestep. The simulator's operational pipeline is described in subsection 5.4.1 and visualised in Figure 5.5. The power flow problem itself is defined in subsection 5.4.2. The principal solver, a Newton-Raphson formulation with active-set switching, is developed in subsection 5.4.3; a Gauss-Seidel baseline used for validation is described in subsection 5.4.4. The comparative behaviour of the two solvers on a representative test network is reported in subsection 5.4.5.

5.4.1. Operational pipeline

The LVDC simulator executes six steps per simulation timestep, visualised in Figure 5.5.

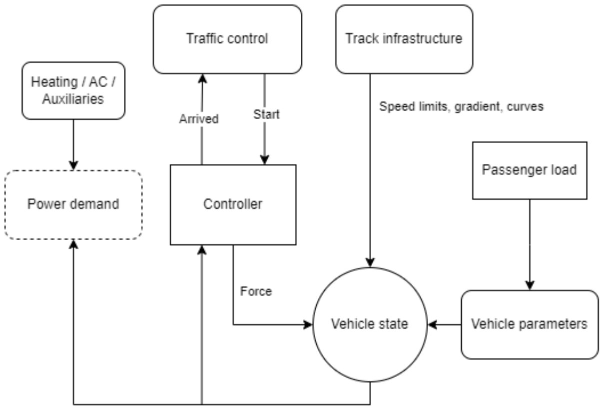

Figure 5.4: Signal interface of the vehicle simulator within the co-simulation. Inputs (left) are received from the traffic controller and consist of the dispatch signal and the station-pair reference. Outputs (right) are the per-step power demand and travelled distance sent to the LVDC feeder, and the arrival pulse returned to the traffic controller for headway enforcement. Internal state (bottom) is retained between steps for the physics integration.

Stage 1 (initialisation). The fixed railway infrastructure is loaded once at simulation start: stations, substations, track sections, feeders, and busbar-to-feeder cables. A static base conductance matrix G_{base} is assembled by stamping all permanent electrical elements. This matrix represents the empty traction network before any vehicles are inserted.

Stage 2 (vehicle location). At each timestep, every active vehicle is electrically located in the network using its current station, next station, and accumulated distance. A breadth-first search on the section graph determines the nearest upstream and downstream substations, accounting for branch topology and substation offsets. The output of this stage is, per vehicle, the identifier of the active feeder or feeder chain on which the vehicle resides.

Stage 3 (branch evaluation). The resolved vehicle position is converted into an electrical branch description. Using the feeder type and return-path resistivity, the simulator computes the distance-dependent conductances linking the vehicle to the surrounding feeder or substation nodes.

Stage 4 (node insertion). Starting from the static base matrix, one dynamic node is inserted per active vehicle. If a vehicle lies on a feeder, the original feeder branch is removed and replaced by conductance segments between the feeder endpoints and the inserted vehicle node. If a vehicle lies on a terminus tail or another non-modelled side segment, a direct conductance is stamped between the vehicle node and the relevant substation busbar. The result is a dynamic conductance matrix G_{eff} that reflects the instantaneous network occupancy.

Stage 5 (problem definition). The node list is extended with the inserted vehicle nodes and the power-demand vector P_{ref} is assembled from the current traction and regenerative setpoints, with zero net power assigned to the remaining non-substation nodes. Active substations are treated as fixed-voltage supply nodes using the AC-side feedback voltage from the previous step (or a default DC voltage at initialisation).

Stage 6 (solver execution and output mapping). The solver returns node voltages, substation currents, chopper dissipation, and DC losses for the current timestep. These values are mapped back onto simulator entities and emitted to the master orchestrator.

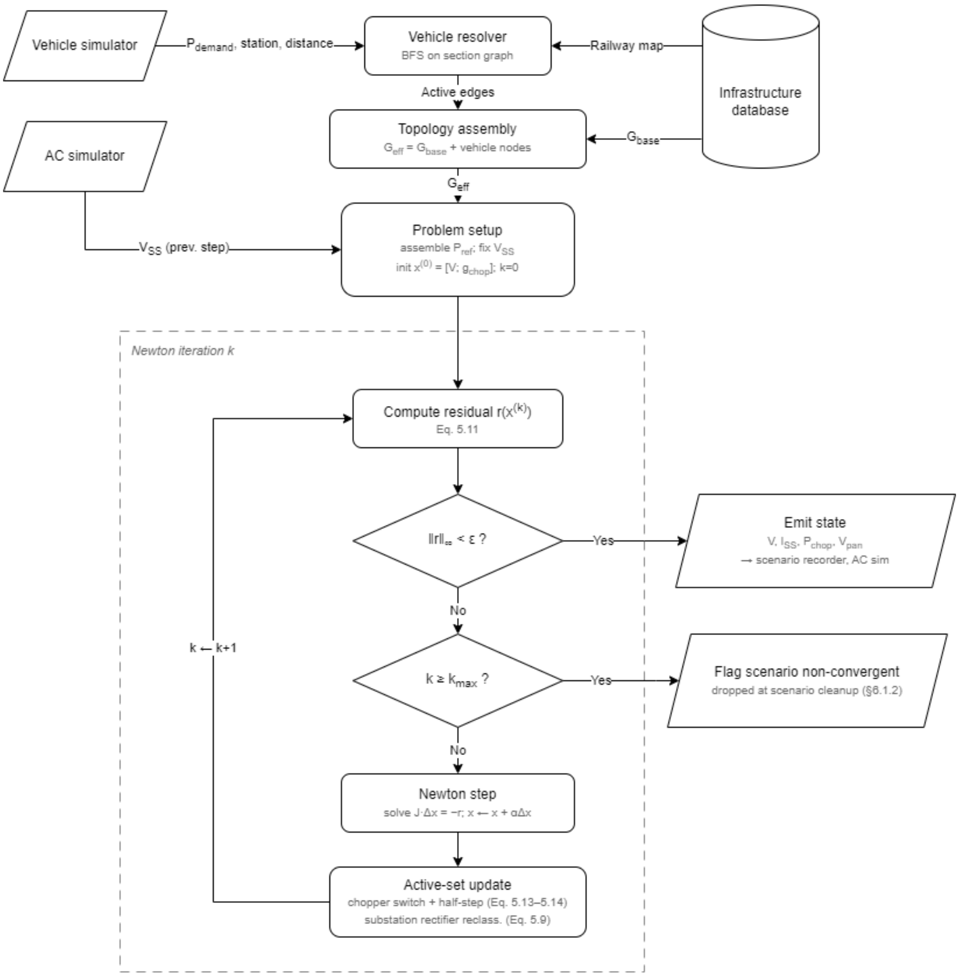

Figure 5.5: Operational pipeline of the LVDC feeder simulator. Infrastructure loading and base conductance matrix construction are performed once at initialisation. The remaining five stages execute at each one-second timestep: per-vehicle electrical location via a breadth-first resolver (within the DC feeder block), evaluation of the vehicle-to-feeder branch conductances, dynamic insertion of the vehicle nodes into the base conductance matrix, assembly of the power-flow residual and configuration of the solver, and execution of the solver followed by the back-mapping of node voltages and currents onto simulator entities .

5.4.2. Power flow problem definition

The DC power flow problem at each timestep is the determination of the node voltages V and the chopper conductances $\mathrm{g}_{\mathrm{chop}}$ that satisfy the network's Kirchhoff constraints under the instantaneous power demand vector. The network nodes are partitioned into four classes: substations $\mathcal{S}$ (fixed-voltage sources subject to the unidirectional rectifier constraint), motoring vehicles $\mathcal{M}$ (constant-power demand, $P_{\mathrm{ref},i} > 0$), regenerative vehicles $\mathcal{G}$ (constant-power source, $P_{\mathrm{ref},i} < 0$, with an active chopper that clamps the local voltage to $V_{\max}$ if exceeded), and passive nodes $\mathcal{P}$ (no injection, $P_{\mathrm{ref},i} = 0$).

The effective conductance matrix combines the base matrix with the active chopper conductances on the regenerative nodes:

$$\mathrm{G}_{\mathrm{eff}} = \mathrm{G}_{\mathrm{base}} + \mathrm{diag}(\mathrm{g}_{\mathrm{chop}}), \qquad I_i = \sum_j G_{\mathrm{eff},ij} V_j. \tag{5.6}$$

The residual at each non-substation node enforces power balance:

$$r_i = V_i I_i + P_{\mathrm{ref},i}, \qquad i \in \mathcal{M} \cup \mathcal{G} \cup \mathcal{P}. \tag{5.7}$$

Convergence is declared at iteration k when the maximum non-substation residual falls below a tolerance ε:

$$\max_{i \in \mathcal{M} \cup \mathcal{G} \cup \mathcal{P}} \left| V_i \sum_j G_{\mathrm{eff},ij} V_j + P_{\mathrm{ref},i} \right| < \varepsilon. \tag{5.8}$$

After convergence, each active substation must additionally satisfy the unidirectional rectifier condition

$$I_{\mathrm{SS},i} = \sum_j G_{\mathrm{eff},ij} V_j \geq I_{\mathrm{thr}}, \qquad i \in \mathcal{S}, \tag{5.9}$$

which excludes operating points in which a substation would draw reverse current. Substations violating Equation 5.9 are reclassified as passive nodes and the solve is restarted from the previous solution.

The two complementary nonlinearities – the chopper activation on the regenerative nodes and the rectifier blocking at substations – together define an active-set problem: the mapping from $\mathrm{P}_{\mathrm{ref}}$ to $(\mathrm{V}, \mathrm{g}_{\mathrm{chop}})$ is piecewise smooth in $\mathrm{P}_{\mathrm{ref}}$, with the active set determined as part of the solve. The two solvers developed in the next subsections differ principally in how this active-set logic is embedded into the iteration.

5.4.3. Newton-Raphson solver with active-set switching

The Newton-Raphson solver collects all free node voltages and all active chopper conductances into a single state vector and updates them simultaneously by solving the fully coupled linearised system at each iteration. The state vector is

$$\mathrm{x} = \begin{bmatrix} \mathrm{V}_{\mathcal{U}_V} \\ \mathrm{g}_{\mathcal{U}_g} \end{bmatrix}, \tag{5.10}$$

where $\mathcal{U}_V$ is the set of voltage unknowns (free voltages: motoring and passive nodes, regenerative nodes with the chopper disabled, substations classified as blocked) and $\mathcal{U}_g$ is the set of chopper-conductance unknowns (regenerative nodes with the chopper active).

The residual components are partitioned by node class:

$$r_i = \begin{cases} I_i, & i \in \mathcal{P} \text{ or blocked substation,} \\ V_i I_i + P_{\mathrm{ref},i}, & i \in \mathcal{M} \text{ or off-chopper regenerative,} \\ V_{\max} I_i + P_{\mathrm{ref},i}, & i \in \mathcal{U}_g. \end{cases} \tag{5.11}$$

The first row enforces zero current at non-injecting nodes; the second enforces power balance at free-voltage nodes; the third enforces power balance at chopper-active nodes with the voltage clamped at $V_{\max}$.

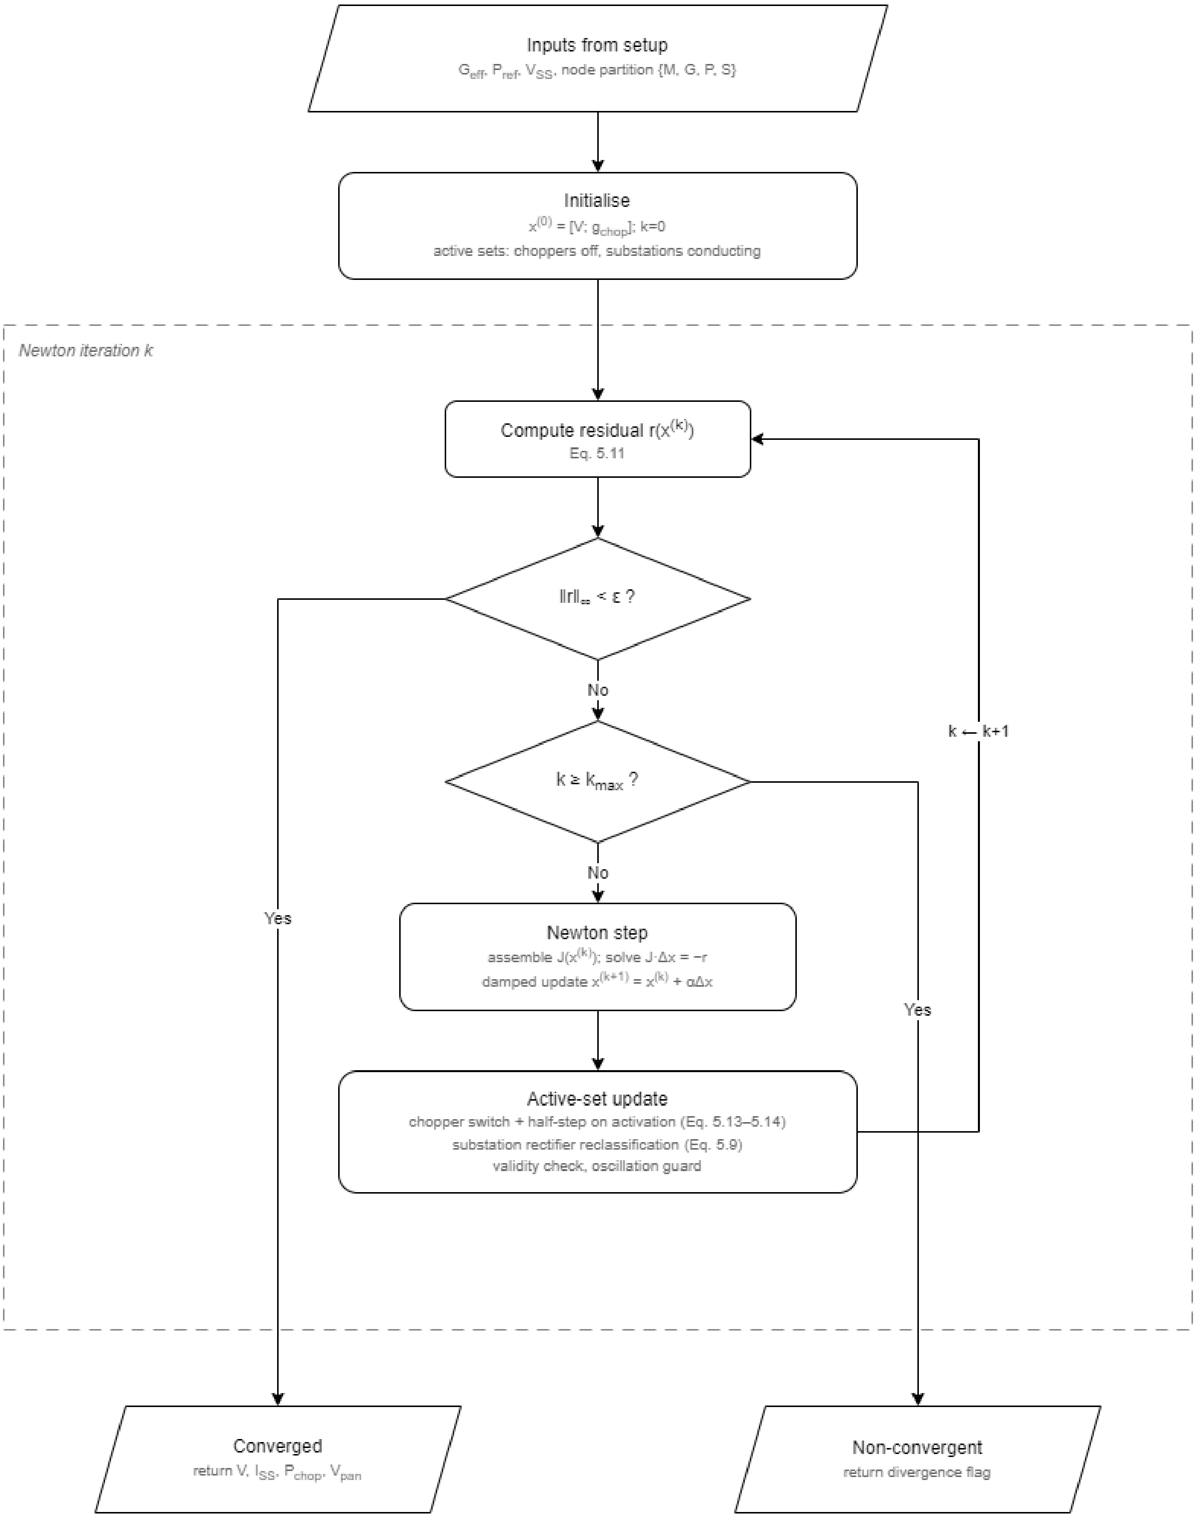

Figure 5.6: Newton Raphson implementation of LVDC traction solver

The Newton update at iteration k is the solution of the linearised system

$$\mathrm{J}(\mathrm{x}^{(k)})\,\Delta\mathrm{x}^{(k)} = -\mathrm{r}(\mathrm{x}^{(k)}), \tag{5.12}$$

with the components of the Jacobian assembled from the partial derivatives of r_i with respect to the unknowns. The fully-coupled solve is the principal distinction from the sequential Gauss-Seidel update of subsection 5.4.4: where Gauss-Seidel applies local scalar corrections at each node in turn, the Newton solver corrects all unknowns simultaneously within one matrix solve, which converges quadratically near the fixed point.

Active-set switching. At the end of each iteration, the active sets are updated. For each regenerative node $i \in \mathcal{G}$, the mode is switched from *off* to *on* if the iterate has produced $V_i > V_{\max}$, and from *on* to *off* if the iterate has produced $g_{\mathrm{chop},i} < 0$. On switching to *on*, the voltage is clamped, $V_i \leftarrow V_{\max}$. On switching to *off*, the chopper is disabled, $g_{\mathrm{chop},i} \leftarrow 0$. For each substation $i \in \mathcal{S}$, the mode is switched from *conducting* to *blocked* if the physical current $I_i^{\mathrm{phys}} = \sum_j G_{\mathrm{base},ij}\,V_j$ falls below the rectifier threshold, and from *blocked* to *conducting* when the voltage rises to match the nominal supply voltage.

Chopper activation half-step. Switching a regenerative node from *off* to *on* introduces a discontinuity into the residual that the next full Newton step would otherwise have to absorb. An optional half-iteration update is applied immediately after the mode switch: with V_k clamped to $V_{\max}$, the intermediate residual is evaluated as

$$r_k^{(k+1/2)} = V_{\max}\, I_k^{(k+1/2)} + P_{\mathrm{ref},k}, \tag{5.13}$$

and one scalar Newton correction is applied in the chopper conductance:

$$\Delta g_k^{(k+1/2)} = -\frac{r_k^{(k+1/2)}}{V_{\max}^2}. \tag{5.14}$$

After clipping to the admissible range and enforcing $g_{\mathrm{chop},k} \geq 0$, the half-step provides an improved starting point for the subsequent fully coupled Newton iteration. The half-step does not replace the full Newton solve.

Stabilisation. To improve robustness, the solver applies componentwise damping of the Newton step, an oscillation guard for nodes that switch repeatedly between active and inactive in successive iterations, and a post-solve validity check that rejects non-physical high-voltage roots above a hard ceiling $\alpha_{\mathrm{hard}} V_{\max}$.

5.4.4. Gauss-Seidel baseline

A Gauss-Seidel solver is retained as a validation baseline. The solver uses an outer-inner iteration: the outer loop maintains the active set of conducting substations, while the inner loop performs sequential per-node updates of the voltages and chopper conductances. Passive nodes are updated from Kirchhoff's current law:

$$V_i = -\frac{\sum_{j\neq i} G_{\mathrm{eff},ij}\, V_j}{G_{\mathrm{eff},ii}}. \tag{5.15}$$

Motoring nodes are updated by a scalar Newton correction

$$V_i \leftarrow V_i - \frac{r_i}{dP_i/dV_i}, \qquad \frac{dP_i}{dV_i} = I_i + G_{\mathrm{eff},ii}\, V_i. \tag{5.16}$$

Regenerative nodes are updated similarly with the chopper disabled. If the resulting voltage exceeds $V_{\max}$, the voltage is clamped and the chopper conductance is recovered explicitly:

$$g_{\mathrm{chop},i} = \mathrm{clip}\left(\frac{-P_{\mathrm{ref},i} - V_i I_i^{\mathrm{net}}}{V_i^2},\ 0,\ g_{\mathrm{chop,max}}\right). \tag{5.17}$$

Inner convergence is declared according to Equation 5.8. After the inner solve, the unidirectional rectifier condition Equation 5.9 is checked for each active substation, and any substation violating the condition is removed from the conducting set in the next outer iteration.

5.4.5. Comparative validation

The two solvers are evaluated on a representative 10-bus benchmark network configured to exercise the principal nonlinearities of the DC traction problem. The benchmark contains three substations (S_0, S_5, S_9 fixed at $V_{\mathrm{nom}} = 750\,\mathrm{V}$), four motoring vehicles (M_1, M_3, M_6, M_8 drawing 1500, 2000, 1000 and $1800\,\mathrm{kW}$), two regenerative vehicles (G_4 and G_7 injecting 2500 and $4000\,\mathrm{kW}$ with chopper threshold $V_{\max} = 900\,\mathrm{V}$), and one passive node (P_2). The two solvers are tested across a sweep of initial voltages between $300\,\mathrm{V}$ and $2500\,\mathrm{V}$.

Both solvers converge to the same fixed point for every tested initialisation, confirming that the physical operating point is unique on this network over the swept range. The Newton-Raphson trajectories (Figure 5.7b) reach the fixed point in 7 to 19 iterations regardless of starting voltage, while the Gauss-Seidel trajectories (Figure 5.7c) require 153 to 177 iterations on the same initialisations. The residual-decay comparison in Figure 5.7a shows the expected quadratic convergence of Newton-Raphson against the approximately linear convergence of Gauss-Seidel.

These results are consistent with the general behaviour expected of Newton-type and successive-iterative solvers on nonlinear power-flow problems and do not, in themselves, constitute a new result. Their purpose here is to verify that the Newton-Raphson implementation is correct (it recovers the same fixed point as the independent Gauss-Seidel implementation) and that the convergence performance is sufficient for the scenario-throughput requirement: at 7 to 19 iterations per solve with a typical per-iteration cost of a few milliseconds, the wall-clock cost of a single power-flow evaluation remains in the tens of milliseconds and the full operating-window of 5 hours (18000 timesteps) simulation completes in approximately 7 minutes, compared with the twelve-hour runtime of the legacy energy-flow simulation toolchain currently in operational use at RET.

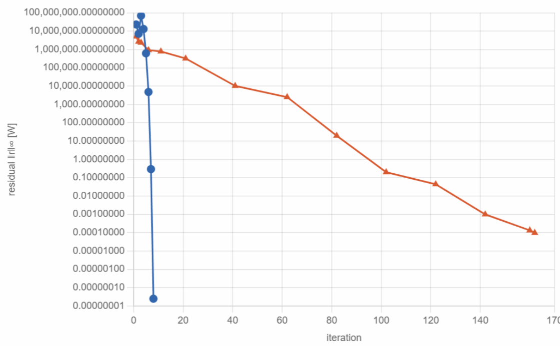

(a) Residual decay per iteration for the Newton-Raphson and Gauss-Seidel solvers from a representative initialisation. The NR residual drops by several orders of magnitude within the first few iterations, whereas the GS residual decays approximately linearly on the same network.

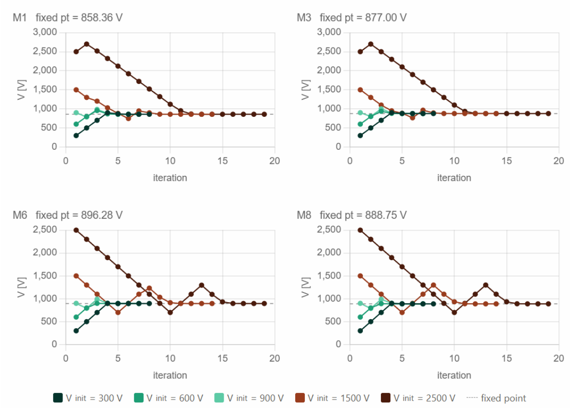

(b) Newton-Raphson voltage trajectories for the swept initialisation range. All trajectories converge to the same fixed point in 7 to 19 iterations regardless of the starting voltage, confirming that the operating point is unique over the swept range.

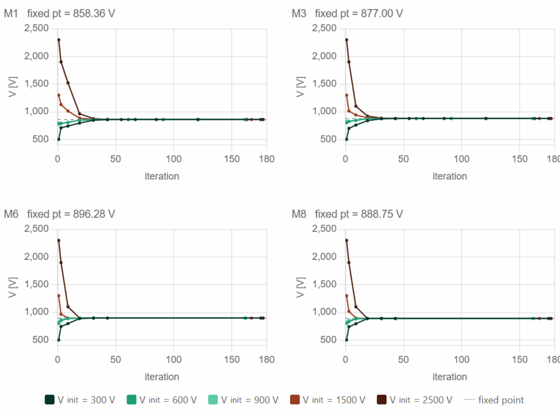

(c) Gauss-Seidel voltage trajectories on the same sweep. The fixed point is recovered for every initialisation but in 153 to 177 iterations.

Figure 5.7: Comparative behaviour of the Newton-Raphson and Gauss-Seidel solvers on the 10-bus benchmark. Both solvers recover the same physical fixed point across the swept initialisation range, confirming uniqueness of the operating point for this network. The convergence rate of the Newton-Raphson solver is approximately an order of magnitude faster (Figure 5.7a, Figure 5.7b, Figure 5.7c).

5.5. Distribution grid model

The RET distribution grid serving the Calandlijn is composed of four feeding clusters, each fed by a single utility connection during normal operation, as described in chapter 2. The clusters are electrically independent on the AC side but remain coupled to one another through the LVDC traction feeder, since vehicles travelling between cluster boundaries can draw power from substations of different clusters within a single journey. The distribution-side model captures only the AC behaviour upstream of the rectifier substations; the DC coupling between clusters is resolved by the LVDC simulator of section 5.4.

The distribution grid is implemented using the open-source pandapower [70] package. The AC network is represented as a fixed-topology, balanced load-flow model in which the railway side appears as time-varying constant-power loads at the rectifier substation buses. The model includes medium-voltage buses at $10.96\,\mathrm{kV}$, with a separate $13.5\,\mathrm{kV}$ supply at the TRP cluster, and twelve rectifier buses at $0.662\,\mathrm{kV}$, corresponding to the transformer secondary before DC rectification. The network is energised at four utility slack points (MSMCP, GSGDW, GSNSL, GSTRP). Transformer ratings are set to either $1.44\,\mathrm{MV\,A}$ or $2.50\,\mathrm{MV\,A}$ per the per-substation data of chapter 2, with fixed tap positions chosen according to the selected operating configuration. Substation loading is imposed as constant-power demand with a fixed power factor of 0.95, so that both the active and reactive loading of the upstream utility points are represented. The pandapower model returns the per-substation busbar voltage to the LVDC simulator and reports the per-utility feeder loading for use in the capacity assessment of chapter 7.

5.6. Scenario generation procedure

The traction-grid model developed in the preceding sections is the engine of the scenario-generation stage of the analysis pipeline. A *scenario*, as defined in section 4.2, is a per-second snapshot of the network state under a specific instantaneous configuration of vehicle positions, vehicle power demands, and substation operating points, together with the resulting solution of the DC power flow. A scenario is generated by running the co-simulation of section 5.2 over a representative operating window of length T_{win} under one realisation of the input distribution; the resulting time-indexed set of per-step records constitutes one element of the operating-condition set Ω.

This section first inventories the quantities recorded per timestep during scenario generation (subsection 5.6.1), then describes the input distribution against which scenarios are sampled, and concludes with the generation procedure itself (subsection 5.6.2).

5.6.1. Recorded data products per scenario

The co-simulation records fourteen distinct quantities at each one-second synchronisation step, organised into three groups by their downstream use in the analysis pipeline. Table 5.1 summarises the full inventory.

Capacity-allocation quantities. The first group comprises the variables on which the multi-objective chance-constrained optimisation of chapter 7 operates directly. The per-substation DC power demand $P_{\mathrm{SS},i}^{\mathrm{DC}}$ is the principal quantity that the CPL allocation modifies: each candidate allocation adds external load to a subset of the twelve substations and the optimiser is responsible for finding the allocation that maximises aggregate CPL while keeping the augmented loading within the operating envelope. The per-substation DC current $I_{\mathrm{SS},i}$ and DC busbar voltage $V_{\mathrm{SS},i}^{\mathrm{DC}}$ are the variables on which the hard feeder-current (g_3) and DC node-voltage (g_1) constraints operate. The substation AC busbar voltage $V_{\mathrm{SS},i}^{\mathrm{AC}}$ captures the upstream support condition; large excursions below nominal indicate that the AC distribution is strained independently of the DC-side loading and that the substation's useful capacity is constrained by voltage support rather than by rated power. The utility connection power $P_{\mathrm{util},k}$ records the contractual loading at each of the four utility points, against which the upstream-side capacity constraint (g_6) is evaluated.

Margin-assessment quantities. The second group comprises the quantities that characterise the energetic state of the network and that distinguish the present problem from a generic

Table 5.1: Quantities recorded per timestep during scenario generation. All quantities are sampled at the synchronisation rate of $1\,\mathrm{Hz}$. The cardinality column states the number of instances per timestep: twelve for per-substation quantities, four for the utility connections of the four feeding clusters, one for the network-aggregate quantities, and per-vehicle for vehicle-indexed quantities (cardinality varies with the traffic configuration of each scenario).

Symbol	Quantity	Units	Cardinality
Substation and utility quantities			
$P^{\mathrm{DC}}_{\mathrm{SS},i}$	Substation DC active power demand	kW	12
$I_{\mathrm{SS},i}$	Substation DC busbar current	A	12
$V^{\mathrm{DC}}_{\mathrm{SS},i}$	Substation DC busbar voltage	V	12
$V^{\mathrm{AC}}_{\mathrm{SS},i}$	Substation AC busbar voltage	V	12
$P_{\mathrm{util},k}$	Utility connection active power	MW	4
Network-aggregate quantities			
$P^{\mathrm{tot}}_{\mathrm{chop}}$	Aggregate chopper dissipation	kW	1
$P^{\mathrm{DC}}_{\mathrm{loss}}$	Aggregate DC transmission losses	kW	1
Per-vehicle quantities			
$V_{\mathrm{pan},j}$	Pantograph voltage	V	per vehicle
v_j	Instantaneous velocity	$\mathrm{m\,s^{-1}}$	per vehicle
$P_{\mathrm{demand},j}$	Signed traction power demand	kW	per vehicle
s_j	Accumulated distance	m	per vehicle
e_j	Active edge identifier	–	per vehicle
σ_j	Current station reference	–	per vehicle
$P_{\mathrm{chop},i}$	Vehicle chopper dissipation	kW	per vehicle

DC OPF. The per-vehicle chopper dissipation $P_{\mathrm{chop},j}$ is the power burnt by the chopper resistor at vehicle j when local regenerative voltage rise forces chopper activation; this quantity is the direct measure of wasted regenerative braking energy at the vehicle level and is a candidate secondary objective for the optimisation. The network-aggregate chopper dissipation $P^{\mathrm{tot}}_{\mathrm{chop}}$ and DC transmission losses $P^{\mathrm{DC}}_{\mathrm{loss}}$ together describe the energy-efficiency profile of the network as a whole: a CPL allocation that increases throughput while reducing aggregate chopper dissipation (by providing a productive sink for recuperated energy near regenerative events) is materially preferable to one that increases throughput without affecting chopper dissipation, even if both have the same aggregate-CPL objective value.

Operational-state quantities. The third group comprises the per-vehicle records needed for dynamic-feasibility validation (section 4.5) and for tracing the instantaneous load distribution back to its source. The pantograph voltage $V_{\mathrm{pan},j}$ is the per-vehicle power quality observable: voltage excursions at the pantograph are the principal operational consequence experienced by the rolling stock and are constrained by EN 50163. The vehicle velocity v_j, accumulated distance s_j, and signed power demand $P_{\mathrm{demand},j}$ together describe the kinematic and energetic state of each vehicle and provide the trajectory data required for chronological replay. The categorical fields e_j (active edge identifier) and σ_j (current station reference) re-localise each vehicle within the network topology at every timestep, enabling the binding-constraint diagnosis of chapter 8 to attribute a constraint violation in the optimised allocation to a specific section of track and a specific traffic configuration.

5.6.2. Generation procedure

For each scenario index $\omega \in \{1, \dots, N_{\text{scen}}\}$, the input variables are sampled jointly from the input distribution. A random seed σ_ω is fixed at the start of the scenario so that the scenario is fully reproducible from (ω, σ_ω) alone. The master scheduler dispatches the vehicle missions according to the sampled traffic configuration, and the co-simulation is run forward for $T_{\text{win}} = 5\,\text{h}$ at one-second resolution. The full set of quantities catalogued in Table 5.1 is recorded at every step, yielding a per-scenario record of $18\,000$ timesteps. The random variable in this instance is the applied by the traffic management simulator, which applies a variable dwell time for departing vehicles.

The scenario set

$$\Omega = \{\omega_1, \omega_2, \dots, \omega_{N_{\text{scen}}}\} \tag{5.18}$$

is the output of the modelling chapter and the input to the scenario reduction stage of chapter 6.

The generation procedure is deliberately decoupled from the optimisation procedure: every scenario in Ω is produced once, stored, and reused as the optimisation explores different formulations and as the reduction stage explores different reduction methods. The decoupled architecture is the principal reason that an optimality-aware reduction step (chapter 6) is needed: without it, the optimiser would either have to evaluate the full Ω at prohibitive cost or rely on a reduced set produced without knowledge of the optimisation it must support.

5.7. Simulation results

This section reports outputs of the traction-grid co-simulation under a single nominal scenario, drawn from Ω for the purpose of demonstrating the model's behaviour and the form of the data products that downstream stages consume. The intent here is illustrative rather than analytical: the optimisation results and the network-wide statistical characterisation of Ω are reported in chapter 8.

Figure 5.8 presents four views of the network loading at a representative snapshot from the nominal scenario, chosen at the five-hour mark of the operating window.

The Newton-Raphson implementation also represents an improvement over the backward-forward sweep solver of Diab et al. [71] in terms of computational cost. The inclusion of the active set switching and newton update step circumvents the expensive per-node evaluation at the cost of solver robustness; 405 out of 18000 (2.3 %) evaluated scenarios remained non-convergent.

5.8. Conclusion

This chapter implemented Design Objective DO1: a co-simulation model coupling the vehicle fleet, the LVDC traction feeder, the AC distribution grid, and a master scheduler at a one-second step, and the procedure that produces the operating-condition set Ω. The vehicle model uses a Davis-resistance point mass with a pulse-and-glide driver policy, which repro-

Table 5.2: Baseline operating characteristics over the 5-hour reference scenario (traffic-only, neutral taps, no configurable load).

Quantity	Baseline (5 h)
Grid efficiency η (Total traction Energy / total utility energy)	82.7 %
Mean traction power	5516 kW
Mean regenerated power	2374 kW
Conductor loss (mean / peak)	348 / 2319 kW
Chopper-dumped regen (mean / peak)	787 / 7115 kW
Regenerated energy dumped	33.1 %
Worst-substation loading (mean / peak)	0.84 / 4.19 pu
Vehicle voltage (min / median)	433 / 845 V
Vehicle undervoltage ($V < 500$ V)	0.002 %
Utility intake (mean / peak)	4.30 / 17.04 MW
Wall clock time (AMD Ryzen 5 7640HS 6C12T)	7 minutes

duces the measured power profile most closely without operational training data. The LVDC feeder is solved by a Newton-Raphson method with active-set switching for chopper activation and rectifier blocking, validated against a Gauss-Seidel baseline that recovers the same unique operating point.

The throughput requirement of section 4.2 is met: a five-hour window of $18\,000$ timesteps completes in about seven minutes, against the twelve-hour legacy runtime at RET, which makes the downstream optimisation tractable. The cost is robustness: 405 of $18\,000$ evaluations (2.3 %) did not converge, a limitation carried to the optimisation and revisited in chapter 8. Generation is decoupled from optimisation: each scenario is produced once and reused as the reduction and optimisation stages vary, which is what makes the optimality-aware reduction of chapter 6 necessary.

In terms of RQ1, this chapter fixes how the operational state is quantified. The fourteen recorded quantities per timestep (Table 5.1) span the substation loading and voltages the constraints act on, the chopper dissipation and losses that measure regenerative-energy waste, and the per-vehicle records that localise each load. The set Ω is the chronological realisation of that state; reducing it to the snapshot representation RQ1 calls for is the task of chapter 6.

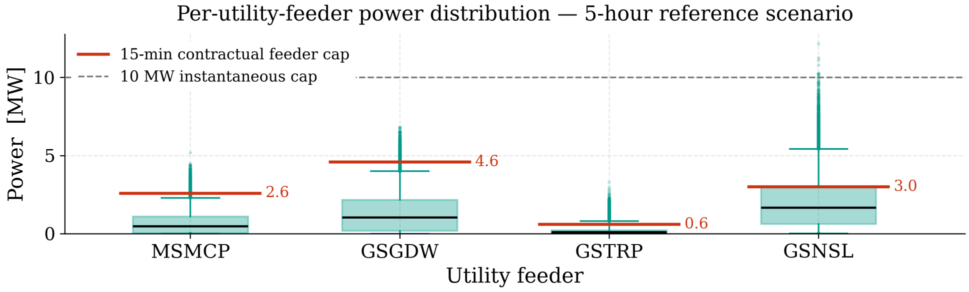

(a) Active power drawn at each of the four utility connection points (MSMCP, GSGDW, GSNSL, GSTRP) at $t = 5\,\mathrm{h}$ into the operating window. The four utility points feed the four clusters of the distribution grid identified in chapter 2.

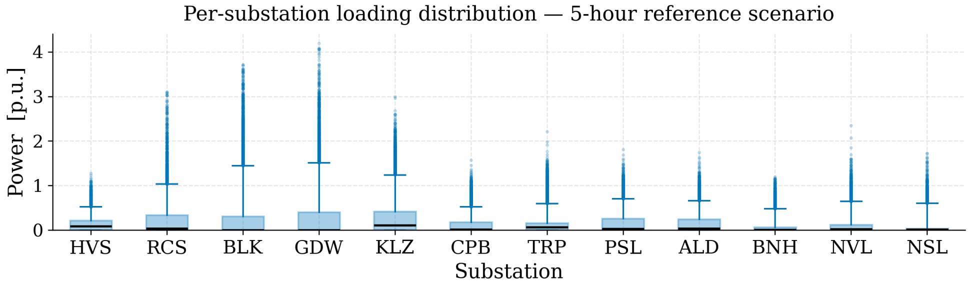

(b) Per-substation active power demand at the same snapshot, showing the distribution of load across the twelve substations.

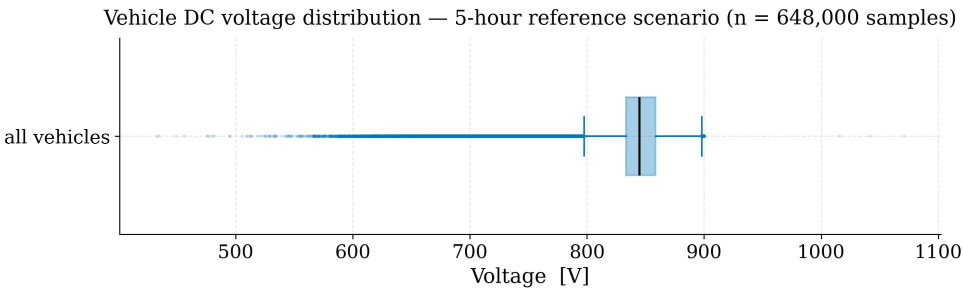

(c) Recorded vehicle voltages across all timesteps, 36 vehicles across 18000 samples

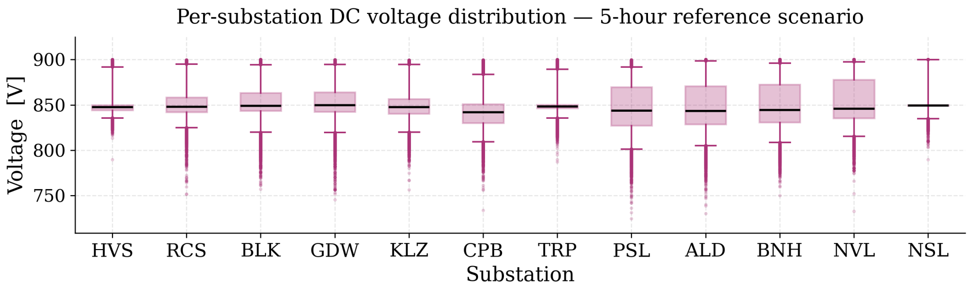

(d) Per-substation DC busbar voltage

Figure 5.8: Network loading at a representative snapshot $t = 5\,\mathrm{h}$ from the nominal scenario: (a) active power at the four utility connections, (b) per-substation active power demand,(c) vehicle voltage distribution (36 recorded vehicles) (d) per-substation DC busbar voltage. The three subfigures together describe the instantaneous distribution of load and voltage across the network at this snapshot.

6

Traction power system scenario reduction

6.1. Introduction

Traction grids are, similar to wind energy and PV penetration in DG grids, subject to largely fluctuating power with a spatiotemporal variation. The scenario representation is illustrated in figure Figure 6.1 which is a simulation of 2364 samples, each corresponding to 1 second with confidence intervals based on the total power distributed from the substations. The samples have been generated by the simulator as implemented in chapter 5.

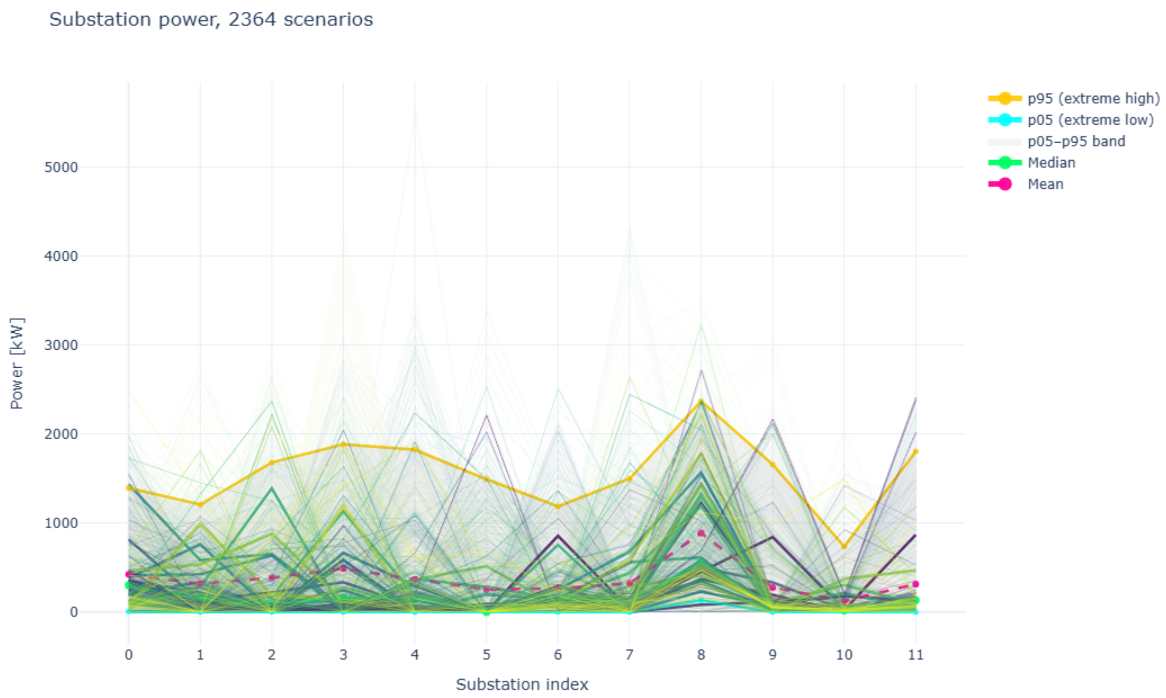

Figure 6.1: Simulator-generated scenarios of the traction grid of 2364 samples, expressed in terms of instantaneous substation power demand. Each line represents a timestamp with the substation power demand.

As mentioned in section 4.4, a multi-objective optimisation formulation is employed in chapter 7 to find the external power load margins at the substation level to preserve the spatial diversity of solutions (evenly distributed allocations across the grid instead of clustered solutions around a single operating point). As scenarios directly serve as evaluation points within the optimiser, the computation scales by the amount of scenarios entered. This chapter explores and evaluates 2 methods to reduce the original simulated set Ω to the reduced set Ω_s, of which the quality is determined by the optimisation scheme.

6.1.1. Selected methods for implementation

Two methods are retained for implementation and comparison. The first is **PCA + k-medoids**, selected as a feasibility-preserving, widely benchmarked baseline that exploits the low-rank structure of traction power flows and is computationally cheap enough to permit sweeps over K and bootstrap stability analysis. GMM clustering is not considered due to the skewed results due to the low average demand (high peaks get smoothed, which causes a loss of information. The second is an **objective-aware scenario reduction** method that clusters scenarios in the space of objective values induced by a fixed reference decision x_0, rather than in raw feature space. Formally, each scenario ξ_i is mapped to a feature vector $\phi_i = \left(f_1(x_0, \xi_i), \dots, f_M(x_0, \xi_i)\right)$, and k-medoids is applied in this objective space. The reference decision x_0 is obtained once from a deterministic-equivalent solve on the mean scenario, avoiding in-the-loop optimisation. This is a computational relaxation of the problem-driven scenario reduction of Zhuang et al. [50] and the decision-based clustering of Keutchayan et al. [60], trading the richness of per-scenario optimality information for tractability and compatibility with a decoupled SR/SO pipeline.

This pairing sets up the central comparison of the chapter: a distribution-preserving method that is blind to the margin objective, against an objective-aware method that injects knowledge of the MOOPF structure into the similarity metric without requiring in-the-loop optimisation. Full problem-driven reduction (Zhuang [50], Keutchayan [60]) is discussed but deferred to future work, since its computational structure (iterative MOOPF solves in an outer SR loop) conflicts with the decoupled SR/SO framing adopted here.

6.1.2. Feature selection - Distribution driven (PCA + k-medoids)

To establish a reference metric for the quality of the problem informed scenario reduction, a distribution oriented approach is required to establish a baseline for the scenario reduction. The quality of the scenario selection depends on the metric of representativeness; the solutions found by optimisation should preserve their validity when configured in the simulation. Hence, this representativeness should be adequately defined before feature extraction and scenario reduction. Feature selection, which entails picking the appropriate metrics derived from the simulation based on their explained variance in the dataset, must be implemented to identify the important metrics and reduce the dataset dimensionality. To this end, the following feature vector is derived from the training set:

$$\phi = \{\phi_{veh}, \phi_{DCfeeder}, \phi_{utility}\} \tag{6.1}$$

In which ϕ_{veh}represents vehicle voltage, location and power demand, chopper dissipation,$\phi_{DCfeeder}$ represents substation power flow, DC node voltages and branch currents and $\phi_{utility}$ represents the distribution grid power flow derived quantities including the supplied energy. As the goal of this reference is to represent scenarios based on the statistical distribution, there are no target variables selected as this would (implicitly) inform the selection method about the objectives. Hence, the clustering method is based on unsupervised learning. Due some metrics being closely correlated, redundancy, which results in a reduction of feature dimensionality, which mitigates clustering ambiguity. The pipeline is therefore implemented in the following manner:

1. Dataset cleanup: non-convergence scenario removal
2. Train-test-validation split: 80/20 split of the scenario set to evaluate reduction quality
3. PCA: feature correlation audit, feature set reduction based on explained variance while keeping substation demand and utility power.
4. K-medoids clustering, in which K is chosen based on the number of the kept scenarios.
5. Wasserstein distance is computed between selected scenarios and the test set, the optimal value for K is extracted for $0.1\% < K < 10\%$ of the training set with overall Wasserstein distance residuals improving less than 5 % from the original update.
6. Vehicle features (ϕ_{veh}) are extracted from selected scenarios and saved as inputs to the optimisation script.

6.1.3. Optimality-aware scenario reduction

The distribution-driven reduction in the previous subsection preserves the empirical distribution of the raw scenario features but is blind to the optimisation it supports: scenarios that are statistically typical can be far from the constraint boundary, and scenarios that bind the constraints can lie in low-density regions of feature space and so be under-represented in the medoid set. The method developed in this section addresses that gap without coupling the reduction to the optimiser. Each scenario is summarised by an *optimality fingerprint*: a feature vector that encodes both the magnitude of the scenario's feasibility margin under a fixed reference allocation and the shape of the constraint that first binds along that allocation. Clustering on this fingerprint over-represents the binding scenarios in proportion to their optimisation relevance.

Feasibility line search. For each scenario $s \in \Omega$, the network is evaluated under a one-parameter family of uniform-scaling CPL allocations

$$P_{\mathrm{CPL},k}(\delta) = \delta \cdot S_{\mathrm{rated},k}, \qquad k = 1, \dots, N_{\mathrm{SS}}, \qquad \delta \in [0, \delta_{\mathrm{hi}}], \tag{6.2}$$

parametrised by a scalar δ. Under this parametrisation every substation is loaded to the same per-unit of its rated power. The feasibility predicate evaluated during the search is the conjunction of the hard, instantaneous constraints of section 4.4,

$$\mathcal{F}^{\mathrm{hard}}(s, \delta) \equiv \left[V_i^{(s)}(\delta) \geq V_{\mathrm{min}} \right] \wedge \left[\lambda_k^{(s)}(\delta) \leq \lambda_{\mathrm{trip}} \right] \wedge \left[|I_{\mathrm{SS},k}^{(s)}(\delta)| \leq I_{\mathrm{max}} \right], \tag{6.3}$$

quantified universally over the relevant index sets. The per-scenario margin is

$$\delta_{\max}(s) = \sup\{\delta \in [0, \delta_{\mathrm{hi}}] \ : \ \mathcal{F}^{\mathrm{hard}}(s, \delta)\}, \tag{6.4}$$

approximated by bisection with $\delta_{\mathrm{hi}} = 4$ and $n_{\mathrm{iter}} = 8$, yielding a resolution of ≈ 0.016 p.u. at ≈ 10 power-flow solves per scenario.

Three design choices in the predicate (6.3) bear stating. *Isotropic search direction*: the uniform-scaling parametrisation (6.2) fixes a single direction in the 12-dimensional CPL space, so $\delta_{\max}(s)$ measures headroom in that direction alone — a scenario with low $\delta_{\max}$ may have more headroom under a non-uniform allocation that the GA is free to find. *Hard constraints only*: the duration-dependent constraints g_4, g_5 and the contractual cap g_6 are excluded because the moving averages they evaluate are undefined on a single-second snapshot; their probabilistic enforcement is handled in chapter 7. *Fixed taps*: t is held at nominal during the search, so $\delta_{\max}$ is a conservative estimate of the headroom in the joint $(\mathrm{P}_{\mathrm{CPL}}, \mathrm{t})$ space. The relative ordering of scenarios — which is what the clustering uses — is preserved as long as the tap-feasibility coupling is similar across scenarios.

Three edge cases are tagged explicitly: *baseline-unsafe* ($\delta_{\max} = 0$, baseline solve converges to an infeasible state), *diverged* (solver fails at $\delta = 0$, scenario excluded), and *unbounded* ($\mathcal{F}^{\mathrm{hard}}$ still holds at δ_{hi}, the bisection saturates). The first two together identify scenarios that the optimiser cannot improve on regardless of allocation; the third identifies scenarios with comfortable slack against every hard constraint.

Fingerprint construction. Each scenario s is summarised by a fingerprint vector

$$\phi 0(s) = \left(\delta_{\max}(s),\ V_{\min}^{\mathrm{veh}\,\star}(s),\ \{\lambda_k^\star(s)\}_{k=1}^{N_{\mathrm{SS}}},\ \{|I_{\mathrm{SS},k}^\star(s)|\}_{k=1}^{N_{\mathrm{SS}}},\ \mathrm{Var}_k(\lambda_k^\star(s)),\ \mathrm{e}_{\iota(s)}\right), \tag{6.5}$$

whose components answer three questions about the scenario: *how much* margin it has against the feasibility boundary ($\delta_{\max}$), *what* the network looks like when that margin is exhausted (the starred state variables at $\delta = \delta_{\max}$), and *which* constraint binds first (the one-hot label $\mathrm{e}_{\iota(s)}$, with $|\iota| \approx 25$ categories covering the voltage floor, per-substation loading-ratio and current limits, and three edge-case tags). Two scenarios with the same $\delta_{\max}$ but different binding substations or constraint types are placed in different regions of fingerprint space, which is the property the clustering will exploit.

The fingerprint has dimension $27 + |\iota| \approx 53$, an order of magnitude smaller than the raw feature vector of subsection 6.1.2. PCA is therefore not needed: k-medoids is applied directly under Euclidean distance after column-wise standardisation on the train split. K is selected by the same Wasserstein-elbow rule as the distribution-driven baseline.

Relation to problem-driven scenario reduction. The construction is a tractable relaxation of the problem-driven scenario reduction of Zhuang et al. [50]. PDSR clusters scenarios by their behaviour under the *optimal* multi-objective decision, which couples the reduction to the optimiser through an outer loop. The fingerprint replaces that loop with a single uniform-scaling line search per scenario (Equation 6.3), and considers only the hard instantaneous constraints; the duration-dependent constraints and multi-objective trade-offs are deferred to the optimiser. The cost is the loss of PDSR's formal optimality-gap guarantee; the benefit is compatibility with the decoupled SR/SO pipeline.

6.2. Reduced set evaluation

The selection of the scenario size (K) is described by the Wasserstein elbow curve. This is illustrated in Figure 6.2. The chosen set size for both methods is 1408 scenarios.

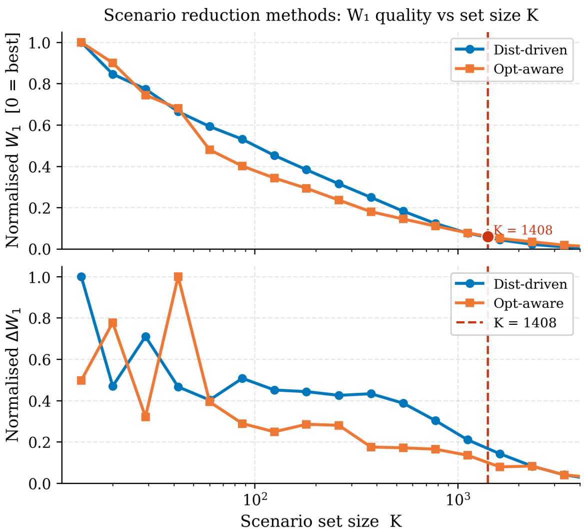

Figure 6.2: K selection based on Wasserstein metric: the top graph indicates the full Wasserstein distance minimisation over K scenarios, while the bottom graph represents the distance which is reduced per step.

The scenario reduction methods yield a distribution of the loading ratio of substations (this metric is chosen as an indicator of a tail-demand) as illustrated in Figure 6.3.

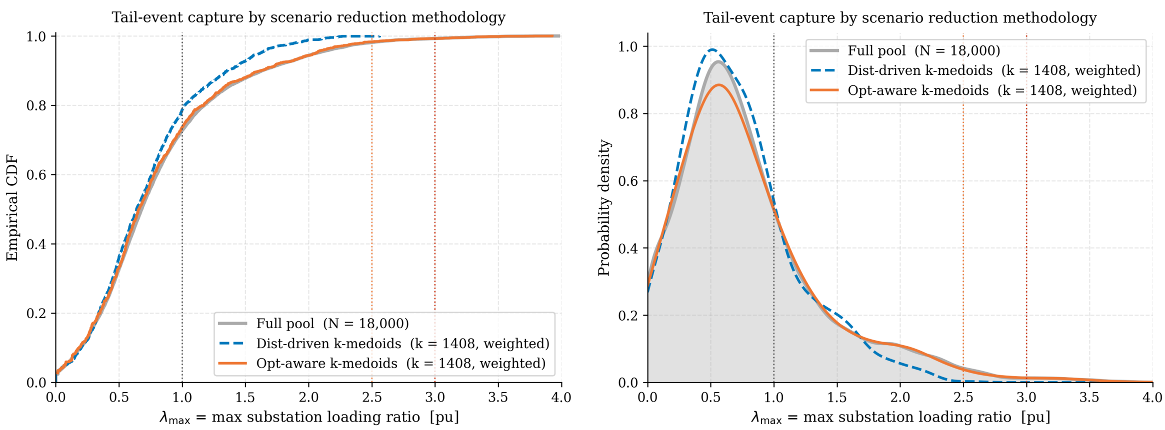

Figure 6.3: Loading ratio CDF and PDF captured by the reduced sets of both SR methods

It is evident that the distribution-driven scenario selection overrepresents the training set for $\lambda_{max} \leq 1$, but fails to capture the tail-end behaviour. For the optimality-aware method, this is the other way around; tail-end events are adequately represented, while the statistic peak of operating conditions is underrepresented. The results on the quality of the scenario reduction after optimisation are further discussed in chapter 8.

6.3. Conclusion

This chapter reduced Ω to a tractable set Ω_s by two methods sized at $K = 1408$ by a common Wasserstein-elbow rule. The distribution-driven method (PCA + k-medoids) preserves the empirical distribution but is blind to the optimisation. The optimality-aware method clusters on a per-scenario fingerprint, built from one uniform-scaling feasibility line search, that encodes each scenario's margin and which constraint binds first; this over-represents the constraint-binding tail and is a tractable relaxation of problem-driven reduction that keeps the reduction-optimisation pipeline decoupled.

The key result is that the two sets capture opposite parts of the loading-ratio distribution: distribution-driven preserves the bulk and misses the tail, optimality-aware preserves the tail and thins the bulk (Figure 6.3). Which serves the certified capacity better is settled only by the optimisation of chapter 8. In terms of RQ2, Ω_s is the representative subset that makes the capacity computation tractable (10×reduction in computational overhead), and these two strategies are the two answers compared there.

7

Traction network optimisation

7.1. The optimisation problem

The multi-objective grid-margin problem is stated formally in section 4.4: the decision vector $\mathrm{x} = [\mathrm{P}_{\mathrm{CPL}}; \mathrm{t}]$ collects the per-substation external-load allocations and rectifier tap positions; the objective vector $\mathrm{f}(\mathrm{x})$ in section 4.4 contains $N_{\mathrm{SS}} + N_{\mathrm{U}} + 2 = 18$ scalar components for the Calandlijn; the constraint set $g_1, \dots, g_8$ of Table 4.1 bounds the admissible operating envelope, with $g_1, \dots, g_5$ enforced strictly and g_6, g_7, g_8 admitting the probabilistic enforcement developed in this chapter. The scenarios over which f and the g_j are evaluated are the representative set $\mathcal{S}_{\mathrm{R}}$ produced by the reduction pipeline of chapter 6, with their associated empirical weights $\{w_s\}$.

The problem is non-convex (the constant-power-load injections are non-linear in voltage), mixed-integer (the tap variables), and evaluated through a non-differentiable simulator whose internal active-set logic has no closed-form derivative. Under these conditions the gradient-based and convex-relaxation methods surveyed in chapter 3 are either inapplicable or require approximations coarse enough to invalidate the result. The chapter commits to two meta-heuristic families, NSGA-III and SMPSO, whose properties are summarised in Table 3.2.3; the emphasis here is on the implementation choices that adapt each canonical algorithm to the TPSS evaluator, and on the chance-constrained reformulation of the duration-dependent constraints.

7.2. Objective formulations evaluated

The optimisation uses the 14-component objective vector

$$\mathrm{f}(\mathrm{x}) = \big(\{f_k^{\mathrm{CPL}}(\mathrm{x})\}_{k=1}^{12},\ f^{\mathrm{imb}}(\mathrm{x}),\ f_{\Sigma}^{\mathrm{util}}(\mathrm{x}) \big), \tag{7.1}$$

with components

$$\begin{aligned}
f_k^{\mathrm{CPL}}(\mathrm{x}) &= -P_{\mathrm{CPL},k} && \text{(per-substation CPL, to be maximised),}\\
f^{\mathrm{imb}}(\mathrm{x}) &= \mathrm{Var}_k\big(\lambda_k(\mathrm{x})\big) && \text{(variance of substation loading ratios),}\\
f_{\Sigma}^{\mathrm{util}}(\mathrm{x}) &= \sum_{u=1}^{4} \overline{P}_{\mathrm{util},u}(\mathrm{x}) && \text{(mean utility intake over the four AC feeds).}
\end{aligned}$$

The same vector is used for both optimisers (NSGA-III and SMPSO); constraints are listed in Table 4.1.

7.3. Empirical constraint handling under uncertainty

The optimisation is performed on the reduced scenario set $\mathcal{S}_{\mathrm{R}}$, whose scenarios carry empirical probability weights w_s satisfying

$$w_s \geq 0, \qquad \sum_{s\in\mathcal{S}_{\mathrm{R}}} w_s = 1. \tag{7.2}$$

For any scenario-dependent quantity $h^{(s)}(\mathrm{x})$, its weighted empirical expectation is

$$\widehat{\mathbb{E}}_{\mathrm{R}}\left[h(\mathrm{x})\right] = \sum_{s\in\mathcal{S}_{\mathrm{R}}} w_s h^{(s)}(\mathrm{x}). \tag{7.3}$$

For constraint j, let

$$g_j^{(s)}(\mathrm{x}) \leq 0 \tag{7.4}$$

denote satisfaction of the corresponding operational limit in scenario s. The weighted empirical violation probability is

$$\widehat{p}_j(\mathrm{x}) = \sum_{s\in\mathcal{S}_{\mathrm{R}}} w_s \mathbb{I}\left[g_j^{(s)}(\mathrm{x}) > 0\right], \tag{7.5}$$

where $\mathbb{I}[\cdot]$ is the indicator function.

A probabilistically enforced constraint is satisfied when

$$\widehat{p}_j(\mathrm{x}) \leq \varepsilon_j, \tag{7.6}$$

where ε_j is the admissible violation probability. The associated nominal satisfaction level is $1 - \varepsilon_j$. The experiments use

$$\varepsilon_j \in \{0.10,\ 0.05,\ 0.01\}, \tag{7.7}$$

corresponding to nominal satisfaction levels of 90 %, 95 %, and 99 %.

7.4. NSGA-III implementation

The NSGA implementation follows the canonical algorithm [35] with three deviations motivated by the TPSS evaluator.

Conservative seeding. Generation zero is sampled from $\mathrm{P_{CPL}} \in [0,\, 500]^{12}$ kW with $\mathrm{t} \in \{1, 2, 3\}^{12}$ (tap settings {-2.5%, 0%, +2.5%}), rather than from the full decision box. The restriction is released from generation one onward. The reason is mechanical: at the full upper bound $\overline{P}_{\mathrm{CPL}}$ a substantial fraction of generation-zero individuals are power-flow-divergent, and divergent individuals carry no useful information for the first non-dominated sort; they enter selection as infeasibility-dominated by every convergent peer regardless of objective vector. The restricted seed produces a generation-zero population that is mostly convergent.

Two-step AC/DC fitness evaluation. Each scenario evaluation iterates the DC power-flow solver and the AC simulator once. The first DC solve assumes nominal substation voltages; its computed substation currents are passed to the AC simulator, which returns updated substation voltages reflecting the AC-side loading; the DC solve is then re-run with these updated voltages, and the second-pass state is used for fitness. The single iteration costs one additional DC solve per scenario but removes the optimistic bias of assuming load-independent rectifier voltages, which is most pronounced in the high-CPL regime where the AC-side reactive demand becomes non-negligible.

7.5. SMPSO implementation

Speed-constrained Multi-objective Particle Swarm Optimisation (SMPSO) [36] is included as a comparator on the lower-dimensional formulations. The canonical velocity update,

$$\mathrm{v}_i^{k+1} = w\,\mathrm{v}_i^k + c_1 r_1(\mathrm{p}_i - \mathrm{x}_i^k) + c_2 r_2(\mathrm{g} - \mathrm{x}_i^k), \tag{7.8}$$

is retained, with p_i the per-particle best, g a leader drawn from the external archive of non-dominated solutions, w the inertia weight, c_1, c_2 the cognitive and social coefficients, and $r_1, r_2 \sim \mathcal{U}(0,1)$. Positions advance as $\mathrm{x}_i^{k+1} = \mathrm{x}_i^k + \mathrm{v}_i^{k+1}$. SMPSO's distinguishing element is the velocity-magnitude clip applied after (7.8): v_i^{k+1} is rescaled per-dimension to a fraction of the decision-space width, which prevents the swarm divergence that uncontrolled velocity can produce on high-dimensional problems [36].

Archive-based leader selection. The leader g is sampled from the external archive with probability inversely proportional to local archive density, computed via crowding distance [72]. The archive size is capped; archive truncation when the cap is reached removes the most crowded entries first, preserving Pareto-front diversity.

Mixed-integer decoding. The tap positions $\mathrm{t} \in \mathcal{T}^{12}$ are discrete and the CPL block is box-constrained. Velocities and positions are clipped to $[\mathrm{x}_{\mathrm{lo}}, \mathrm{x}_{\mathrm{up}}] = [0;\ \mathrm{t}_{\min}1_{12}]\times[\overline{P}_{\mathrm{CPL}}1_{12};\ \mathrm{t}_{\max}1_{12}]$ after each update. The tap block is rounded to the nearest integer at evaluation time only; the underlying continuous representation is preserved through the velocity update, which avoids aliasing the swarm dynamics onto a discrete lattice.

The CCP constraint aggregation of section 7.3 carries over unchanged: it modifies how feasibility is scored, not how the swarm searches. The two-step AC/DC fitness evaluation of section 7.4 likewise applies without modification.

7.6. Conclusion

This chapter set up the optimisation that allocates external load across the twelve substations. The problem is non-convex, mixed-integer, and evaluated through a non-differentiable simulator, which rules out gradient-based and convex-relaxation methods and motivates the two metaheuristics used: NSGA-III and SMPSO. Both search the same decision vector of per-substation CPL allocations and tap positions against the objective vector of per-substation CPL, loading-ratio imbalance, and total utility intake.

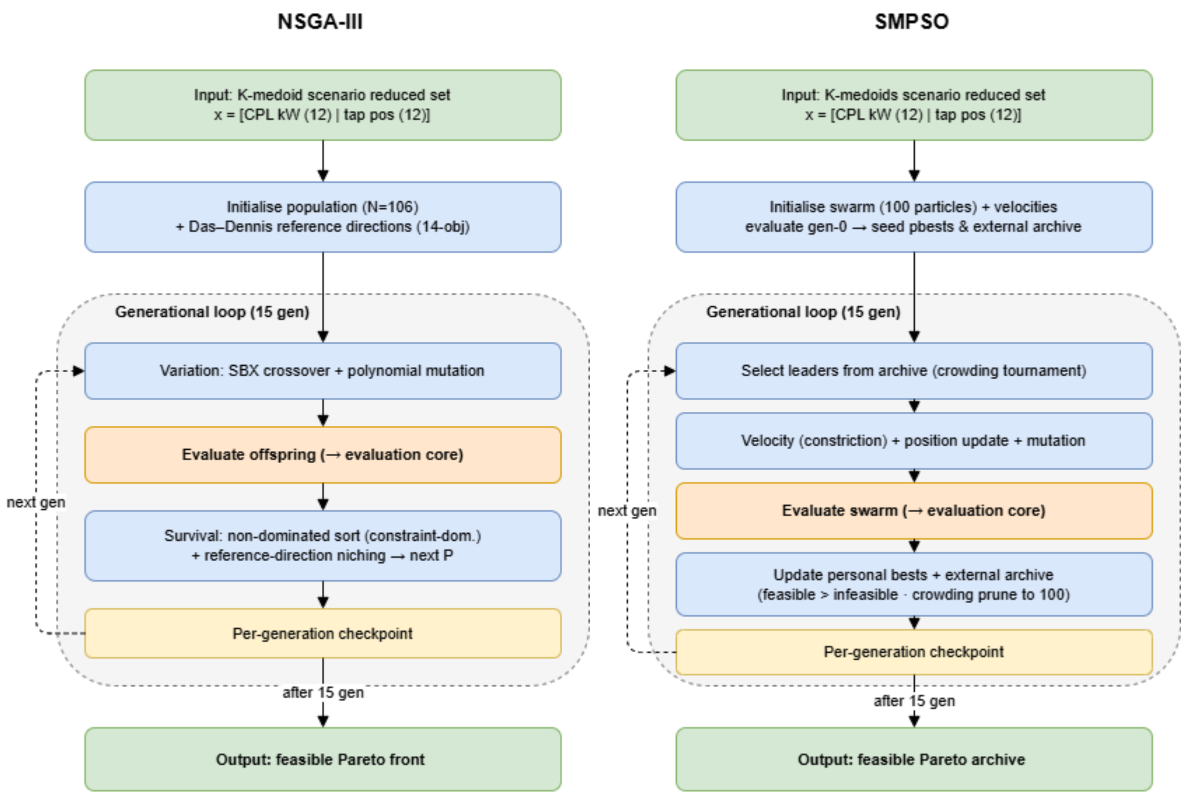

Figure 7.1: NSGA and SMPSO systematic overview

The duration-dependent and contractual constraints are enforced probabilistically through chance constraints at three confidence levels (90%, 95%, 99%), chosen to span the operator's plausible risk tolerance while keeping the optimisation tractable. Two implementation choices adapt the canonical algorithms to the traction evaluator and apply to both optimisers: conservative seeding, which keeps generation zero mostly convergent by restricting the initial CPL range, and a two-step AC/DC fitness evaluation, which removes the optimistic bias of assuming load-independent rectifier voltages at the cost of one extra DC solve per scenario.

This chapter produces the optimised allocations but not their verdict. Whether the fronts are converged, how the two optimisers and the two reduced sets compare, and what the achievable capacity is are evaluated in chapter 8.

8
Results

8.1. Introduction

This chapter presents the findings of the feasibility analysis pipeline as implemented in the preceding chapters. The results are presented as a quantitative comparison between decision variables (scenario reduction type, optimisation algorithm, risk-factor) across a set of metrics. Based on the research questions, formulated in Section 1.3, the following metrics are evaluated:

1. Total permissible constant power load connected to the RET Calandlijn traction grid
2. Spatial distribution of power allocation across the grid
3. Violation probability within the power system
4. Energy flows within the traction power system with CPL perturbation and reconfigured tap settings - Utility energy flow, regenerative braking energy losses, distribution losses, voltage/current profiles

The decision space consists of 2 scenario reduction methods (distribution driven, optimality-aware), 2 optimisation algorithms (genetic, swarm optimisation), and 3 confidence levels for chance-constrained programming (0.90, 0.95, 0.99). This yields 12 combinations of the decision variables, which in turn yield 641 feasible Pareto (nondominated) designs: 141 from NSGA-III and 500 from SMPSO, 362 evaluated on the distribution-driven (distrib) medoid set and 279 on the optimality-aware (optaware) set. The result set consists of 2 steps; first, optimisation is applied and the CPL and tap settings are recorded. These results are then **validated** in a 3 hour simulation after which constraint violations and energy flows are recorded. Unless stated otherwise, the validation runs are evaluated at a 95 % confidence level ($\varepsilon = 0.05$).

8.2. Scenario reduction

8.2.1. Power allocation - quantitative comparison

As illustrated in Figure 6.3, optimality-aware scenario reduction overrepresents the scenario set at tail-end behaviour (peak substation loading) and underrepresents the median. The distribution driven method does the opposite, with higher representation near the median, but less tail end awareness. The optimisation (pre-validation) indicates that the distribution driven set explores a total CPL range of (629-2583) with the median at 1540 kW, while the optimality-aware method peaks at 2401 kW, but contains a higher median at 1870 kW (Figure 8.1a).

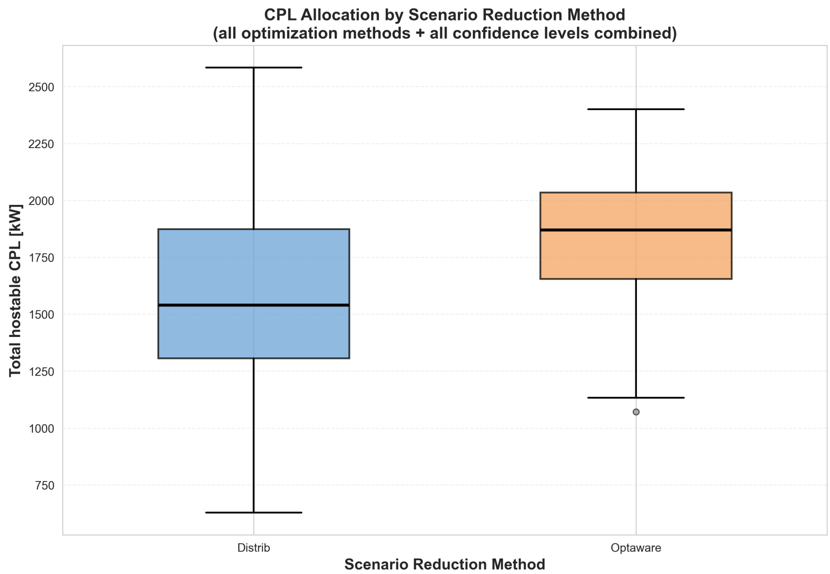

(a) Aggregated CPL per scenario reduction method

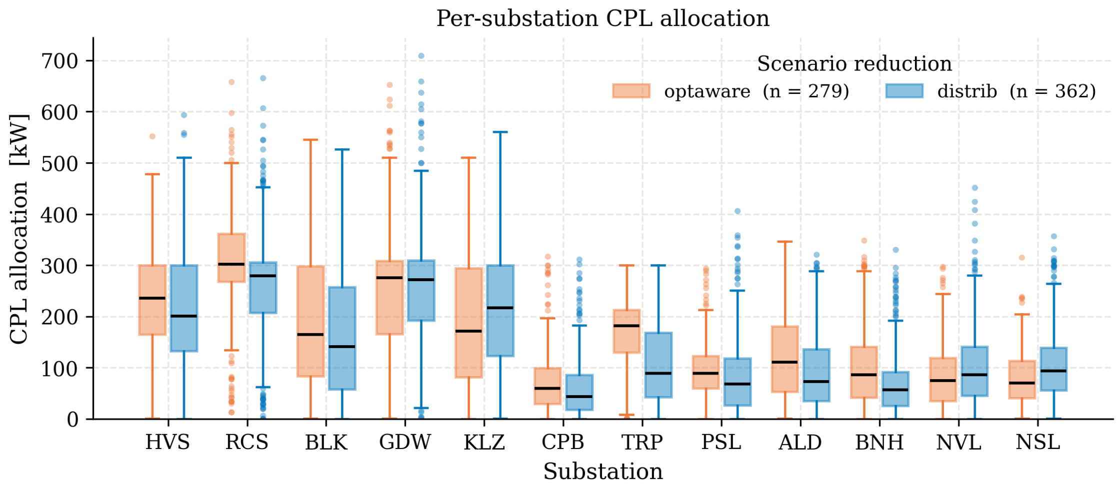

(b) Per-substation representation of results

Figure 8.1: CPL results by scenario reduction, optimality-aware vs distribution driven. Each box represents the interquartile range (IQR), the line represents the median, and the whiskers represent 1.5 × the IQR. This setting is applied to all subsequent boxplots.

Upon inspection of the per-substation comparison presented in Figure 8.1b both methods tend to explore a solution range within the same bounds. Except for some substations (NVL, PSL, TRP) both methods produce a distribution with the median spaced at most 30 kW within each other. This behaviour extends to the extreme / bounding solutions: the distribution driven method finds a bound difference greater than 100 kW at 2 out of 12 substations, otherwise all extremes fall within a range within 30 kW.

8.2.2. Validation results

Both methods produce constraint violations within 0.29 % of each other, with the optimisation aware method at the lower end of 0.04% less violations in minimum vehicle voltage and utility violations. Comparing this to the baseline, the optimality aware method performs on average the same as the baseline of no extra host load Table 8.1.

Table 8.1: Constraint-violation comparison between scenario-reduction methods, aggregated over the six validated designs per method (2 optimisers $\times\ \varepsilon \in \{0.01, 0.05, 0.1\}$): mean and worst dynamic exceedance rate per constraint, and number of designs honouring the constraint at their design ε. Baseline column: exceedance rate of a no-CPL co-simulation against the same limits.

	Baseline	Distribution-driven			Optimality-aware		
Constraint	rate	mean rate	worst rate	pass	mean rate	worst rate	pass
g_1 voltage	0.09%	0.13%	0.19%	6/6	0.09%	0.18%	6/6
g_3 current	0.00%	0.00%	0.00%	6/6	0.00%	0.00%	6/6
g_4 LR 1-min	0.00%	0.00%	0.00%	6/6	0.00%	0.00%	6/6
g_6 utility	0.06%	0.19%	0.29%	6/6	0.15%	0.28%	6/6

Table 8.2: Energy-flow comparison between scenario-reduction methods over the 3-h co-simulation, against the no-CPL baseline. Each method column lists the six validated best-CPL designs per method (2 optimisers $\times\ \varepsilon \in \{0.01, 0.05, 0.1\}$), sorted ascending; these are the individual validated results, not a confidence interval. RBE waste is regenerative braking energy dissipated in the on-board choppers.

Metric	Baseline	Distribution-driven	Optimality-aware
CPL delivered [MW]	0	1.86, 2.06, 2.36, 2.39, 2.58, 2.75	1.13, 2.11, 2.32, 2.40, 2.42, 2.59
CPL energy [MWh]	0.00	5.59, 6.20, 7.09, 7.16, 7.75, 8.23	3.40, 6.32, 6.97, 7.20, 7.25, 7.77
Utility energy [MWh]	12.60	17.68, 18.09, 19.12, 19.30, 19.96, 20.45	15.51, 18.01, 18.90, 19.04, 19.19, 19.96
RBE waste [MWh]	2.20	1.40, 1.41, 1.41, 1.52, 1.53, 1.53	1.34, 1.38, 1.43, 1.45, 1.51, 1.63
DC losses [MWh]	1.00	1.13, 1.16, 1.17, 1.22, 1.26, 1.30	1.07, 1.13, 1.16, 1.16, 1.18, 1.25
Min vehicle voltage [V]	415	407, 422, 429, 438, 442, 456	429, 430, 436, 452, 468, 468

8.2.3. Utility feeder

From the utility (DSO) intake constraint perspective, both methods are within 5% of each other (Figure 8.2). The NSL feeder proves to be the binding constraint; both methods violate this near the 25th minute in simulation. The instantaneous violation plots are deferred to section A.4.

8.2.4. Verdict: scenario reduction

Optimality-aware scenario reduction provides a more faithful picture of the constraint-binding tail events, and this is visible in the validated designs: dynamic exceedance rates are at or near the no-CPL baseline (0.09 % on vehicle voltage, Table 8.1), vehicle voltage margins are

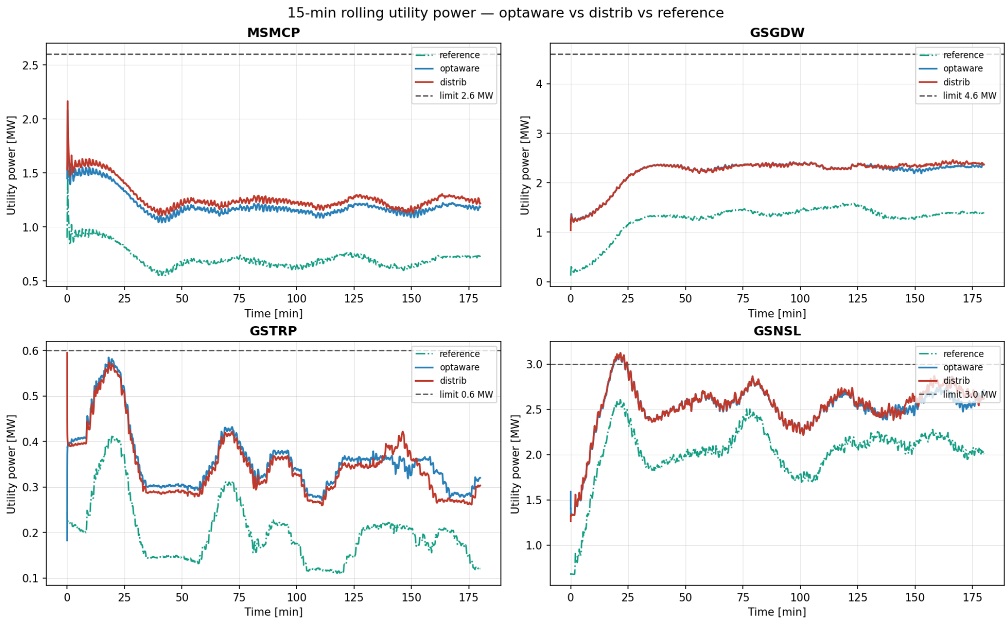

Figure 8.2: Utility feeder - time series analysis for best CPL design for scenario reduction methods against reference (0 CPL)

on average 15 V higher, and DC losses 5 % lower (Table 8.2). The price of this conservatism is hosting capacity: the optimality-aware designs deliver on average 0.17 MW (7 %) less CPL than their distribution-driven counterparts. The distribution-driven method, informed predominantly by median operating conditions, places larger allocations and sails closer to the limits; yet never breaches them: all six designs per method honour every enforced constraint at their design ε, so the difference between the methods is one of margin, not of feasibility.

The choice therefore reduces to risk posture rather than validity. Where the objective is to certify a conservative hosting capacity without relying on downstream verification, optimality-aware reduction is preferable: its overrepresentation of peak-loading scenarios means the optimiser sees the binding events that a dynamic replay would otherwise have to reveal. Where maximum capacity is the priority and a validation stage is part of the workflow anyway, distribution-driven reduction extracts roughly 7 % more CPL from the same network at a still-negligible violation prevalence (≤0.29 % worst case). Notably, the headline system benefit is insensitive to the choice: both methods recover the same ~ 33 % of otherwise-dissipated regenerative braking energy, confirming that RBE recovery is a property of hosting CPLs at all, not of how the scenario set was reduced.

Examining the aggregate and per-substation CPL metrics together, the clearest distinction between the two reductions is how they explore the solution space. The optimality-aware reduction works in a tighter feasible region, keeping a higher median allocation per solution with fewer outliers than the distribution-driven method, so it favours a more homogeneous allocation that is not necessarily less aggressive. The per-substation CPL boxplots show a statistically significant difference between the two reductions, even though the validated energy flows are only marginally more efficient for the optimality-aware designs. That efficiency gain is of the order of 7 %, which is small given the limited set of validated designs, so the methods are better separated on their tail-end behaviour than on their mean energy flows.

The instantaneous over-loading and low-voltage plots in section A.4 make this difference visible. The distribution-driven reduction produces the less reliable designs: transient peak over-loading occurs about 30 % more often, and vehicle-voltage excursions below 500 V are 0.25–3× more frequent than for the optimality-aware designs. These differences are present across the whole low-voltage range rather than at a single point, so they are taken to be significant, even though both methods honour the enforced constraints at their design ε.

8.3. Optimisers

8.3.1. Power allocation - quantitative comparison

Figure 8.3 compares the feasible design pools produced by the two optimisers, with both scenario-reduction methods and all three confidence levels pooled together. NSGA-III contributes the final non-dominated fronts from its runs, giving 141 designs across the six cells. SMPSO contributes its 100-slot archives, giving 501 designs in total. Since the SMPSO archive also contains interior points, its lower pooled mean CPL (1.73 $\pm$ 0.38 MW compared with 2.05 $\pm$ 0.21 MW for NSGA-III) should not be interpreted as weaker performance. It is mainly a consequence of the different set definitions. At the high-capacity end, the two optimisers are very similar: the validated best-CPL designs average 2.28 $\pm$ 0.15 MW for NSGA-III and 2.22 $\pm$ 0.61 MW for SMPSO.

Both optimisers also arrive at broadly the same allocation pattern. Most of the CPL capacity is assigned to HVS and to the four 1440 kW substations RCS, BLK, GDW and KLZ, each receiving around 190–290 kW. By contrast, the six substations connected to the GSNSL utility feeder receive only 60–110 kW each, despite their larger ratings. The limiting mechanism therefore differs by feeder: the GSNSL substations are held back by their shared utility-intake cap rather than by their own ratings, whereas the heavily loaded 1440 kW substations on the MSMCP and GSGDW feeders are pushed close to their traction loading limits (cf. section 8.5). The clearest difference between the two optimisers is at BLK, where NSGA-III allocates 273 $\pm$ 70 kW, compared with 140 $\pm$ 110 kW for SMPSO. The strictest optimality-aware case, $\varepsilon = 0.01$, also shows a difference in robustness: SMPSO retained only one feasible archive member, whereas NSGA-III still produced 21 feasible front designs. This is the only cell in which one optimiser effectively failed to form a useful front.

8.3.2. Pareto front tap setting comparison

Figure 8.4 shows the tap-setting distribution across all substations and all pooled feasible designs. The main result is that neither optimiser uses the tap settings in a strongly systematic way. The pooled means remain close to 0 %, with a spread of roughly $\pm$2 %. SMPSO is more concentrated around the nominal setting, with 0 % being the modal tap at ten of the twelve substations. NSGA-III, in contrast, spreads its designs more evenly over the available $\pm$2.5 % tap band.

The tap settings are also only weakly related to local CPL loading. Across the twelve validated designs, the correlation between tap setting and per-substation loading is $r = +0.06$ (Figure 8.9, bottom right). In practical terms, the optimisers are not preferentially raising the tap at the most heavily loaded substations. The tap position is therefore used more as an

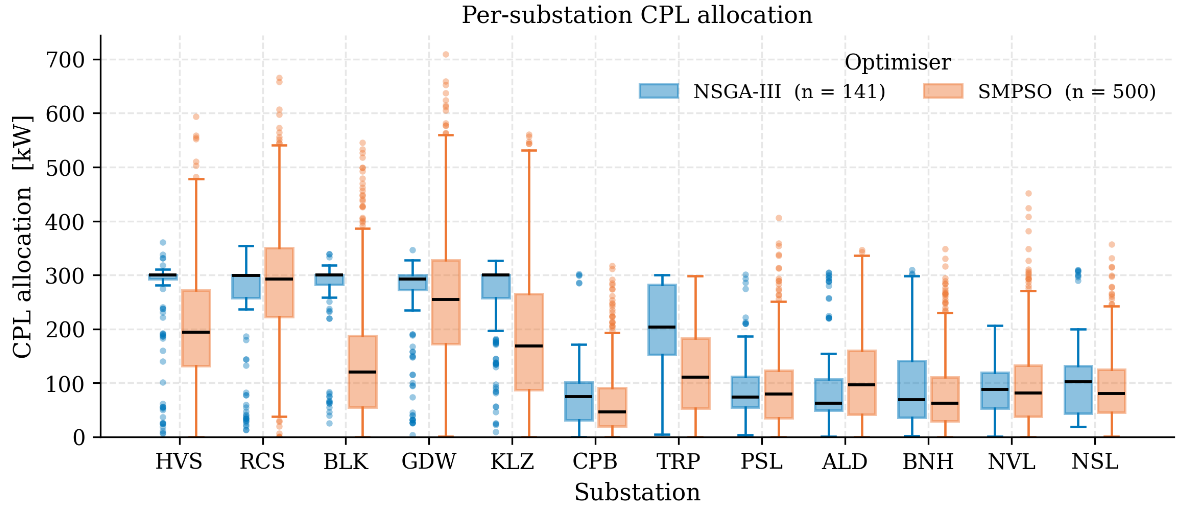

(a) Per-substation CPL allocation by optimiser

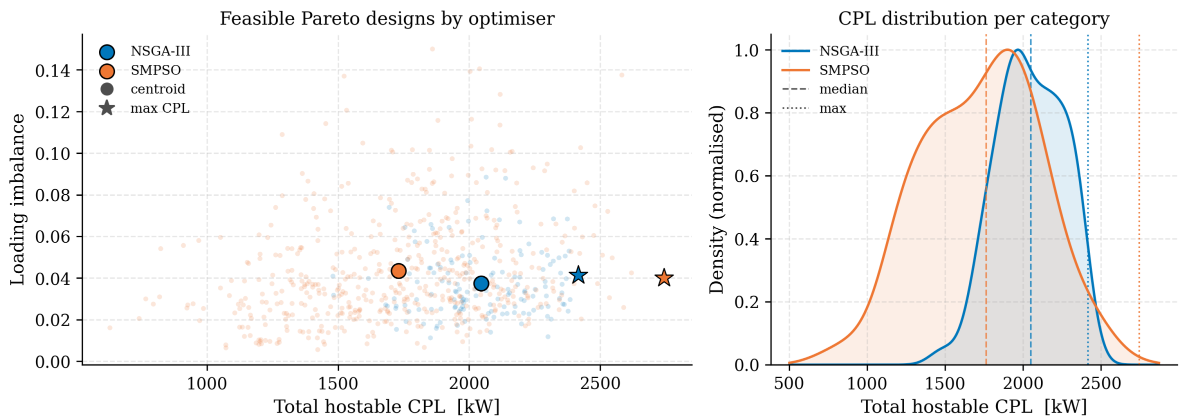

(b) Pareto fronts: total CPL vs allocation imbalance

Figure 8.3: CPL results by optimiser, comparing NSGA-III and SMPSO over both scenario-reduction methods and all confidence levels

independent degree of freedom than as a direct complement to the power allocation. This is consistent with the large voltage margins observed in the dynamic validation (Table 8.6), which leave the optimisers little need to use tap changes to satisfy the voltage constraint.

8.3.3. Validation metrics per optimiser

Table 8.3 compares the validated energy flows for the two optimisers. Utility energy is increased by 49 % for both optimiser. The energy losses (regenerative braking waste and DC losses) are generally higher among the SMPSO results, with the braking energy recovery (acting as CPL utility) are in the order of 10% higher for the NSGA result. This implies that for the tested configurations,$\approx 11\%$ of NSGA load is supplied by braking energy compared to $\approx 10\%$ for the SMPSO results. The recorded vehicle voltage is higher on average for the SMPSO results.

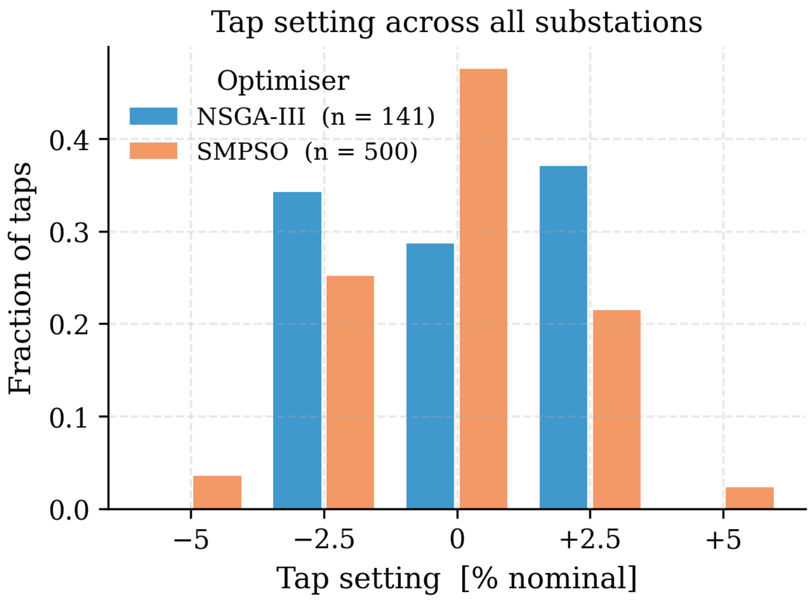

Figure 8.4: Tap-setting distribution across all substations and pooled feasible Pareto designs, by optimiser

Table 8.3: Energy flow comparison by optimiser over the 3-h validation co-simulations, against the no-CPL baseline. Each optimiser column lists its six validated best-CPL designs (2 SR methods $\times$ $\varepsilon \in \{0.01, 0.05, 0.1\}$), sorted ascending; these are the individual validated results, not a confidence interval. RBE recovered is the reduction in chopper-dissipated braking energy relative to baseline.

Metric	Baseline	NSGA-III	SMPSO
CPL delivered [MW]	0	2.06, 2.11, 2.32, 2.36, 2.39, 2.42	1.13, 1.86, 2.40, 2.58, 2.59, 2.75
CPL energy [MWh]	0.00	6.20, 6.32, 6.97, 7.09, 7.16, 7.25	3.40, 5.59, 7.20, 7.75, 7.77, 8.23
Utility energy [MWh]	12.60	18.01, 18.09, 18.90, 19.04, 19.12, 19.30	15.51, 17.68, 19.19, 19.96, 19.96, 20.45
RBE waste [MWh]	2.20	1.34, 1.40, 1.41, 1.43, 1.51, 1.52	1.38, 1.41, 1.45, 1.53, 1.53, 1.63
RBE recovered [MWh]	0.00	0.68, 0.69, 0.77, 0.79, 0.80, 0.86	0.57, 0.67, 0.67, 0.75, 0.79, 0.82
DC losses [MWh]	1.00	1.13, 1.16, 1.16, 1.16, 1.17, 1.22	1.07, 1.13, 1.18, 1.25, 1.26, 1.30
Min vehicle voltage [V]	415	407, 429, 430, 438, 442, 468	422, 429, 436, 452, 456, 468

8.3.4. Verdict: optimisers

Based on the validated CPL designs and their energy-flow metrics, NSGA-III dominates SMPSO on useful regenerative energy, on the reduction in wasted energy, and on the median allocation within the discovered feasible region. The convergent NSGA-III fronts reflect the implicit removal of valid solutions in earlier generations, whereas SMPSO maintains diversity by retaining the best directions found across its whole evaluation history. That diversity does not make SMPSO a clear winner: although it reaches the single highest feasible allocation, the validation shows a significant difference in violation frequency between the two methods (section A.4). The conclusion is therefore that NSGA-III gives the better local optimisation of tap settings and energy conservation, while SMPSO gives the more diverse front.

8.4. Risk analysis

8.4.1. Power allocation - confidence levels

Figure 8.5 and Table 8.4 show the capacity cost of increasing confidence. Hosting capacity decreases monotonically as the chance constraint is tightened. The validated best-CPL designs average 2.53 MW at 90 % confidence, 2.42 MW at 95 %, and 1.79 MW at 99 %. Moving from 90 % to 99 % confidence therefore reduces hosting capacity by roughly 0.7 MW, or about 29 %. The full-front statistics show the same trend, with median CPL values of 2.04, 1.91 and 1.55 MW. The number of feasible designs also falls from 262 at $\varepsilon = 0.1$ to 132 at $\varepsilon = 0.01$. Thus, the strictest confidence level not only limits the maximum capacity, but also gives the optimisers a much smaller feasible region to work with. The relatively large spread in the validated $\varepsilon = 0.01$ column, $\pm$0.39 MW, is mainly due to the single-feasible SMPSO cell discussed in the optimiser comparison.

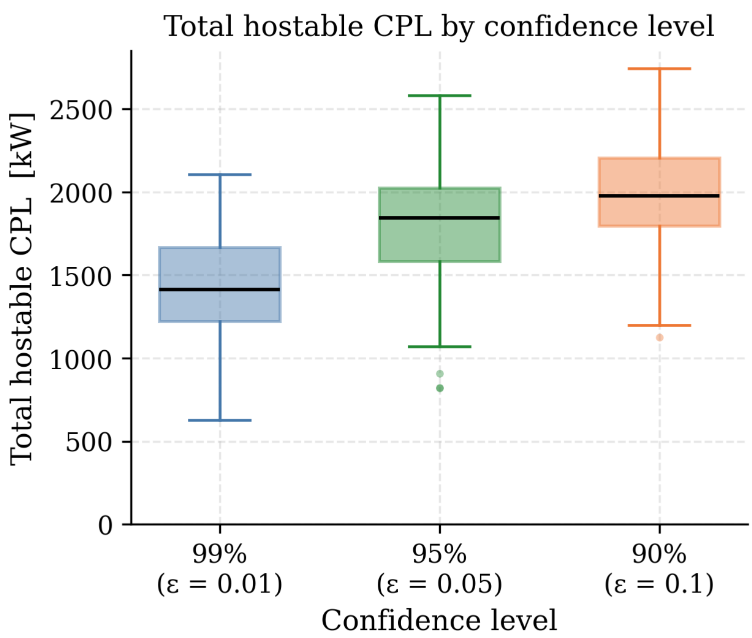

Figure 8.5: Total hostable CPL across the pooled feasible Pareto designs, by confidence level

8.4.2. Utility limits

Table 8.5 reports the dynamic utility-intake metrics against the 10 MW instantaneous limit on the worst feeder. The realised $(1 - \varepsilon)$-quantile changes with the chosen risk level, from 7.60 MW at $\varepsilon = 0.01$ to 5.61 MW at $\varepsilon = 0.1$. The latter is almost the same as the no-CPL baseline value of 5.56 MW. Absolute peaks, however, still exceed the 10 MW limit at every confidence level, reaching 12.9–13.5 MW. The no-CPL baseline also peaks above the limit, at 12.2 MW, which indicates that these exceedances are short traction transients already present in the timetable rather than a direct consequence of adding CPL. Their duration is very small: at most 0.29 % of timesteps, well within the permitted design tolerance of 1–10 %.

Table 8.4: Hosting capacity vs confidence level. Pareto-front statistics pool the four optimisation cells per ε (2 optimisers × 2 SR methods; cell medians averaged to weight cells equally); the validated column lists the four 3-h-validated best-CPL designs per ε (individual values, not a confidence interval).

ε (confidence)	feasible designs	max CPL [MW]	median CPL [MW]	validated CPL [MW]
0.01 (99%)	132	2.11	1.55	1.13, 1.86, 2.06, 2.11
0.05 (95%)	248	2.58	1.91	2.32, 2.36, 2.40, 2.58
0.1 (90%)	262	2.74	2.04	2.39, 2.42, 2.59, 2.75

Table 8.5: Dynamic utility-intake metrics vs confidence level (limit $P_{\text{util}} \leq 10$ MW on the instantaneous worst feeder). Each row lumps the four validated designs per ε; rate = fraction of timesteps above the limit (PASS at rate ≤ ε). The $Q_{1-\varepsilon}$ and peak columns list the four validated designs per ε (individual values, not a confidence interval); rate is the mean and worst rate the maximum over those designs.

ε	$Q_{1-\varepsilon}$ [MW]	peak [MW]	rate	worst rate
Baseline (no CPL)	5.56	12.16	0.06%	–
0.01	7.25, 7.64, 7.66, 7.84	12.39, 12.69, 12.85, 14.57	0.11%	0.13%
0.05	5.92, 5.95, 6.65, 6.81	13.20, 13.32, 13.62, 13.71	0.20%	0.29%
0.1	5.23, 5.60, 5.79, 5.81	12.03, 12.55, 13.29, 13.86	0.21%	0.29%

8.4.3. Grid voltage stability analysis

Vehicle-voltage stability is summarised in Table 8.6. Because the binding constraint is the utility feeder rather than vehicle voltage, the voltage chance constraint is slack at every confidence level: the realised fraction of timesteps below the 500 V floor stays at about 0.1 %, well under even the strictest budget of $\varepsilon = 0.01$, and barely changes with ε. Measured consistently as the 1st-percentile minimum vehicle voltage $V_{1\%}$, headroom is essentially flat across ε (578–584 V) rather than improving; the apparent rise in the earlier Q_ε formulation was an artefact of reading a higher quantile at larger ε. The absolute minima (431–448 V) bracket the no-CPL baseline (438 V), so the deepest dips are driven by coincident vehicle accelerations and hosting up to about 2.7 MW of CPL neither meaningfully worsens nor improves them. All twelve validated designs honour the voltage floor at their design ε.

From the risk-dependent voltage violations in Table 8.6 and the vehicle voltage distribution, there is about a 0.03% difference in violation rate, and in the order of 0.05 % of absolute 99th percentile of lowest voltage differences. This is of little *practical* significance, as voltage fluctuations can range in the order of 30 %. Hence, regarding voltage stability, the fully constrained optimisation is not intrinsically limited by voltage collapses.

Table 8.6: Vehicle-voltage stability vs confidence level (floor $V_{\min} \geq 500$ V; the chance constraint permits a fraction ε of timesteps below it). $V_{1\%}$ is the 1st-percentile per-timestep minimum vehicle voltage, computed identically for every design so it is comparable across ε; rate = fraction of timesteps below the floor (PASS at rate ≤ ε). The $V_{1\%}$ and abs. min columns list the four validated designs per ε (individual values, not a confidence interval); rate is the mean and worst rate the maximum over those designs.

ε	$V_{1\%}$ [V]	abs. min [V]	rate	worst rate
Baseline (no CPL)	583	438	0.09%	–
0.01	580, 581, 588, 588	407, 430, 456, 468	0.09%	0.16%
0.05	573, 574, 584, 585	429, 442, 452, 468	0.12%	0.19%
0.1	575, 575, 578, 584	422, 429, 436, 438	0.12%	0.18%

8.4.4. Violation assessment

Table 8.7 summarises the dynamic exceedance rates across constraints. The realised rates are far below the corresponding design limits. At the strictest level, $\varepsilon = 0.01$, the mean exceedance rates are 0.09 % for voltage and 0.11 % for utility, against a permitted 1 %. At $\varepsilon = 0.1$, they are 0.12 % and 0.21 % against a permitted 10 %. The current and 1-minute loading ratio constraints do not register any dynamic exceedances. All validated designs pass the enforced constraints at their design value of ε under the implemented validation scorecard and among the successfully evaluated dynamic states. It should be noted that vehicle undervoltage is not a dealbreaker under the conditions that it is short lived, and does not hinder the operation. The power flow solver mentioned in [10] decouples load flow from the vehicle simulator, as it assumes that due to the rolling momentum, instantaneous voltage fluctuations have minimal effect on kinematics.

This gives two main observations. First, among the successfully evaluated states, the dynamic replay produces lower physical exceedance rates than the nominal chance-constraint budgets. This indicates conservatism within the solver-converged scenario population. It does not establish unconditional conservatism over the complete generated population, because the physical status of the solver-filtered snapshots is unknown. The medoid scenario set appears to over-represent tail events compared with the 3-h co-simulation, so the realised risk is much lower than the certified bound. Second, the small number of remaining exceedances is largely already present in the baseline traction load: the no-CPL case has a 0.09 % voltage violation rate and a 0.06 % utility violation rate. The hosted CPL therefore adds little additional violation prevalence. In this setting, the limiting issue is hosting capacity, not whether the validated designs satisfy the dynamic constraints.

8.4.5. Verdict: risk level

The confidence level acts on two things at once: hosting capacity and the realised violation prevalence. Its effect on capacity is strongly non-linear. Tightening the chance constraint from 90 to 95 % confidence costs little (2.53 to 2.42 MW), but moving to 99 % removes a further

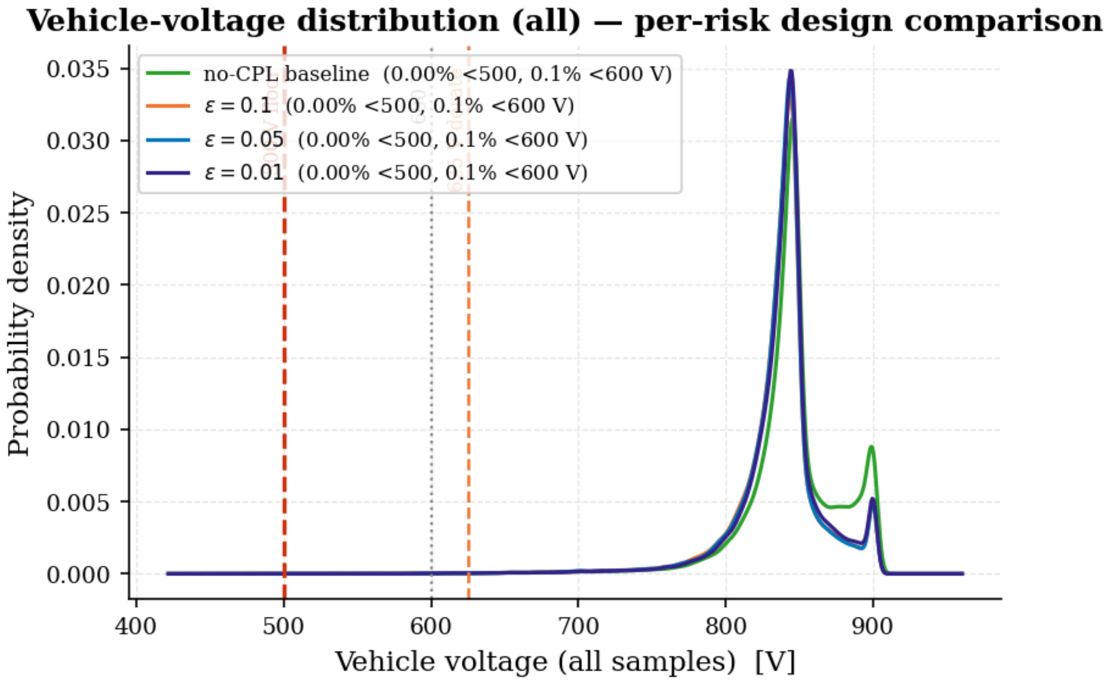

Figure 8.6: KDE representation of vehicle voltage distribution under different confidence design levels. A distinction in the 900V peak voltage between the optimised and baseline configurations reveals that vehicle peak voltage prevalence reduces by $\approx 40\%$ when interfaced with constant power loads, while the 600V useful voltage boundary is breached less than 0.1 % in all configurations.

Table 8.7: Violation assessment vs confidence level: mean dynamic exceedance rate per constraint over the four validated designs per ε, and designs passing at their design ε. Baseline column: no-CPL co-simulation.

	Baseline	$\varepsilon = 0.01$		$\varepsilon = 0.05$		$\varepsilon = 0.1$	
Constraint	rate	rate	pass	rate	pass	rate	pass
g_1 voltage	0.09%	0.09%	4/4	0.12%	4/4	0.12%	4/4
g_3 current	0.00%	0.00%	4/4	0.00%	4/4	0.00%	4/4
g_4 LR 1-min	0.00%	0.00%	4/4	0.00%	4/4	0.00%	4/4
g_6 utility	0.06%	0.11%	4/4	0.20%	4/4	0.21%	4/4

0.6 MW and roughly halves the feasible design count, so most of the sellable capacity sits at moderate confidence.

The risk level also changes how hard the network is driven. The enforced constraints are honoured at every level. The 1-minute loading ratio and the substation current never violate, and the voltage floor is breached in only about 0.1 % of timesteps, but the transient stress scales clearly with ε. The instantaneous peak loading above 2.5 pu falls from about 8.3 % of timesteps at $\varepsilon = 0.1$ to 7.1 % at $\varepsilon = 0.05$ and 4.5 % at $\varepsilon = 0.01$ (section A.4), so the strictest setting roughly halves the transient over-loading relative to the loosest, and the low-voltage tail follows the same ordering. These excursions concentrate at the same four 1440 kW substations and remain sub-second transients on the instantaneous peak rather than breaches of the enforced 1-minute constraint, but they are a real, ε-dependent measure of how close the network runs to its limits. The binding quantity differs by feeder: on the GSNSL and GSTRP connections it is the utility intake, against whose tail the confidence level reserves headroom, while on the MSMCP and GSGDW feeders it is the traction-side loading of the small 1440 kW substations. Tightening ε produces more conservative allocations that reduce this transient stress. The choice of ε is therefore both a commercial trade-off between certified capacity and conservatism and a lever on the realised over-loading.

8.5. Spatial distribution of load allocation

Figure 8.7 presents the general relation between substation pairs (fed by a corresponding utility connection) and their hosting capacity. The upper bound is defined by CPL-stratified Pareto fronts, and indicates the trade-off in hosting capacity when configuring a desired load allocation. The corresponding plots can be found in the appendix under Figure A.4.

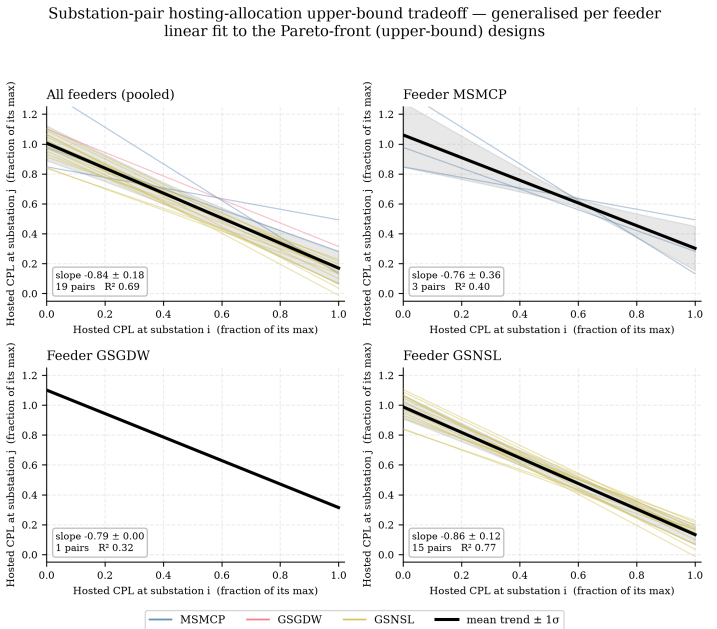

Figure 8.7: Spatial relation between substation pairs per feeder. Each substation pair is regressed against another interfaced by a common utility feeder. The axes represent the fraction of the total theoretical value found for CPL allocation at which a substation can host external loads. The regression line coefficient for the upper bound sits between -0.7 and - 0.9 per substation pair, which implies that a increasing the host allocation by 10 % potentially shaves 8% of the capacity of the other. This is derived from the pareto fronts as illustrated in Figure A.4.

8.6. Technical feasibility

The preceding sections report hosting capacity under the full constraint set. In that formulation, the utility feeder limits represent contractual intake ceilings agreed with the DSO, rather than physical limits of the traction network itself. To separate what the infrastructure can technically host from what the supply contracts permit, this section considers the SMPSO optimisation with the utility constraints removed. The remaining constraints are therefore only the traction-side limits: the vehicle voltage floor, substation current, and loading-ratio envelope.

Figure 8.8a shows the resulting CPL distribution across the 100-design Pareto front. Without the feeder ceilings, the network hosts a median of 3.36 MW and a mean of 3.38 MW, with solutions ranging from 1.6 to 4.7 MW. This is roughly 0.7–2 MW above the feeder-constrained optimum of 2.74 MW. The extra capacity is not spread evenly: it comes almost entirely from the GSNSL substations whose utility contract was binding, while the heavily loaded MSMCP and GSGDW substations are already close to their traction limits and gain little. This is also supported by the dynamic replays of the two max-CPL reference designs, at 3.25 and 3.84 MW (Figure 8.9). These cases still operate with absolute minimum vehicle voltages of 436–448 V, which lie within the same range as the validated feeder-constrained designs. The main difference is loading: in the unconstrained case, the worst-loaded substation operates at or above rated power on average.

Figure 8.8b shows how this technical capacity is distributed across substations. The result differs clearly from the feeder-constrained allocation in Figure 8.3a, where the six GSNSL-feeder substations are restricted to only 60–110 kW. Once the feeder constraints are removed, these GSNSL substations rise sharply, to mean allocations between 230 and 353 kW and individual values reaching the 800 kW design bound; NSL, for example, averages 353 kW with a maximum of 799 kW and BNH averages 331 kW. The allocations on the MSMCP and GSGDW feeders change little by comparison, because the small 1440 kW substations there already run close to their traction loading limits under the contracts. The contractually suppressed capacity is therefore concentrated on the GSNSL group, and to a lesser extent GSTRP.

The practical implication is that the gap between the technical hosting capacity, around 3.4 MW median, and the permitted capacity, around 2.7 MW, is contractual on the GSNSL and GSTRP feeders but not elsewhere. Additional capacity is most cheaply found by raising the GSNSL and GSTRP intake caps; the MSMCP and GSGDW substations are already traction-limited and would need traction-side reinforcement rather than a contract change to host more. There is still a limit to the value of this extra headroom. Beyond roughly 3 MW, the chopper–conductor crossover in Figure 8.9 indicates diminishing energetic returns: additional CPL is then hosted in the transport-limited regime, with increasing conduction losses and very little remaining loading margin.

8.7. Conclusion

8.7.1. Objective findings

The quantitative results of the chapter are listed below and collected in Table 8.8 and Table 8.9.

- All twelve validated best-CPL designs pass the four enforced constraints (g_0 vehicle voltage, g_1 substation current, g_2 1-minute loading ratio, g_4 utility intake) at their design confidence level.

- Validated hosting capacity is 2.53 MW at 90 % confidence, 2.42 MW at 95 %, and 1.79 MW at 99 %, with a pooled feasible maximum of 2.74 MW. The number of feasible designs falls from 262 at $\varepsilon = 0.1$ to 132 at $\varepsilon = 0.01$.

- Dynamic exceedance rates stay at or below 0.21 % for every constraint and confidence level: vehicle voltage 0.09–0.12 %, utility intake 0.11–0.21 %, current and 1-minute loading ratio 0 %. The no-CPL baseline already carries 0.09 % voltage and 0.06 % utility

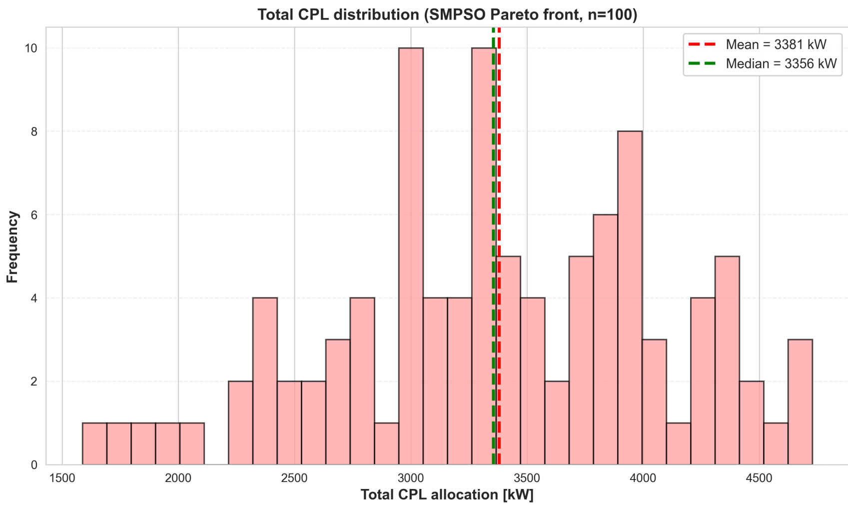

(a) Total hostable CPL distribution across the SMPSO Pareto front with the utility feeder constraints removed (n = 100): technical hosting capacity of the traction network under voltage, current and loading limits only

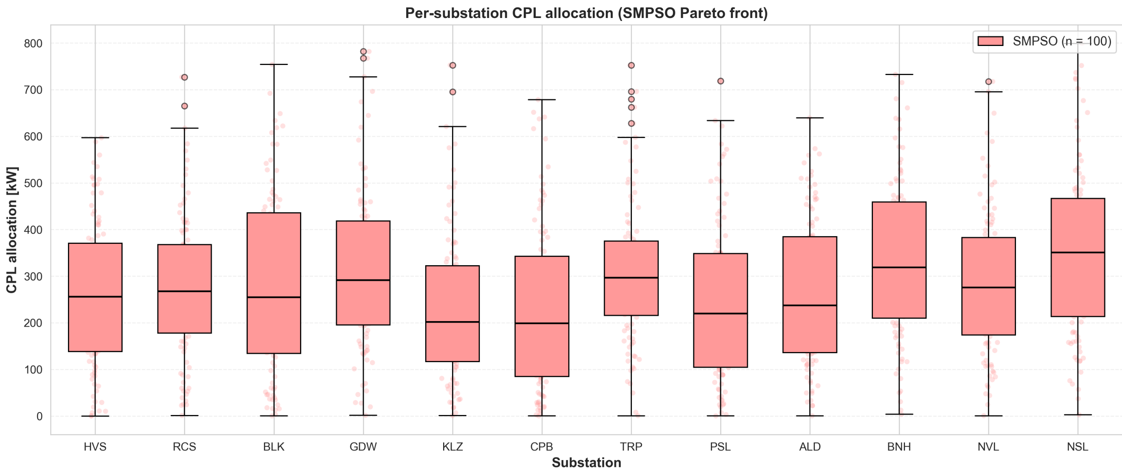

(b) Per-substation CPL allocation across the same feeder-unconstrained SMPSO Pareto front: all twelve substations are technically capable hosts, with allocations reaching the 800 kW per-substation design bound

Figure 8.8: Technical hosting capacity of the traction network with the utility feeder constraints removed (SMPSO, 100-design Pareto front): total hostable CPL (a) and its per-substation distribution (b).

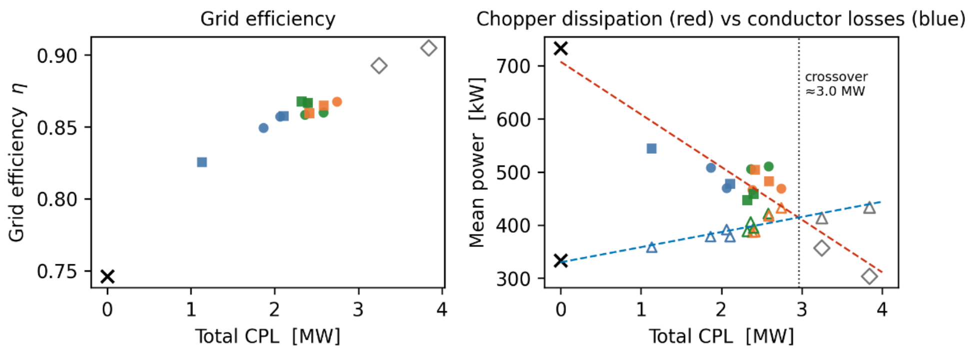

Figure 8.9: The same trade-offs under the chance-constrained formulation with feeder limits active, based on the validated 3-h co-simulations. Hollow diamonds indicate the feeder-unconstrained max-CPL reference designs, which lie outside the feeder-constrained feasible set.

exceedances.

- Instantaneous peak loading above 2.5 pu occurs in 8.3 % of timesteps at $\varepsilon = 0.1$, 7.1 % at $\varepsilon = 0.05$, and 4.5 % at $\varepsilon = 0.01$. These transient excursions occur almost entirely at the four 1440 kW substations BLK, GDW, KLZ and RCS.
- The distribution-driven reduction delivers about 7 % more capacity than the optimality-aware reduction (2.33 vs 2.16 MW) but shows about 30 % more transient over-loading; the optimality-aware runs hold about 15 V more voltage margin and 5 % lower DC losses. Both recover about 33 % of the regenerative braking energy.
- NSGA-III and SMPSO validate to 2.28 and 2.22 MW. NSGA-III recovers about 10 % more braking energy, and at the strictest optimality-aware cell SMPSO retains one feasible design against 21 for NSGA-III.
- Grid efficiency rises from 0.75 without CPL to 0.87 across the validated range, and the chopper–conductor loss crossover lies near 3 MW.
- With the feeder constraints removed the network hosts a median of 3.36 MW (range 1.6–4.7 MW); the gain is concentrated on the GSNSL substations, which rise from 60–110 kW under the contracts to 230–353 kW, while the MSMCP and GSGDW substations change little.

Table 8.8: Influence of the three design factors on the comparison criteria, with the preferred setting for each. Capacities are validated best-CPL figures.

Design factor	Hosting capacity	Spatial allocation	Violations / risk	Energy	Preferred setting
Scenario reduction (distrib vs optaware)	distrib ~7 % higher (2.33 vs 2.16 MW)	distrib higher per-substation peaks	optaware at baseline rates; distrib ~30 % more transient over-loading	both recover ~33 % RBE	optaware to certify; distrib for maximum capacity
Optimiser (NSGA-III vs SMPSO)	near-equal (2.28 vs 2.22 MW); SMPSO holds the single highest	same feeder-driven pattern	NSGA-III more robust in the strictest cell	NSGA-III recovers ~10 % more braking energy	NSGA-III; SMPSO for diversity / verification
Confidence level ($\varepsilon = 0.10/0.05/0.01$)	2.53 / 2.42 / 1.79 MW; non-linear cost	essentially unchanged	enforced constraints pass; transient over-loading scales with ε (~8.3 to 4.5 % >2.5 pu, $\varepsilon = 0.1 \to 0.01$)	essentially unchanged	90–95 % (best capacity per unit risk)

Table 8.9: Contractual (feeder-constrained) versus technical (feeder-unconstrained) hosting capacity, by comparison criterion.

Criterion	Feeder-constrained (contractual)	Feeder-unconstrained (technical)
Hosting capacity	2.4–2.5 MW at 90–95 % confidence (best 2.74 MW)	3.4 MW median (1.6–4.7 MW)
Binding constraint	utility intake on GSNSL/GSTRP; traction loading on MSMCP/GSGDW	traction-side loading ratio
Per-substation allocation	GSNSL substations capped at 60–110 kW	all twelve substations 230–353 kW
Energy regime	below the chopper–conductor crossover; net loss reduction	at or above the ∼3 MW crossover; diminishing returns
Route to more capacity	renegotiate feeder caps (GSNSL, GSTRP)	traction-side reinforcement (limited value)

8.7.2. Interpretation and implications

The binding resource is not the same across the network. On the GSNSL and GSTRP feeders the limit is the contractual utility intake: these substations sit well below their ratings and gain the most when the caps are lifted. On the MSMCP and GSGDW feeders the limit is the traction side, where the small 1440 kW substations run close to their loading limits and host similar amounts with or without the contracts. Hosting capacity is therefore utility-limited for part of the network and infrastructure-limited for the rest, and within each feeder it behaves as a shared resource. Connection requests can be assessed feeder by feeder, and the cheapest additional capacity lies in renegotiating the GSNSL and GSTRP intake caps; the MSMCP and GSGDW substations would instead need traction-side reinforcement. The feeder-unconstrained study confirms the split: lifting the caps raises the GSNSL allocations sharply while leaving the already-loaded substations little changed.

The confidence level is both the main capacity lever and a control on how hard the network is driven. Relaxing it buys capacity cheaply at first and expensively near certainty, and it raises the prevalence of transient over-loading; tightening it does the reverse. The confidence setting is therefore a deliberate trade-off between certified capacity, conservatism and the transient stress the network experiences, not a feasibility switch, since every certified design stays within the enforced limits.

The scenario-reduction method sets the risk posture rather than the feasibility. The optimality-aware reduction is informed by the loading tail and yields the more reliable designs, with fewer transient over-loads and deeper voltage margins, while the distribution-driven reduction reaches further into the high-capacity region and still satisfies every enforced constraint. The choice is one of intent: optimality-aware reduction for conservative certification, and distribution-driven reduction for maximum capacity when a validation stage is available to catch the difference.

The optimiser choice barely affects the certified outcome. NSGA-III converges to tighter, energy-efficient allocations and is the more dependable under the strictest constraint, whereas

SMPSO maintains a broader, more diverse front. NSGA-III is the natural primary solver and SMPSO a useful independent check; disagreement between them is most telling as a sign that a cell is over-constrained.

For the operator, the practical message is that a substantial share of external load can be hosted under the existing contracts with little added violation exposure and a net reduction in the traction system's own energy losses. Where the remaining headroom can be unlocked depends on the feeder: on GSNSL and GSTRP it is a commercial question at the utility intake, whereas on the MSMCP and GSGDW feeders the substations are already traction-limited and would need infrastructure reinforcement to host more. The energetic benefit persists only while the added load keeps absorbing regenerative braking energy; past that point the network shifts into a transport-limited regime where losses, not contracts, set the limit.

8.8. Limitations

Model accuracy

The traction grid model is not accurate by itself; many simplifications were made to keep the pipeline computationally tractable. The power flow is quasi-static at 1 Hz, vehicles are point loads on lumped feeder impedances, and the rectifier, chopper and regenerative-braking behaviour are idealised characteristics rather than measured converter models. Timetable-driven traffic stands in for real operation. None of these simplifications has been checked against measurements: external validation against operational SCADA data is out of scope (section 4.5), so any systematic model error propagates into the certified capacity unchecked. In addition, the DC power-flow solver diverges on about 2 % of evaluations. Diverged snapshots are treated as infeasible, which pushes the search in the conservative direction but adds noise to the constraint quantiles near the feasibility boundary.

Optimisation budget

The optimisation was not evaluated until convergence. Each cell ran for roughly 1 000–1 100 evaluations (about ten generations at the population sizes used); running the matrix to 50 generations was computationally impractical within the pipeline. The reported fronts are therefore snapshots of an unconverged search, the capacity figures are lower bounds on what the formulation admits, and the optimiser comparison reflects behaviour under an equal, limited budget rather than asymptotic quality. The feasibility starvation observed for SMPSO at the strictest confidence level may partly be a budget effect.

Validation scope

Dynamic validation rests on a single 3-h chronological window per design under one timetable and headway. Seasonal demand variation, degraded operating modes and substation outages are not exercised. The realised exceedance rates of 0–0.3 % are estimated from 10 800 one-second timesteps, so they are resolved to roughly ± 0.03 % at best, and the verdicts at $\varepsilon = 0.01$ rest on tens of exceedance events. The no-CPL violation baseline comes from a 5-h replay; exceedance rates are duration-normalised and comparable, but peak values between runs of different length are not strictly so. The 15-minute contractual feeder cap was enforced during optimisation but validated only through the instantaneous feeder limit in the dynamic scorecard.

The hosted load

The CPL is modelled as an ideal, time-invariant constant-power sink. Real candidate loads, EV charging foremost, are stochastic, diurnal and potentially controllable. Treating them as constant ignores both the extra capacity that load flexibility could unlock and the coincidence risk of correlated charging peaks. Protection coordination under reverse-flow conditions at high CPL is not examined.

Constraint handling under optimisation

The constraint handling under uncertainty applied during optimisation is blind to the meaning of the constraint violations. Constraint rates are required to fall within one single defined confidence level, while the severity and practical prevalence of each constraint are not taken into account. An overcurrent violation, which interrupts operation by tripping the switchgear, cannot be treated the same way as a vehicle undervoltage, where vehicle power may simply be restricted temporarily. The loading profiles of substations may also carry different thermal implications across locations, even when the metrics are numerically identical. A contextual constraint background per individual installation would therefore give a clearer picture of what actually limits the system. Violating a contractual limit, for instance, may incur a monetary penalty that is potentially outweighed by the economic benefit of a particular load allocation, without directly reflecting on technical operability.

Numerical reliability of the power-flow solution

About 2% of the initial snapshots failed to yield a converged DC power-flow solution and were removed before scenario reduction; the remaining scenario weights were then normalised. The chance constraints thus reflect the empirical distribution conditional on successful numerical solution.

Power-flow convergence was not tracked by design, timestep, or other attributes after optimisation, so retrospective analysis of non-convergence rates by CPL allocation, tap setting, optimiser, scenario-reduction method, or confidence level is impossible. We therefore cannot tell whether excluded cases were isolated numerical artefacts or systematically linked to electrically stressed states.

Non-convergence does not prove physical infeasibility, and discarding a non-convergent state does not prove feasibility. The reported violation rates and confidence levels apply only within the numerical domain successfully solved by the chosen algorithm. This caveat is especially important for the nominal 99% case, since the initially discarded scenario mass exceeds its 1% risk allowance.

9
Conclusion

This chapter concludes the research questions of section 1.3 using the findings of chapter 8.

9.1. Problem and objectives

This thesis investigated how the load hosting capacity of a traction power supply system can be determined under uncertainty and duration-dependent constraints. The objective was to quantify how much external constant-power load such a network can host without violating its operational and contractual limits, and to identify what limits that capacity. The traction demand is stochastic and the thermal limits are duration-dependent, so the capacity cannot be read off a single worst-case condition; it has to be established across the distribution of operating states the network actually sees.

This objective was pursued through a framework of four stages: modelling the traction network to generate operational scenarios, reducing those scenarios to a representative subset that captures the underlying distribution of operational states, solving a multi-objective chance-constrained optimisation to allocate constant-power load under uncertainty, and validating the resulting allocations against dynamic replays of the network.

9.2. Answering the research questions

RQ1: How can the state of a traction grid be quantified and represented in terms of the energetic quantities that reflect the operability of the traction power supply system?

The state of the traction grid was described using fourteen quantities recorded at each timestep. These include the power, current, and voltage at each substation, which are directly related to the voltage and feeder constraints. Chopper dissipation and transmission losses were included to describe regenerative-energy losses and network efficiency. Vehicle position, speed, and pantograph quantities were used to identify the location and behaviour of each traction load.

The resulting dataset was represented as a collection of static operating snapshots. Each snapshot was supplemented with short-term rolling quantities for substation loading and feeder intake. This retains the information needed to assess loading, voltage, and regenerative-energy losses, without carrying the complete time series into the later stages of the analysis.

The representation therefore provides enough information for scenario reduction and optimisation while keeping the computational problem manageable.

RQ2: Given that determining reliable load hosting capacity under every possible traction network state is infeasible, how can a reduced, representative scenario set be extracted for tractable computation of load allocation limits?

Both scenario-reduction methods produced subsets that led to valid load allocations. Every design that satisfied the optimisation constraints also passed the dynamic replay validation (Table 8.1). The validation procedure was intended to identify solutions that were feasible in the static optimisation but failed during dynamic replay. This failure mode was not observed.

The main difference between the two methods is the amount of margin left in the final allocation. The optimality-aware method selects more states from the constraint-binding tail of the distribution. As a result, it produces more conservative designs, with violation rates close to those of the no-CPL baseline. The distribution-driven method retains a broader sample of the original operating states and generally allows slightly more capacity. Despite this difference, the resulting designs still passed all validation checks.

The effect of the stronger tail weighting is most visible at the 99% confidence level. At this level, the feasible region becomes much smaller for the swarm-based optimiser than for the evolutionary optimiser. The reduction method therefore affects not only the conservatism of the result, but also the difficulty of the optimisation problem. This becomes particularly relevant when the confidence requirement is strict.

RQ3: How can load hosting capacity be computed under hard and probabilistic enforcement of duration-dependent operational constraints, and what does the resulting allocation reveal about the factors that limit it?

The network can host approximately 2.4–2.5 MW of constant-power load at confidence levels of 90–95,%, and approximately 1.8 MW at 99% confidence (Table 8.4). The results show that the hosting capacity is not limited by the same constraint throughout the network. The dominant limitation depends on the feeder and on the substations connected to it.

For the GSNSL and GSTRP connections, the main restriction is the rolling utility-intake limit. During the optimisation, almost all infeasible candidates on these feeders violated this constraint. The six GSNSL substations therefore remain between 60 and 110 kW at every confidence level, even though their installed ratings are relatively large. Their available technical capacity cannot be fully used under the present contractual limits.

The MSMCP and GSGDW feeders are limited by the traction system itself. The smaller 1440 kW substations operate close to their loading limits and host similar amounts of external load with and without the contractual constraints. An increase in contracted intake would therefore have little effect at these locations unless the traction-side infrastructure were also reinforced.

When the feeder constraints are removed, the median technically feasible hosting capacity increases to approximately 3.4 MW. Most of this increase occurs at the GSNSL substations, where the contractual limits restrict capacity that is technically available. The substations that already operate close to their traction-side loading limits show little change (Figure 8.8b).

This distinction between contractual and technical capacity is important. The purpose of the thesis is to determine how much additional load the traction power system can accommodate. A result based only on the current contracts does not fully describe the capability of the physical system. Contractual limits and market arrangements may change, while the technical characteristics of the network remain largely the same. It is therefore useful to report both the contractually available capacity and the technically feasible capacity.

The technically feasible result can support discussions about future contracted capacity and network planning. However, a higher contractual limit cannot be assessed from normal operating conditions alone. The DSO must also consider distribution-level N−1 conditions, including outages and maintenance. These situations may reduce the power available to the traction system and may require lower intake limits than those suggested by the normal operating case. Contingency analysis is therefore needed before a higher contractual limit can be considered reliable.

The choice of optimiser and scenario-reduction method has less influence on the final capacity than the structure of the network itself. Additional load can be accommodated within the operational and contractual limits, but the limiting factor differs by feeder. On the GSNSL and GSTRP feeders, further capacity mainly depends on the contractual intake limits. On the MSMCP and GSGDW feeders, further gains would require reinforcement of the traction-side infrastructure.

9.3. Recommendations

This feasibility study presents the theoretical limit that can be applied to the traction grid case, but several questions remain open.

First, the traction simulation model should be improved. This means accounting for exact track grading and a more detailed rolling-stock representation, since real operation uses different train types, each with its own load-torque curve, and it requires verification against the electrical infrastructure. The constraints should also be mapped more completely: the thermal stability of components does not depend on instantaneous load alone.

Second, on the infrastructure itself, some results show a peak loading above 4 p.u. for individual substations even with no added load. These substations are rated for a peak of only 2.5 p.u., so it is in the operator's interest to schedule them for replacement under a revised programme of requirements. Voltage collapse is not the limiting factor, and tap-setting reconfiguration does not translate into meaningful additional capacity, so this configuration variable is not worth pursuing further in this traction case.

Third, grid protection settings and contingency (N−1) analysis were outside the scope of this study, although they pose an important engineering question: a load allocation must remain reliable under alternative network configurations. This reflects directly on the contractual agreements, which are sized to accommodate present N−1 situations. A case for relieving grid congestion can be made where congestion exists in parallel with the RET grid; the existing congestion on the DSO grid that supplies the traction power system, upstream in the supply chain, cannot be relieved by the measures studied here.

Bibliography

[1] G. Falchetta and M. Noussan, “Electricity consumption patterns within cities: Application of a data-driven settlement characterization method,” *International Journal of Digital Earth*, vol. 12, no. 7, pp. 840–858, 2019, Establishes that urban areas consume approximately 75% of global primary energy supply. DOI: 10.1080/17538947.2018.1556355.

[2] Statistics Netherlands (CBS), “Electrification in the Netherlands 2017–2021,” Centraal Bureau voor de Statistiek, The Hague, Tech. Rep., 2022, Electrification by EVs and heat pumps added 1% to Dutch electricity consumption over 2017–2021. [Online]. Available: https://www.cbs.nl/en-gb/longread/aanvullende-statistische-diensten/2022/electrification-in-the-netherlands-2017-2021 (visited on 05/14/2026).

[3] Enlit World, *Comparing grid crunches and flexibility in the Netherlands and Germany*, Estimated EUR 200B grid investment through 2040., Feb. 2026. [Online]. Available: https://www.enlit.world/library/lessons-from-grid-bottlenecks-in-germany-and-the-netherlands (visited on 05/14/2026).

[4] Taylor Wessing, *Grid capacity in the Dutch energy sector*, Industry insight publication, Peak demand projected 19 GW (current) to 27 GW by 2030., May 2025. [Online]. Available: https://www.taylorwessing.com/en/insights-and-events/insights/2025/05/grid-capacity-in-the-dutch-energy-sector (visited on 05/14/2026).

[5] NL Times, *Power grid in Rotterdam and surrounds now also at max capacity*, Stedin announces grid in Rotterdam-Zuid and surrounding municipalities reached max capacity; reinforcement expected by 2029., Nov. 2024. [Online]. Available: https://nltimes.nl/2024/11/05/power-grid-rotterdam-surrounds-now-also-max-capacity (visited on 05/14/2026).

[6] International Energy Agency, *Grid congestion is posing challenges for energy security and transitions*, IEA Commentary, Approximately 10,000 large users and 7,500 generation projects waitlisted for grid connection in the Netherlands as of early 2025., Mar. 2025. [Online]. Available: https://www.iea.org/commentaries/grid-congestion-is-posing-challenges-for-energy-security-and-transitions (visited on 05/14/2026).

[7] Strategic Energy Europe, *Businesses on a waiting list for power: The Netherlands launches 100 measures to ease grid congestion*, Summary of the LAN Second Progress Report (March 2025): >12,000 companies waitlisted; 100+ targeted actions., Mar. 2025. [Online]. Available: https://strategicenergy.eu/the-netherlands-grid-congestion/ (visited on 05/14/2026).

[8] M. Bartłomiejczyk, L. Jarzebowicz, and R. Hrbáč, “Application of traction supply system for charging electric cars,” *Energies*, vol. 15, no. 4, 2022, ISSN: 1996-1073. DOI: `10.3390/en15041448`. [Online]. Available: `https://www.mdpi.com/1996-1073/15/4/1448`.

[9] I. Diab, G. R. C. Mouli, and P. Bauer, “Increasing the integration potential of ev chargers in dc trolleygrids: A bilateral substation-voltage tuning approach,” in *2022 International Symposium on Power Electronics, Electrical Drives, Automation and Motion (SPEEDAM)*, 2022, pp. 264–269. DOI: `10.1109/SPEEDAM53979.2022.9841989`.

[10] I. Diab, “Drop the droop: Assessing the limitations of droop control in dc traction grids with energy storage and third-party loads,” in *2025 IEEE Transportation Electrification Conference and Expo, Asia-Pacific (ITEC Asia-Pacific)*, 2025, pp. 1–5. DOI: `10.1109/ITECAsia-Pacific63742.2025.11345065`.

[11] K. van der Horst, I. Diab, G. R. Chandra Mouli, and P. Bauer, “Methods for increasing the potential of integration of ev chargers into the dc catenary of electric transport grids: A trolleygrid case study,” *eTransportation*, vol. 18, p. 100 271, 2023, ISSN: 2590-1168. DOI: `https://doi.org/10.1016/j.etran.2023.100271`. [Online]. Available: `https://www.sciencedirect.com/science/article/pii/S2590116823000462`.

[12] International Electrotechnical Commission, *IEC 60076-7:2018 Power Transformers – Part 7: Loading Guide for Mineral-Oil-Immersed Power Transformers*, International Standard, Geneva, Switzerland, 2018.

[13] CENELEC, “EN 50122-1: Railway applications — Fixed installations — Electrical safety, earthing and the return circuit — Part 1: Protective provisions against electric shock and stray current corrosion,” European Committee for Electrotechnical Standardization, Brussels, Standard, 2022.

[14] CENELEC, “EN 50163: Railway applications — Supply voltages of traction systems,” European Committee for Electrotechnical Standardization, Brussels, Standard, 2004, Permissible voltage bands; for 750 V DC: $U_{\mathrm{min1}} = 500$ V, $U_{\mathrm{max1}} = 900$ V (permanent), $U_{\mathrm{max2}} = 1000$ V (non-permanent). Equivalent in scope to IEC 60850.

[15] NEN, “NEN-EN 50388-1: Railway applications — Fixed installations and rolling stock — Technical criteria for the coordination between electric traction power supply systems and rolling stock to achieve interoperability — Part 1: General,” Nederlands Normalisatie-instituut, Delft, Standard, 2022, Specifies the minimum mean useful voltage at the pantograph.

[16] CENELEC, “EN 50123: Railway applications — Fixed installations — D.C. switchgear,” European Committee for Electrotechnical Standardization, Brussels, Standard, 2003.

[17] “IEC 60076-12: Power Transformers – Part 12: Loading Guide for Dry-Type Power Transformers,” International Electrotechnical Commission, Geneva, Switzerland, Standard IEC 60076-12, 2008.

[18] K. Smith, L. Hunter, S. Galloway, C. Booth, C. Kerr, and M. Kellett, “Integrated charging of evs using existing lvdc light rail infrastructure: A case study,” in *2019 IEEE Third International Conference on DC Microgrids (ICDCM)*, 2019, pp. 1–7. DOI: `10.1109/ICDCM45535.2019.9232726`.

[19] M. Lukianov, I. Verbitsky, E. R. Cadaval, and R. Strzelecki, “An overview of bidirectional ev chargers: Empowering traction grid-powered chargers,” in *Power Systems Research and Operation: Selected Problems III*, O. Kyrylenko, S. Denysiuk, R. Strzelecki, I. Blinov, I. Zaitsev, and A. Zaporozhets, Eds. Cham: Springer Nature Switzerland, 2024, pp. 191–230, ISBN: 978-3-031-44772-3. DOI: `10.1007/978-3-031-44772-3_9`. [Online]. Available: `https://doi.org/10.1007/978-3-031-44772-3_9`.

[20] H. Hu, Y. Liu, Y. Li, *et al.*, “Traction power systems for electrified railways: Evolution, state of the art, and future trends,” *Railway Engineering Science*, vol. 32, pp. 1–19, 2024. DOI: `10.1007/s40534-023-00320-6`.

[21] I. Diab, G. R. C. Mouli, and P. Bauer, “A traction substation state estimator for integrating smart loads in transportation grids without the need for additional sensors,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 25, no. 3, pp. 2669–2680, 2024. DOI: `10.1109/TITS.2023.3317375`.

[22] Y. Chen, M. Chen, L. Xu, and Z. Liang, “Chance-constrained optimization of storage and pfc capacity for railway electrical smart grids considering uncertain traction load,” *IEEE Transactions on Smart Grid*, vol. 15, no. 1, pp. 286–298, 2024. DOI: `10.1109/TSG.2023.3276198`.

[23] G. Pierrou and G. Hug, *Integrating optimal ev charging in the energy management of electric railway stations*, 2023. arXiv: `2308.01145 [eess.SY]`. [Online]. Available: `https://arxiv.org/abs/2308.01145`.

[24] Y. Chen, Y. Chen, Y. Zhou, W. Gan, M. Chen, and X. Chen, “Robust energy management of electrified railway traction substation providing secondary frequency regulation,” *Energy*, vol. 335, p. 138 231, 2025, ISSN: 0360-5442. DOI: `https://doi.org/10.1016/j.energy.2025.138231`. [Online]. Available: `https://www.sciencedirect.com/science/article/pii/S0360544225038733`.

[25] M. Bollen and F. Hassan, “Integration of distributed generation in the power system - a power quality approach,” Nov. 2008, pp. 1–8. DOI: `10.1109/ICHQP.2008.4668746`.

[26] N. Qammar, A. Arshad, R. J. Miller, K. Mahmoud, and M. Lehtonen, “Probabilistic hosting capacity assessment towards efficient pv-rich low-voltage distribution networks,” *Electric Power Systems Research*, vol. 226, p. 109 940, 2024, ISSN: 0378-7796. DOI: `https://doi.org/10.1016/j.epsr.2023.109940`. [Online]. Available: `https://www.sciencedirect.com/science/article/pii/S0378779623008283`.

[27] S. Munikoti, M. Abujubbeh, K. Jhala, and B. Natarajan, “A novel framework for hosting capacity analysis with spatio-temporal probabilistic voltage sensitivity analysis,” *International Journal of Electrical Power Energy Systems*, vol. 134, p. 107 426, 2022, ISSN: 0142-0615. DOI: `https://doi.org/10.1016/j.ijepes.2021.107426`. [On-

line]. Available: https://www.sciencedirect.com/science/article/pii/S0142061521006657.

[28] M. J. Chihota, B. Bekker, and T. Gaunt, "A stochastic analytic-probabilistic approach to distributed generation hosting capacity evaluation of active feeders," *International Journal of Electrical Power Energy Systems*, vol. 136, p. 107 598, 2022, ISSN: 0142-0615. DOI: https://doi.org/10.1016/j.ijepes.2021.107598. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0142061521008322.

[29] H. Liu, J. He, L. Deng, J. Xu, and Y. Chen, "Assessment of distribution network hosting capacity incorporating multiple uncertainty factors," *Electronics*, vol. 14, no. 22, 2025, ISSN: 2079-9292. DOI: 10.3390/electronics14224495. [Online]. Available: https://www.mdpi.com/2079-9292/14/22/4495.

[30] Q. Gao, H. Cheng, S. Zheng, S. Wan, and W. Li, "A method for predicting traction load of electrified railways considering spatiotemporal correlation characteristics," *IET Electrical Systems in Transportation*, vol. 2025, no. 1, p. 5 516 562, 2025. DOI: https://doi-org.tudelft.idm.oclc.org/10.1049/els2/5516562. eprint: https://ietresearch-onlinelibrary-wiley-com.tudelft.idm.oclc.org/doi/pdf/10.1049/els2/5516562. [Online]. Available: https://ietresearch-onlinelibrary-wiley-com.tudelft.idm.oclc.org/doi/abs/10.1049/els2/5516562.

[31] Z. Li, X. Li, Y. Wei, C. Lu, and X. Bai, "Fast quasi-optimal power flow of flexible dc traction power systems," *IEEE Transactions on Power Systems*, vol. 39, no. 1, pp. 1555–1567, 2024. DOI: 10.1109/TPWRS.2023.3255236.

[32] A. Babiker, S. S. Ahmad, I. Ahmed, M. Khalid, M. A. Abido, and F. Saleh Al-Ismail, "Optimal power flow: A review of state-of-the-art techniques and future perspectives," *IEEE Access*, vol. 13, pp. 60 012–60 039, 2025. DOI: 10.1109/ACCESS.2025.3556168.

[33] K. Suslov, A. Kryukov, P. Ilyushin, A. Kryukov, and O. Shepovalova, "Modeling of modes of traction power supply systems equipped with renewable energy sources," *Energy Reports*, vol. 9, pp. 447–461, 2023, Technologies and Materials for Renewable Energy, Environment and Sustainability, ISSN: 2352-4847. DOI: https://doi.org/10.1016/j.egyr.2023.06.034. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2352484723010545.

[34] P. Martínez Fernández, I. Villalba Sanchís, V. Yepes, and R. Insa Franco, "A review of modelling and optimisation methods applied to railways energy consumption," *Journal of Cleaner Production*, vol. 222, pp. 153–162, 2019, ISSN: 0959-6526. DOI: https://doi.org/10.1016/j.jclepro.2019.03.037. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0959652619307267.

[35] K. Deb and H. Jain, "An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part I: Solving problems with box constraints," *IEEE Transactions on Evolutionary Computation*, vol. 18, no. 4, pp. 577–601, 2014. DOI: 10.1109/TEVC.2013.2281535.

[36] A. Nebro, J. Durillo, J. Garcia-Nieto, C. Coello Coello, F. Luna, and E. Alba, “Smpso: A new pso-based metaheuristic for multi-objective optimization,” in *2009 IEEE Symposium on Computational Intelligence in Multi-Criteria Decision-Making(MCDM)*, 2009, pp. 66–73. DOI: `10.1109/MCDM.2009.4938830`.

[37] I. Das and J. E. Dennis, “Normal-boundary intersection: A new method for generating the pareto surface in nonlinear multicriteria optimization problems,” *SIAM Journal on Optimization*, vol. 8, no. 3, pp. 631–657, 1998. DOI: `10.1137/S1052623496307510`.

[38] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, “A fast and elitist multiobjective genetic algorithm: NSGA-II,” *IEEE Transactions on Evolutionary Computation*, vol. 6, no. 2, pp. 182–197, 2002. DOI: `10.1109/4235.996017`.

[39] K. Deb, “An efficient constraint handling method for genetic algorithms,” *Computer Methods in Applied Mechanics and Engineering*, vol. 186, no. 2–4, pp. 311–338, 2000. DOI: `10.1016/S0045-7825(99)00389-8`.

[40] E. Zitzler, M. Laumanns, and S. Bleuler, “A tutorial on evolutionary multiobjective optimization,” in *Metaheuristics for Multiobjective Optimisation*, X. Gandibleux, M. Sevaux, K. Sörensen, and V. T’kindt, Eds., Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 3–37, ISBN: 978-3-642-17144-4.

[41] A. Charnes and W. W. Cooper, “Chance-constrained programming,” *Management Science*, vol. 6, no. 1, pp. 73–79, 1959. DOI: `10.1287/mnsc.6.1.73`.

[42] R. T. Rockafellar and S. Uryasev, “Optimization of conditional value-at-risk,” *Journal of Risk*, vol. 2, no. 3, pp. 21–41, 2000. DOI: `10.21314/JOR.2000.038`.

[43] P. Ju, H. Li, X. Pan, C. Gan, Y. Liu, and Y. Liu, “Stochastic dynamic analysis for power systems under uncertain variability,” *IEEE Transactions on Power Systems*, vol. 33, no. 4, pp. 3789–3799, 2018. DOI: `10.1109/TPWRS.2017.2777783`.

[44] N. Rehman, M.-d. Mufti, and N. Gupta, “Power flow analysis in a distribution system penetrated with renewable energy sources: A review,” *International Journal of Ambient Energy*, vol. 45, no. 1, p. 2 305 701, 2024. DOI: `10.1080/01430750.2024.2305701`. eprint: `https://doi.org/10.1080/01430750.2024.2305701`. [Online]. Available: `https://doi.org/10.1080/01430750.2024.2305701`.

[45] Z. Kaseb, S. Orfanoudakis, P. P. Vergara, and P. Palensky, “Physics-informed neural network with adaptive activation for power flow,” *International Journal of Electrical Power Energy Systems*, vol. 174, p. 111 525, 2026, ISSN: 0142-0615. DOI: `https://doi.org/10.1016/j.ijepes.2025.111525`. [Online]. Available: `https://www.sciencedirect.com/science/article/pii/S0142061525010737`.

[46] Y. Yuan, S. H. Low, O. Ardakanian, and C. J. Tomlin, “Inverse power flow problem,” *IEEE Transactions on Control of Network Systems*, vol. 10, no. 1, pp. 261–273, 2023. DOI: `10.1109/TCNS.2022.3199084`.

[47] Y. Cai, M. Irving, and S. Case, “Iterative techniques for the solution of complex dc-rail-traction systems including regenerative braking,” *IEE Proceedings - Generation, Transmission and Distribution*, vol. 142, pp. 445–452, 5 1995. DOI: `10.1049/ip-gtd:19952098`. eprint: `https://digital-library.theiet.org/doi/pdf/10.1049/ip-gtd%3A19952098`. [Online]. Available: `https://digital-library.theiet.org/doi/abs/10.1049/ip-gtd%3A19952098`.

[48] H. Heitsch and W. Römisch, “Scenario reduction algorithms in stochastic programming,” *Computational optimization and applications*, vol. 24, no. 2, pp. 187–206, 2003.

[49] Y. Zhang, Y. Tang, J. Hu, *et al.*, “High-dimensional probability preserving scenario reduction method using adaptive selection with sampling for large-scale power system,” *Applied Energy*, vol. 402, p. 126 898, 2026.

[50] Y. Zhuang, L. Cheng, N. Qi, M. R. Almassalkhi, and F. Liu, “Problem-driven scenario reduction framework for power system stochastic operation,” *IEEE Transactions on Power Systems*, vol. 40, no. 4, pp. 3232–3246, 2025. DOI: `10.1109/TPWRS.2024.3523220`.

[51] A. Gonzalez-Castellanos and D. Pozo, “An optimal scenario reduction method for stochastic power system problems,” in *2021 IEEE Madrid PowerTech*, IEEE, 2021, pp. 1–6.

[52] K. Bruninx and E. Delarue, “Scenario reduction techniques and solution stability for stochastic unit commitment problems,” in *2016 IEEE International Energy Conference (ENERGYCON)*, 2016, pp. 1–7. DOI: `10.1109/ENERGYCON.2016.7514074`.

[53] J. Dupacová, N. Gröwe-Kuska, and W. Römisch, “Scenario reduction in stochastic programming,” *Mathematical programming*, vol. 95, no. 3, p. 493, 2003.

[54] H. Teichgraeber and A. R. Brandt, “Clustering methods to find representative periods for the optimization of energy systems: An initial framework and comparison,” *Applied energy*, vol. 239, pp. 1283–1293, 2019.

[55] T. Schütz, M. H. Schraven, M. Fuchs, P. Remmen, and D. Müller, “Comparison of clustering algorithms for the selection of typical demand days for energy system synthesis,” *Renewable energy*, vol. 129, pp. 570–582, 2018.

[56] I. Blanco and J. M. Morales, “An efficient robust solution to the two-stage stochastic unit commitment problem,” *IEEE Transactions on Power Systems*, vol. 32, no. 6, pp. 4477–4488, 2017.

[57] I. Chopra and D. Selvamuthu, “Scenario generation in stochastic programming using principal component analysis based on moment-matching approach,” *Opsearch*, vol. 57, no. 1, pp. 190–201, 2020.

[58] K. Høyland, M. Kaut, and S. W. Wallace, “A heuristic for moment-matching scenario generation,” *Computational optimization and applications*, vol. 24, no. 2, pp. 169–185, 2003.

[59] M. Kaut and S. W. Wallace, *Evaluation of scenario-generation methods for stochastic programming*, J. L. Higle, W. RÃ¶misch, and S. Sen, Eds. Humboldt-UniversitÃ¤t zu Berlin, Mathematisch-Naturwissenschaftliche FakultÃ¤t II, Institut fÃ¼r Mathematik, 2003. DOI: `http://dx.doi.org/10.18452/8296`.

[60] J. Keutchayan, J. Ortmann, and W. Rei, “Problem-driven scenario clustering in stochastic optimization,” *arXiv preprint arXiv:2106.11717*, 2021.

[61] J. Faraji, H. Hashemi-Dezaki, and A. Ketabi, “Optimal probabilistic scenario-based operation and scheduling of prosumer microgrids considering uncertainties of renewable energy sources,” *Energy Science & Engineering*, vol. 8, no. 11, pp. 3942–3960, 2020.

[62] P. Cheng, H. Kong, J. Ma, and L. Jia, “Overview of resilient traction power supply systems in railways with interconnected microgrid,” *CSEE Journal of Power and Energy Systems*, vol. 7, no. 5, pp. 1122–1132, 2021. DOI: 10.17775/CSEEJPES.2020.02110.

[63] M. B. Wafaa and L.-A. Dessaint, “Multi-objective stochastic optimal power flow considering voltage stability and demand response with significant wind penetration,” *IET Generation, Transmission & Distribution*, vol. 11, no. 14, pp. 3499–3509, 2017. DOI: https://doi.org/10.1049/iet-gtd.2016.1994. eprint: https://ietresearch.onlinelibrary.wiley.com/doi/pdf/10.1049/iet-gtd.2016.1994. [Online]. Available: https://ietresearch.onlinelibrary.wiley.com/doi/abs/10.1049/iet-gtd.2016.1994.

[64] I. K. Iliev, A. V. Kryukov, K. V. Suslov, *et al.*, “Modeling the operating conditions of electric power systems feeding dc and ac traction substations,” *Energies*, vol. 17, no. 18, 2024, ISSN: 1996-1073. DOI: 10.3390/en17184692. [Online]. Available: https://www.mdpi.com/1996-1073/17/18/4692.

[65] C. Fiori, K. Ahn, and H. A. Rakha, “Power-based electric vehicle energy consumption model: Model development and validation,” *Applied Energy*, vol. 168, pp. 257–268, 2016, ISSN: 0306-2619. DOI: https://doi.org/10.1016/j.apenergy.2016.01.097. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S030626191630085X.

[66] D. Morea, S. Elia, C. Boccaletti, and P. Buonadonna, “Improvement of energy savings in electric railways using coasting technique,” *Energies*, vol. 14, no. 23, 2021, ISSN: 1996-1073. DOI: 10.3390/en14238120. [Online]. Available: https://www.mdpi.com/1996-1073/14/23/8120.

[67] Q. Zou, H. Li, and R. Zhang, “Inverse reinforcement learning via neural network in driver behavior modeling,” in *2018 IEEE Intelligent Vehicles Symposium (IV)*, 2018, pp. 1245–1250. DOI: 10.1109/IVS.2018.8500666.

[68] R. Ghandour, A. J. Potams, I. Boulkaibet, B. Neji, Z. A. Barakeh, and A. S. Karar, “Machine learning methods for driver behaviour classification,” in *2021 4th International Conference on Bio-Engineering for Smart Technologies (BioSMART)*, 2021, pp. 1–4. DOI: 10.1109/BioSMART54244.2021.9677801.

[69] P. Lukaszewicz, “Driving techniques and strategies for freight trains,” in *Computers in Railways VII*, C. A. Brebbia, J. Allan, R. J. Hill, G. Sciutto, and S. Sone, Eds., ser. WIT Transactions on the Built Environment, vol. 50, Southampton, UK: WIT Press, 2000, pp. 1065–1073, ISBN: 1-85312-826-0. DOI: 10.2495/CR001041. [Online]. Available:

https://www.witpress.com/Secure/elibrary/papers/CR00/CR00104FU.pdf.

[70] L. Thurner, A. Scheidler, F. Schafer, *et al.*, “Pandapower - an open source python tool for convenient modeling, analysis and optimization of electric power systems,” *IEEE Transactions on Power Systems*, 2018, ISSN: 0885-8950. DOI: 10.1109/TPWRS.2018.2829021. [Online]. Available: https://arxiv.org/abs/1709.06743.

[71] I. Diab, A. Saffirio, G. R. C. Mouli, A. S. Tomar, and P. Bauer, “A complete dc trolleybus grid model with bilateral connections, feeder cables, and bus auxiliaries,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 10, pp. 19 030–19 041, 2022. DOI: 10.1109/TITS.2022.3157080.

[72] C. A. Coello Coello, G. T. Pulido, and M. S. Lechuga, “Handling multiple objectives with particle swarm optimization,” *IEEE Transactions on Evolutionary Computation*, vol. 8, no. 3, pp. 256–279, 2004. DOI: 10.1109/TEVC.2004.826067.

A

RET Calandlijn technical specifications

A.1. Rolling stock

Table A.1: Rolling stock parameters for the Calandlijn.

Parameter	**Value**	**Notes**
Vehicle type	SG3	RandstadRail consist
Tare mass (per unit)	$65\,\mathrm{t}$	Empty vehicle
Mass at passenger loading	$100\,\mathrm{t}$	Standard service loading
Rated traction power (per carriage)	$1.04\,\mathrm{MW}$	Continuous
Rotating mass	$6680\,\mathrm{kg}$	Addition to effective weight under acceleration
Maximum braking power (per carriage)	$1.04\,\mathrm{MW}$	Limited by chopper rating
Regenerative braking efficiency, η_{regen}	0.80	Below $5\,\mathrm{m\,s^{-1}}$: mechanical only
Auxiliary load, P_{aux}	$50\,\mathrm{kW}$	HVAC, lighting, on-board systems
Davis coefficient A	$1000\,\mathrm{N}$	Constant resistance
Davis coefficient B	$30\,\mathrm{N\,s\,m^{-1}}$	Bearing and flange
Davis coefficient C	$2\,\mathrm{N\,s^{2}\,m^{-2}}$	Aerodynamic
Maximum operating speed	$100\,\mathrm{km\,h^{-1}}$	Line maximum varies by section
Maximum acceleration	$1.2\,\mathrm{s^{-2}}$	-
Maximum deceleration	$1.2\,\mathrm{s^{-2}}$	-
Maximum traction effort	$136\,\mathrm{kN}$	Max direct input by driver

A.2. Traction grid specifications

Table A.2: Calandlijn traction substations.

ID	Name (location)	Rating	Cluster (utility)	Function
HVS	Havenstraat	2.88 MV A	MSMCP	Line-end terminus
RCS	Rochussenstraat	1.44 MV A	MSMCP	Intermediate feeding
BLK	Blaak	1.44 MV A	MSMCP	Intermediate feeding
GDW	Gerdesiaweg	1.44 MV A	GSGDW	Cluster head
KLZ	Kralingse Zoom	1.44 MV A	GSGDW	Intermediate feeding
CPB	Capelsebrug	2.50 MV A	GSNSL	Junction feeding
TRP	Terp	1.44 MV A	GSTRP	Cluster head
PSL	Prinsenlaan	2.50 MV A	GSNSL	Intermediate feeding
ALD	Alexander	2.88 MV A	GSNSL	Intermediate feeding
BNH	Binnenhof	2.50 MV A	GSNSL	Junction feeding
NVL	Nieuw Verlaat	2.88 MV A	GSNSL	Intermediate feeding
NSL	Nesselande	2.88 MV A	GSNSL	Cluster head / line-end

Table A.3: Conductor parameters of the Calandlijn 750 V DC traction network used in the power-flow model. Each feeder run enters the conductance matrix as a lumped branch whose per-metre resistance is the sum of the positive-conductor and return-rail values listed below; substation-to-feeder cables enter as separate branches between the substation busbar and the feeder-endpoint node.

Conductor / element	ρ [mΩ/m]	Application	Notes
Third rail (positive)	0.007	HVS–CPB main trunk; CPB–TRP branch; TTN–NSL terminus	composite steel/aluminium; continuous along route
Overhead catenary (BVL)	0.060	CPB–PSL–ALD–NVL with BNH branch; NVL–TTN (eastern section)	bilaterally fed; EN 50119, EN 50122
Running rail (return)	0.040	all feeder runs	two 54E1 rails in parallel; bonded $\leq$ 300 m per EN 50122-1
Substation cable, std†	0.015	busbar–to–feeder cables at HVS, RCS, BLK, GDW, KLZ, CPB, TRP	Cu; runs typically 50–165 m
Substation cable, heavy	0.008	busbar–to–feeder cables at PSL, ALD, BNH, NVL, NSL	Cu, larger cross-section; runs up to 445 m

A.3. Normative references

A.2 Reference time period over which values can be averaged or integrated

Table A.1 — Integration periods

Function	Duration	Comments
Protection	< 100 ms up to 300 ms for a.c. < 30 ms up to 100 ms for d.c.	Clearance of short-circuit faults depending on the intensity of the short-circuit current.
Back up protection	200 ms to 3 s	Clearance of fault on back-up protection.
Peak load current	1 s to 5 s	At supply points and for individual overhead contact lines. Applicable to computer simulations, effect on feeding arrangements and assessment of back-up protection.
	1 min	r.m.s. average current for calculation of thermal loading on equipment.
Thermal overloading	30 s to 20 min preferred value: 10 min	r.m.s. average current for calculation of thermal loading on equipment.
Energy metering period	10 min, 15 min or 30 min	Average current or power for definition of the demand on the supply authority.
Long-term loading	1 h, 2 h	Conventional parameter to define long-term railway load.
Energy consumption	1 day, 1 week, 1 month, 1 year	Average energy consumption, used by electricity suppliers to study (for example) railway power requirements or when purchasing electricity.

Figure A.1: NEN EN 50388 timescale table for timestep integration

Timescales for integration periods of values within TPSS (NEN-EN 50388)

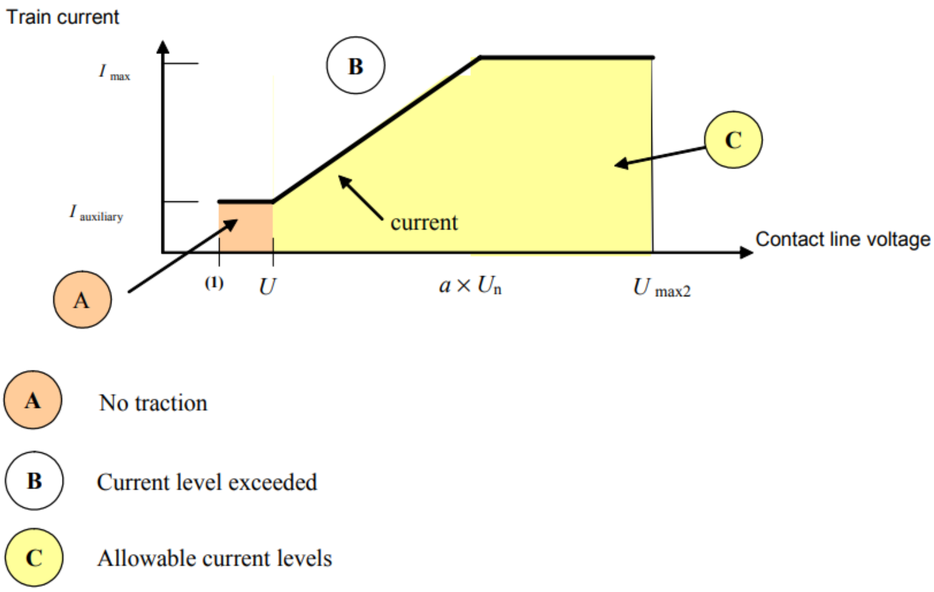

Key

U contact line voltage according to EN 50163

I_{max} is the maximum current consumed by the train at nominal voltage.

(1) With regard to the setting values of the under-voltage releases, see EN 50163:2004, 4.1, Note 2.

Figure 1 — Maximum train current against voltage

Table 2 — Value of factor a

Power supply system	Value of a
a.c. 25 000 V 50 Hz	0,9
a.c. 15 000 V 16,7 Hz	0,95
d.c. 3 000 V	0,9
d.c. 1 500 V	0,9
d.c. 750 V	0,8

Figure A.2: Allowable train current with corresponding voltage scaling factor, EN 50388

Nominal alllowable train current at permissible voltage levels

Nominal voltage range and permissible limits in values and duration, NEN-EN 50163

4 Voltages and frequencies of traction systems

4.1 Voltages

The characteristics of the main voltage systems (overvoltages excluded) are specified in Table 1 below.

Table 1 – Nominal voltages and their permissible limits in values and duration

Electrification system	Lowest non-permanent voltage U_{min2} V	Lowest permanent voltage U_{min1} V	Nominal voltage U_n V	Highest permanent voltage U_{max1} V	Highest non-permanent voltage U_{max2} V
d.c. (mean values)	400	400	600 [a]	720	800
	500 [c]	500	750	900 [c]	1 000
	1 000	1 000	1 500	1 800 [c]	1 950
	2 000	2 000	3 000	3 600	3 900 [b]
a.c. (r.m.s. values)	11 000	12 000	15 000	17 250	18 000
	17 500 [c]	19 000 [c]	25 000	27 500 [c]	29 000
Special national conditions for France, see Annex B.					

[a] Future d.c. traction systems for tramways and local railways should conform with system nominal voltage of 750 V, 1 500 V or 3 000 V.

[b] Special national conditions for Belgium, see Annex B.

[c] Special national conditions for United Kingdom, see Annex B.

The following requirements shall be fulfilled:

a) the duration of voltages between U_{min1} and U_{min2} shall not exceed 2 min;

b) the duration of voltages between U_{max1} and U_{max2} shall not exceed 5 min;

c) the voltage of the busbar at the substation at no load condition shall be less than or equal to U_{max1}. For d.c. substations it is acceptable to have this voltage at no load condition less than or equal to U_{max2}, knowing that when a train is present, the voltage at this train's pantograph (s) shall be in accordance with Table 1 and its requirements;

d) under normal operating conditions, voltages shall lie within the range $U_{min1} \leq U \leq U_{max2}$;

e) under abnormal operating conditions the voltages in the range $U_{min2} \leq U \leq U_{min1}$ in Table 1 shall not cause any damages or failures.

NOTE 1 The use of train power limitation devices on board may limit the presence of low voltage on the overhead line (see EN 50388).

f) if voltages between U_{max1} and U_{max2} are reached, it shall be followed by a level below or equal to U_{max1}, for an unspecified period.

Voltages between U_{max1} and U_{max2} shall only be reached for non permanent conditions such as

– regenerative braking,

– move of voltage regulation systems such as mechanical tap changer.

g) lowest operational voltage: under abnormal operating conditions U_{min2} is the lowest limit of the contact line voltage for which the rolling stock is intended to operate.

NOTE 2 Recommended values for undervoltage tripping: The setting of undervoltage relays in fixed installations or on board rolling stock should be set from 85 % to 95 % of U_{min2}.

Figure A.3: Permissible duration dependent voltage limits of traction feeders, per NEN EN 50163

A.4. Constraint violation results: Scenario reduction, optimiser choice, confidence level

This section provides figures on constraint violations post validation, as well as a table at the end comparing substation overloading prevalence

A.4.1. Scenario reduction

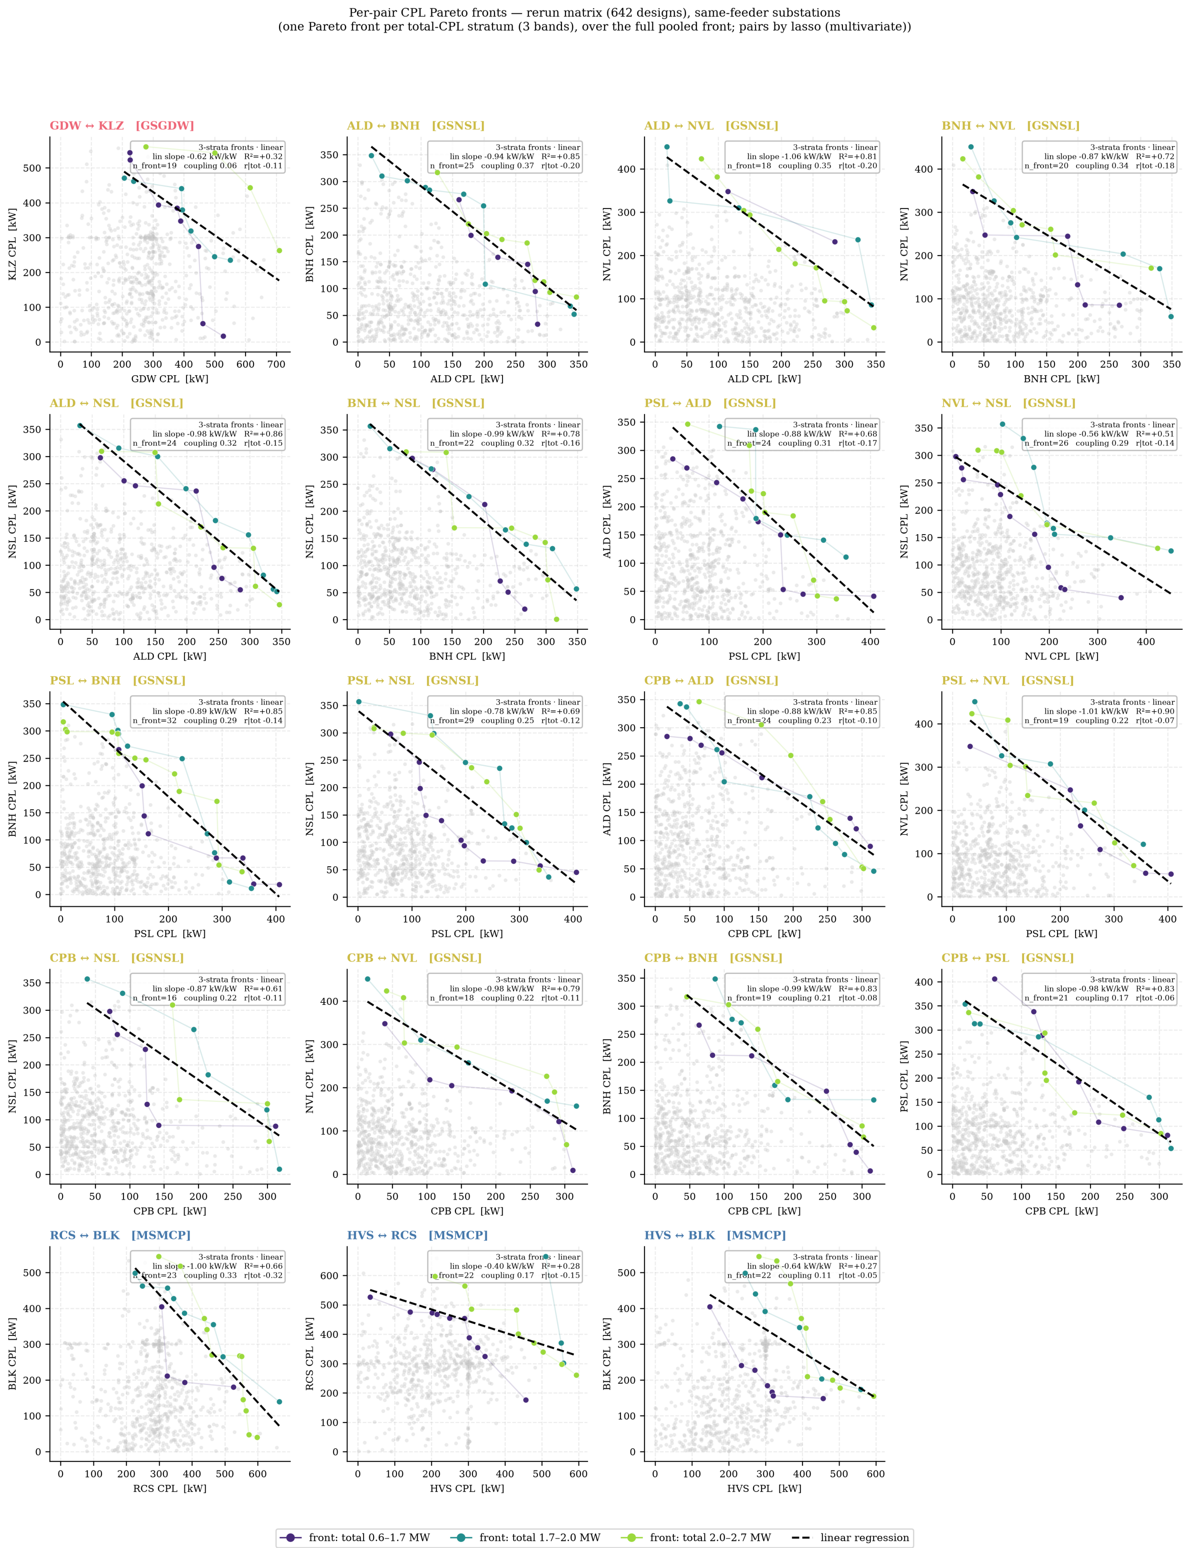

Figure A.4: Trade-off power allocation between substations connected to the same utility feeder. This figure illustrates the spatial distribution of load allocation at the Pareto front between 2 substations. The regression line represents the upper bound of the combined power allocation of these substations, while the pareto fronts represent the different strata (based on total CPL) under which the system may exist.

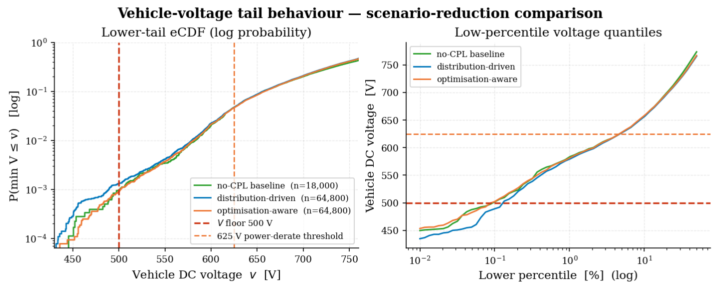

Figure A.5: Vehicle-voltage lower-tail eCDF and low-percentile quantiles, by scenario-reduction method.

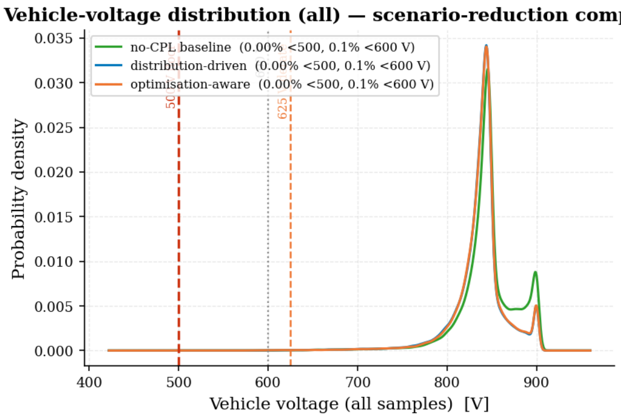

Figure A.6: Vehicle-voltage distribution (Gaussian KDE, all samples), by scenario-reduction method.

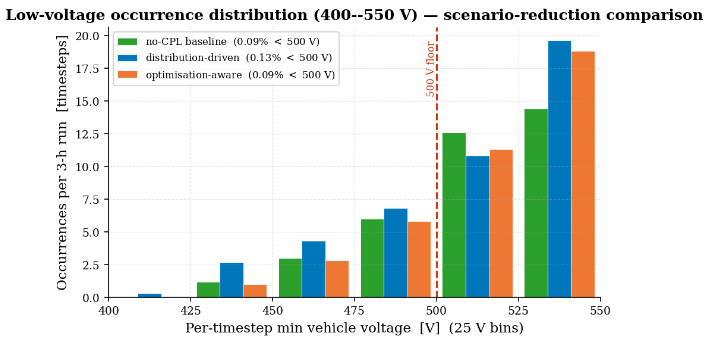

Figure A.7: Low-voltage occurrence distribution (per-timestep minimum, 400–550 V), by scenario-reduction method.

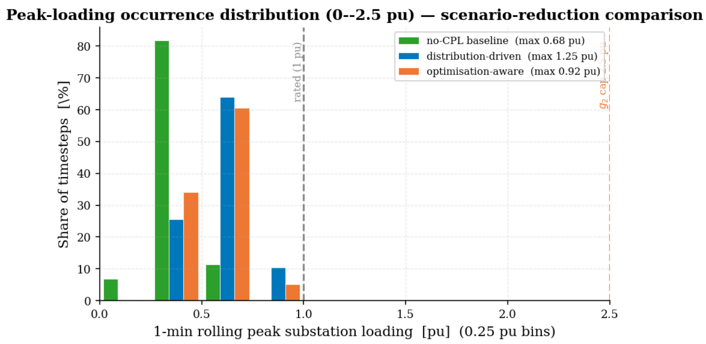

Figure A.8: 1-min rolling peak substation loading distribution (the enforced g_2 quantity), by scenario-reduction method.

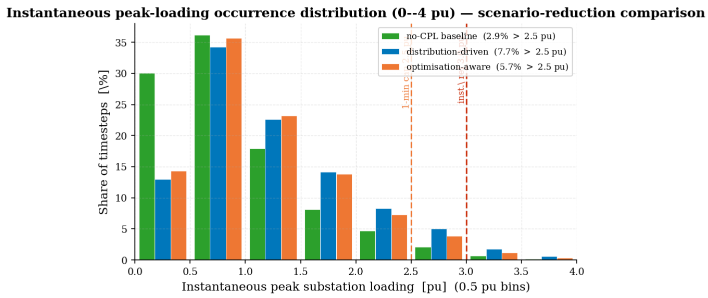

Figure A.9: Instantaneous peak substation loading distribution (0–4 pu), by scenario-reduction method.

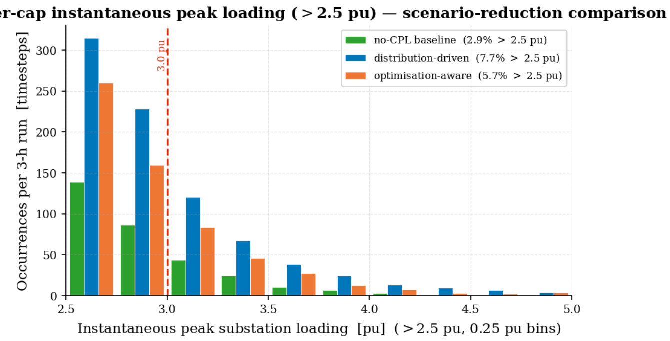

Figure A.10: Over-cap instantaneous peak loading (> 2.5 pu transient excursions), by scenario-reduction method.

A.4.2. Optimiser choice

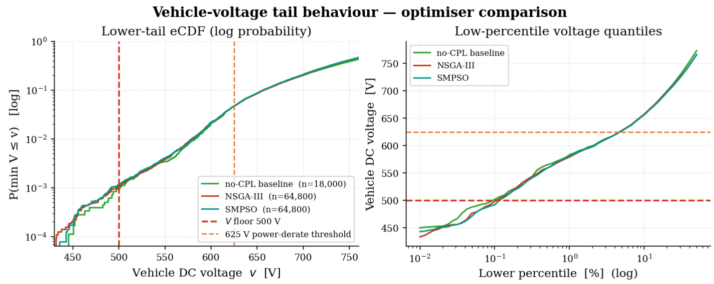

Figure A.11: Vehicle-voltage lower-tail eCDF and low-percentile quantiles, by optimiser (NSGA-III vs SMPSO).

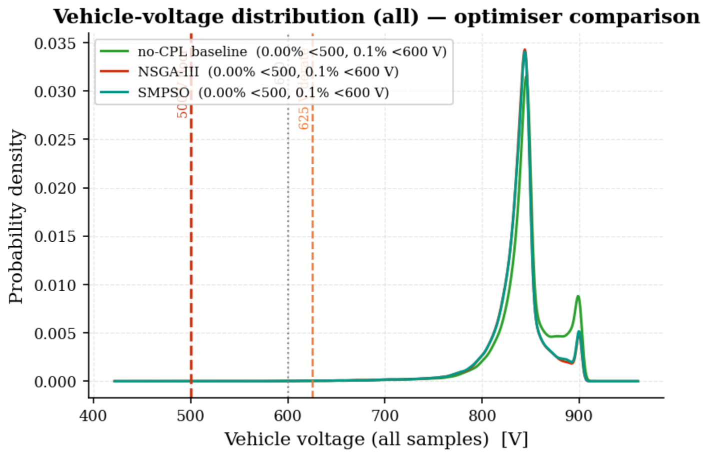

Figure A.12: Vehicle-voltage distribution (Gaussian KDE, all samples), by optimiser.

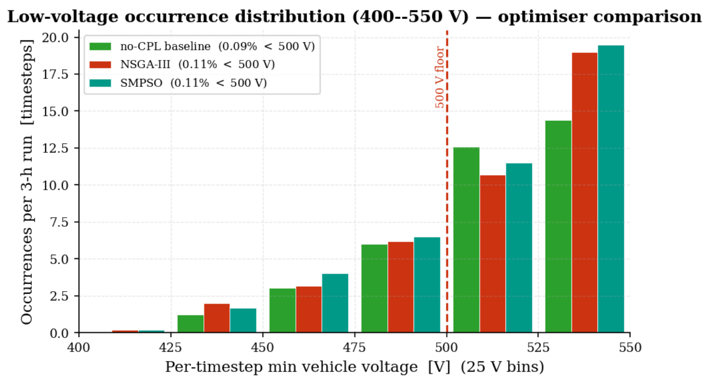

Figure A.13: Low-voltage occurrence distribution (per-timestep minimum, 400–550 V), by optimiser.

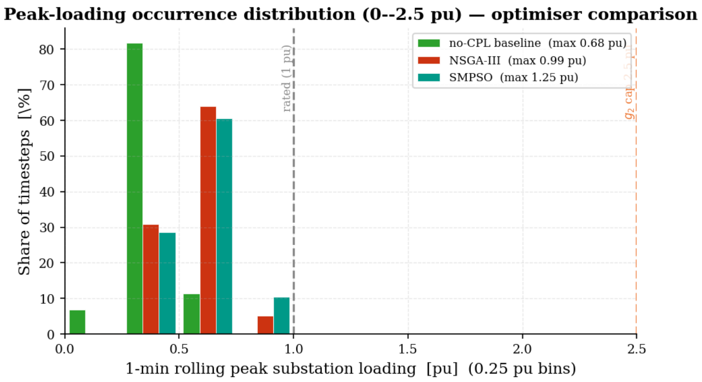

Figure A.14: 1-min rolling peak substation loading distribution (the enforced g_2 quantity), by optimiser.

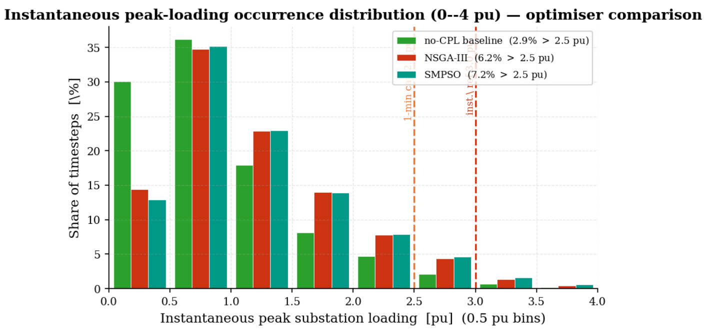

Figure A.15: Instantaneous peak substation loading distribution (0–4 pu), by optimiser.

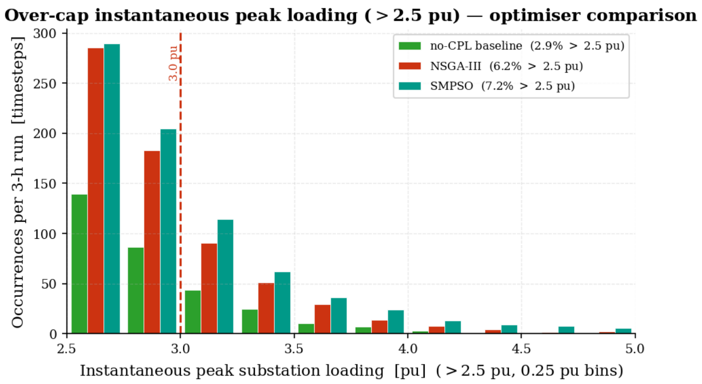

Figure A.16: Over-cap instantaneous peak loading (> 2.5 pu transient excursions), by optimiser.

A.4.3. Confidence level

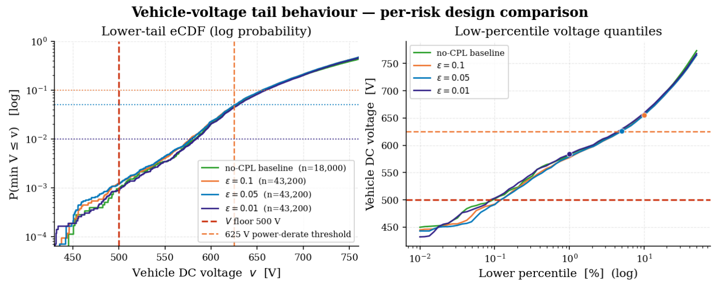

Figure A.17: Vehicle-voltage lower-tail eCDF and low-percentile quantiles, by confidence level ε (dotted lines mark each design's permitted allowance).

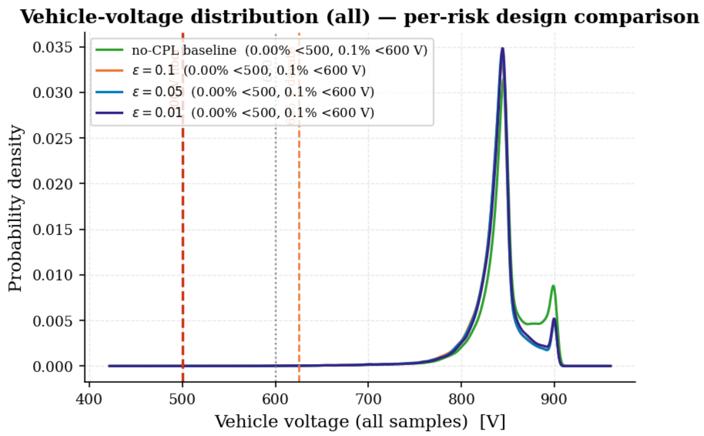

Figure A.18: Vehicle-voltage distribution (Gaussian KDE, all samples), by confidence level ε.

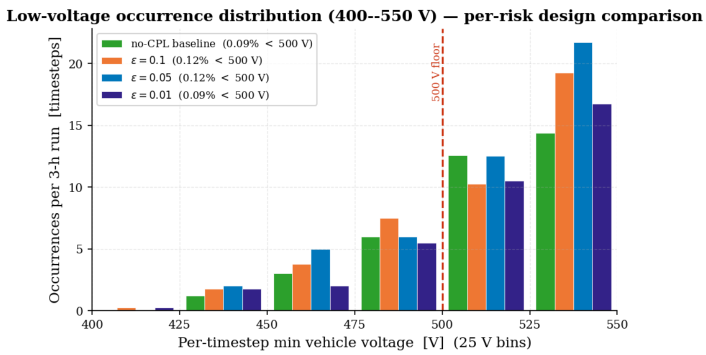

Figure A.19: Low-voltage occurrence distribution (per-timestep minimum, 400–550 V), by confidence level ε.

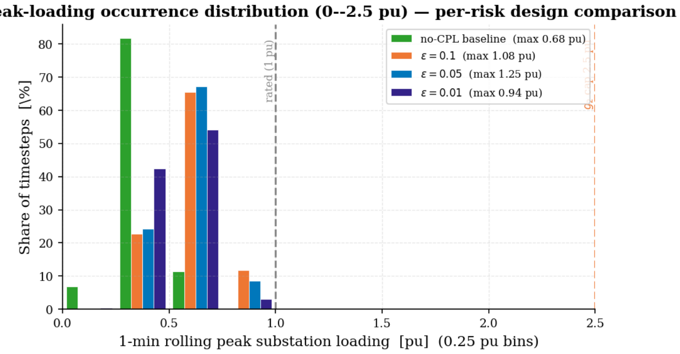

Figure A.20: 1-min rolling peak substation loading distribution (the enforced g_2 quantity), by confidence level ε.

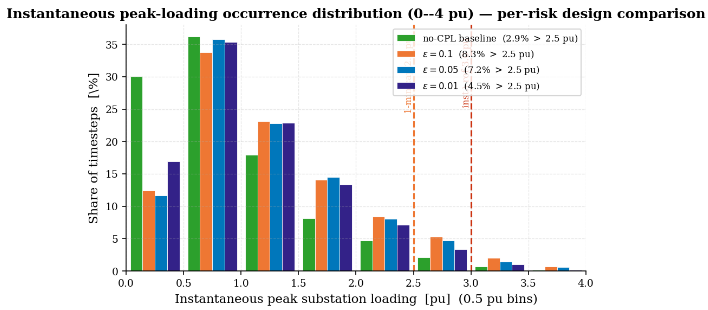

Figure A.21: Instantaneous peak substation loading distribution (0–4 pu), by confidence level ε.

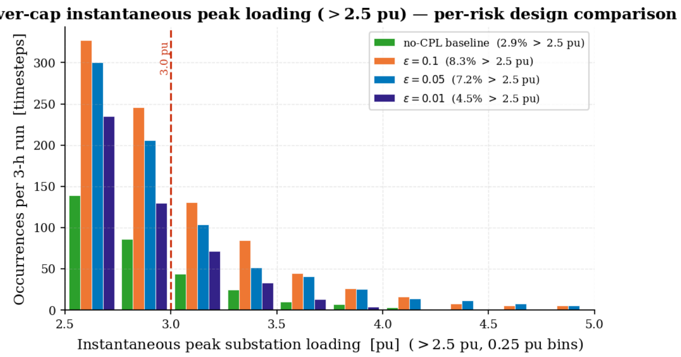

Figure A.22: Over-cap instantaneous peak loading (> 2.5 pu transient excursions), by confidence level ε.

Table A.4: Per-substation instantaneous over-loading. For each substation: the fraction of timesteps at which its instantaneous loading ratio exceeds 2.5 pu, and its share of the system-wide instantaneous peak on the timesteps where that peak is above 2.5 pu. Pooled over the twelve validated 3-h co-simulations ($129\,600$ timesteps). Every over-cap excursion occurs at one of the four 1440 kW substations; the larger 2500–2880 kW substations never exceed 2.5 pu. These are sub-second transients on the *instantaneous* peak, not breaches of the enforced 1-min rolling constraint g_2.

Substation	Feeder	Rated [kW]	> 2.5 pu [% time]	Share of peaks [%]
GDW	GSGDW	1440	2.37	32.0
BLK	MSMCP	1440	2.30	34.2
KLZ	GSGDW	1440	1.27	18.6
RCS	MSMCP	1440	1.10	15.2
TRP	GSTRP	1440	0.00	0.0
HVS	MSMCP	2880	0.00	0.0
CPB	GSNSL	2500	0.00	0.0
PSL	GSNSL	2500	0.00	0.0
ALD	GSNSL	2880	0.00	0.0
BNH	GSNSL	2500	0.00	0.0
NVL	GSNSL	2880	0.00	0.0
NSL	GSNSL	2880	0.00	0.0