Power Hardware-in-The-Loop Setup for Phase Balancing Using Coordinated EV Charging

Master Thesis

Randa Ramana



TU Delft

Power Hardware-in-The-Loop Setup for Phase Balancing Using Coordinated EV Charging

Master Thesis

by

Randa Ramana

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Tuesday August 18, 2026 at 09:30 AM.

Student number:	6147186	
Project duration:	November 14, 2025 – August 18, 2026	
Thesis committee:	Dr. Ir. G. R. Chandra Mouli,	TU Delft, supervisor
	Dr. Ir. M. Ghaffarian Niasar,	TU Delft
	B. I. Priambodo, MSc.	TU Delft, daily supervisor

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

TU Delft

Preface

This thesis concludes my Master of Science in Sustainable Energy Technology at Delft University of Technology. I am grateful to those who supported me during the programme and during the work presented here.

I would like to thank my supervisors, Dr. ir. G. R. Chandra Mouli and Bagas Priambodo, M.Sc., for the opportunity to carry out this thesis within the DC Systems, Energy Conversion & Storage (DCE&S) group, and for their guidance throughout. I also thank Dr. ir. M. Ghaffarian Niasar for taking part in the thesis committee.

The anonymised low-voltage grid topology data used in this work was provided by Stedin Netbeheer. I thank Bart Kers and Tim Slangen for making it available, which allowed the case study to be based on a real Dutch distribution feeder.

I am grateful to the staff of the ESP laboratory for their supervision and assistance during the hardware experiments, which were essential to the experimental part of this thesis.

I am thankful to LPDP and PT.PLN (Persero) who funded my study in TU Delft which enable me to pursue this degree in the first place.

I thank my wife Shofi and my daughter Karissa for their support and patience over the past two years, without which I would not have completed this programme.

Finally, I thank my fellow students in Sustainable Energy Technology for their company throughout the programme.

Randa Ramana
Delft, August 2026

Summary

The uptake of rooftop photovoltaics and electric vehicles in European low-voltage networks is structurally single-phase, and their combined growth drives voltage unbalance beyond the 2 % limit of EN 50160. Coordinated phase switching of EV chargers has been proposed as a software-defined alternative to hardware-based compensation, but the existing literature treats phase balance, charging cost, and photovoltaic-driven flicker as separate problems, and validates the resulting algorithms exclusively in simulation.

This thesis develops a coordinated charging framework for a public charging facility equipped with phase-switching chargers, local PV, and a shared installation fuse, and validates it on physical charger hardware. The framework is decomposed across two timescales: a 15-minute mixed-integer linear program that assigns phases and base charging currents while weighing charging cost against a linearized load-imbalance term, and a 20-second linear program that absorbs PV fluctuations through a common-mode current shift constrained to be phase-symmetric, so that flicker mitigation cannot undo the balance decision. A three-phase forward-backward sweep load flow evaluates the resulting voltage unbalance on a real anonymized Dutch distribution feeder.

Across four scenarios, phase balancing reduces average load imbalance from 5.27 A under uncontrolled charging to 1.23 A, and the instantaneous voltage unbalance factor at the facility bus from a maximum of 2.41 % under cost-only optimization to 0.187 %, with zero unmet charging demand. Cost-only smart charging is shown to worsen worst-case unbalance by 80 % relative to no control, because it synchronizes the fleet. Once phase-switching freedom is held constant between baselines, the cost premium of balancing is small and centered near zero, ranging from $+7.5$ % to -4.5 %. The fast loop reduces visible flicker occurrence from 0.847 % to 0.523 % and remains effective up to roughly 60 kWp of installed PV, beyond which the connected fleet lacks the authority to absorb the fluctuation.

The schedules are executed on a Power Hardware-in-the-Loop testbed comprising a real-time simulator, a power amplifier, and a commercial phase-switching charger. Every current and phase command is executed as scheduled, establishing executability, while an open measurement path and a 7.3 % energy discrepancy attributable to unmodelled switching transients and conversion losses bound the claim short of closed-loop grid-service validation.

Contents

Nomenclature **xiii**

1 Introduction **1**
1.1 Context: Energy Transition and Low Voltage (LV) Grid Stress . . . 1
1.2 Problem: Voltage Unbalance and Its Consequences . . . 2
1.3 Research Gap and Motivation . . . 3
1.4 Research Objective and Questions . . . 4
1.5 Thesis Outline . . . 5

2 Literature Review **7**
2.1 Voltage Unbalance Compensation: Overview of Strategies . . . 7
2.2 EV Charging as a Phase-Balancing Resource . . . 8
2.3 Extended Objectives and Constraints in EV Phase Balancing . . . 8
2.3.1 Cost Optimization . . . 8
2.3.2 PV Integration . . . 9
2.4 Optimization-Based Phase-Balancing Strategies . . . 9
2.4.1 VUF Targeting vs Proxy Solution . . . 9
2.4.2 Optimization Algorithm and Load Flow . . . 9
2.5 Control Architecture . . . 11
2.6 Hardware Validation in EV Charging Research . . . 11
2.7 Comparative Overview . . . 13

3 Mathematical Model and Solution Method **15**
3.1 Control Framework and Scheduling Architecture . . . 15
3.2 Outer Loop . . . 16
3.2.1 EV Charging Model . . . 17
3.2.2 Installation Constraints . . . 19
3.3 Inner Loop . . . 20
3.4 Three-Phase Load Flow Model . . . 21
3.4.1 Network Representation . . . 21
3.4.2 Forward-Backward Sweep Algorithm . . . 22
3.4.3 Voltage Unbalance Factor Computation . . . 23

4 Simulation Setup and Case Studies **25**
4.1 Performance Metrics . . . 25
4.1.1 Load Imbalance . . . 25
4.1.2 VUF . . . 25
4.1.3 Per-Phase Voltage Magnitude . . . 26
4.1.4 EV Energy Delivery and SOC Fulfillment . . . 26
4.1.5 Charging Cost and Cost Premium . . . 26
4.1.6 Voltage Ramps/Solar Flicker . . . 26
4.2 Public Charging Facility Installation . . . 27
4.3 Test Network . . . 27
4.4 Input Data . . . 28
4.4.1 PV Supply . . . 29
4.4.2 Residential Load Profiles . . . 30
4.4.3 EV Fleet Schedules and Parameters . . . 30
4.4.4 Electricity Price Profile . . . 31
4.5 Scenario Design and Objective Weight Tuning . . . 31
4.6 Objective Weighting Tuning and Sensitivity Analysis . . . 32

5 Results and Discussion **33**
5.1 Baseline Scenarios . . . 33
5.1.1 Scenario 1 : Uncontrolled Charging . . . 33
5.1.2 Scenario 2 : Cost-Only Optimization . . . 34
5.2 VUF Optimized Results . . . 36
5.2.1 Load Balance and VUF . . . 36
5.2.2 Per-Phase Voltage Magnitude . . . 38
5.2.3 EV Energy Delivery and SOC Fulfillment . . . 38
5.2.4 Charging Cost and Cost Premium . . . 38
5.2.5 Voltage Ramps . . . 39
5.2.6 Other Observations . . . 40
5.3 Objective Weighting Tuning . . . 40
5.4 Sensitivity to Operating Conditions . . . 41
5.4.1 Charging Cost . . . 41
5.4.2 Load Imbalance and Voltage Unbalance . . . 42
5.4.3 Voltage Ramps . . . 43

6 Power Hardware-in-the-Loop Validation **45**
6.1 PHIL Setup Description . . . 45
6.1.1 Setpoint Scheduler . . . 46
6.1.2 Real-Time Simulator . . . 47
6.1.3 Power Amplifier and Interface Algorithm . . . 47
6.1.4 Device Under Test: Zaptec Pro Charger . . . 47
6.1.5 On-Board Charger Emulator . . . 48
6.2 Feasibility Results . . . 49
6.2.1 Optimizer Re-run for PHIL simulation . . . 49
6.2.2 DUT Command Execution Performance . . . 49
6.2.3 Setpoint versus Measured Power . . . 49
6.2.4 Schedule Execution over the Full Run . . . 50
6.2.5 Charged Energy Discrepancy . . . 51

7 Conclusions **53**
7.1 Summary of Findings . . . 53
7.1.1 RQ1: Joint formulation of cost, phase balance and flicker mitigation . . . 53
7.1.2 RQ2: Trade-offs and the cost premium of grid-friendly charging . . . 54
7.1.3 RQ3: Hardware-in-the-loop construction and the scope of its feasibility claims . . . 54
7.2 Scientific Contributions . . . 55
7.3 Recommendations . . . 56
7.3.1 For public charging facility operators . . . 56
7.3.2 For distribution system operators and regulators . . . 56
7.3.3 For implementation and further experimental work . . . 56
7.4 Limitations . . . 56
7.5 Future Work . . . 57

A Appendix : Simulation Setup Data **59**
A.1 Installation and Parking Data . . . 59
A.2 Test Network . . . 60

B Appendix B : Simulation Results Data **63**
B.1 Voltage comparison of main simulation . . . 63
B.2 EV Charging Load Profile Effect Towards Charging Cost . . . 63

C Appendix : PHIL Testbed Equipment **67**

References **70**

List of Figures

1.1 Global electric car sales rises rapidly in the last decade [1] . . . 1
1.2 Middle-voltage grid capacity map shows the extent of grid congestion across The Netherlands. Yellow: limited transport capacity available without queue, Orange: Area under investigation with queue, Red: Shortage of transport capacity with queue [2] . . . 2

2.1 Proposed switch phase flowchart by [3] . . . 10
2.2 Centralized/ remote control re-phasing illustration [4] . . . 11
2.3 [5] experimental setup for the voltage unbalance testing[5] . . . 12
2.4 Available power routing options (a) AC Load (b) grid-connected (c) circulating power[6] 13

3.1 Proposed model flowchart consisting two loops . . . 16
3.2 Illustration of EV charging KPI adapted from [7], as reproduced in [8]. . . . 18
3.3 CC/CV charging profile [9] . . . 19

4.1 Illustration of public charging facility . . . 27
4.2 Radial low-voltage test grid (BU06370601, Stedin). Node colors indicate phase connection: blue = L1, purple = L2, orange = L3, green = three-phase. The transformer substation is shown at the root. . . . 28
4.3 Power output profile of a three-phase 50 kWp rooftop PV system on a representative summer day. . . . 29
4.4 Total load of the simulated grid per phase without public charging facility load . . . 30
4.5 NL day-ahead electricity price profile for 24 July 2018, retrieved from the ENTSO-E Transparency Platform [10]. Prices are in EUR/MWh at hourly resolution. . . . 31

5.1 Per-phase current of scenario 1, indicative of uncontrolled charging behavior. The headroom is expanded during the day due to PV generation. . . . 34
5.2 Scenario 1 results: voltage profile and voltage unbalance factor. . . . 35
5.3 Example of detection of voltage flicker at boundaries(switching) and within it(solar flicker) 35
5.4 Per-phase current of Scenario 2, illustrating load concentration during low-price windows. 36
5.5 Scenario 2 results: voltage profile and voltage unbalance factor. . . . 37
5.6 Load balance comparison of 4 scenarios. Scenario 3 and 4 trace almost entirely similar balance during the simulation. . . . 37
5.7 VUF at tested bus of all scenarios. Sc3 and Sc4 track very similar profile, improving VUF during charging window . . . 38
5.8 Aggregate charging load of Sc4 reacts to changes of headroom/PV supply. . . . 40
5.9 Phase assignment of EVs in Sc4 . . . 40
5.10 Sensitivity of the main performance metrics to objective weighting in Scenario 3. . . . 41
5.11 Charging cost, cost premium against Sc2b, and PV utilization across the 5×5 fleet-size $\times$ PV-capacity grid, for Sc3 (top) and Sc4 (bottom). Cost and PV utilization share a color scale within their column so the two scenarios are directly comparable, and the premium panels share a scale centered on zero, with green indicating a case cheaper than Sc2b. 42
5.12 VUF sensitivity results for Scenarios 3 and 4. . . . 43
5.13 Flicker sensitivity result heatmaps. . . . 44

6.1 Overview of proposed PHIL testbed setup. . . . 46
6.2 Changing charging automatically from three-phase to single-phase with Zaptec chargers[8] 47
6.3 The setup of On-Board Charger Emulator, receiving single phase AC supply and feeding back the grid though three-phase connection . . . 48
6.4 Charger setpoint and load measurement during hardware test . . . 51

6.5 Cumulative charged energy over the comparison window: energy scheduled by the optimizer against energy measured at the DC load. The lower panel shows the corresponding instantaneous power. 51

B.1 Voltage profile comparison of the 4 scenarios in main simulation 63
B.2 Total EV charging power Sc3 : 11 EVs PV capacity 50kWp. 64
B.3 Total EV charging power Sc4 : 11 EVs PV capacity 50kWp. 64
B.4 Total EV charging power Sc3 : 9 EVs PV capacity 70kWp. 65
B.5 Total EV charging power Sc4 : 9 EVs PV capacity 70kWp. 65

C.1 Overall view of the PHIL testbed. 67
C.2 OPAL-RT digital real-time simulator running the three-phase load flow. 68
C.3 Cinergia GE-40 power amplifier. 68
C.4 Zaptec Pro charger, the device under test. 69
C.5 Walther-Werke EV signal tester. 69
C.6 Back-to-back Delta Elektronika supply pair forming the on-board charger emulator and power circulation unit. 70

List of Tables

2.1 Comparison of EV-based phase-balancing, VUF mitigation, or other related studies. . . 14

4.1 Low-voltage network characteristics. Extended detail in Appendix A.2 28
4.2 Overview of the four simulation scenarios. 32

5.1 Load imbalance and facility-bus VUF across scenarios. The Sc2 maximum VUF exceeds the EN 50160 limit of 2%. 37
5.2 Charging cost, cost premium, and PV utilization at the nominal operating point. The cost premium of the balancing scenarios is measured against Sc2b, the cost-only scenario with phase switching enabled, so that phase-switching freedom is held constant and the premium reflects the balance objective alone. 39
5.3 Voltage ramp and flicker statistics at the facility bus across scenarios. Event counts are decomposed by their timing relative to the 15-minute scheduling boundaries. Each phase is counted separately, so the counts are taken over phase-windows and are consistent with the occurrence percentages. 39
5.4 Range of facility-bus load imbalance and VUF across the 5×5 EV-count $\times$ PV-capacity grid. 43
5.5 Visible-flicker occurrence [% of time, facility bus] versus PV capacity, averaged over the EV-count axis. 44

6.1 Equipment used in the PHIL testbed. 46
6.2 Comparison of measured quantities at a single operating point. 50

A.1 EV Charging Installation Configuration – Site Alpha 59
A.2 Parking Spot Directory – Site Alpha 60
A.3 EV Charging Session Schedule – 14-Spot Scenario (23 July 2018) 60
A.4 EV Charging Session Schedule – 11-Spot Scenario 61
A.5 Low-voltage network characteristics 61

Nomenclature

Sets and Indices

$\mathcal{C}$	Set of EVs present at the installation during the scheduling horizon	
$\mathcal{C}_b$	Subset of EVSEs connected to branch b	
$\mathcal{P}_k$	Subset of EVs participating in the inner-loop modulation during window k	
$\mathcal{K}$	Set of long-period (15-minute) intervals, $\mathcal{K} = \{1, \dots, K\}$	
$\mathcal{S}$	Set of short-period (20-second) steps within a long-period interval	
$\mathcal{F}$	Set of phases, $\mathcal{F} = \{1, 2, 3\}$	
$\mathcal{N}$	Set of network buses	
$\mathcal{L}$	Set of branch segments	
c	Index of an individual EV	
k	Long-period time index	
s	Short-period time index	
ϕ	Phase index	
b	Branch index	
n	Bus index	
m	Forward-backward sweep iteration counter	

Decision Variables

$I_{c,k}$	Charging current setpoint of EV c at interval k	A
$s_{c,k}$	On/off binary; 1 if EV c is actively charging at interval k	
$z_{c,\phi,k}$	Phase-assignment binary; 1 if EV c is connected to phase ϕ at interval k	
$E_{c,k}$	Cumulative battery energy of EV c at the end of interval k	kWh
ξ_c	Unmet state-of-charge shortfall of EV c at departure	kWh
$i_{c,s}$	Inner-loop charging current of EV c at short step s	A
Δ_s	Common-mode current shift applied to every phase at short step s	A

Charging and Installation Parameters

Δt_1	Long-period interval duration, 15 min (0.25 h)	
Δt_2	Short-period interval duration, 20 s	
T	Total simulation horizon	h
t_{rt}	Real-time simulation timestamp	s

t_c^{arr}, t_c^{dep}	Arrival and departure time of EV c	
k_c^{arr}, k_c^{dep}	First and last interval index of the parking window of EV c	
$I^{\min}$	Minimum charging current permitted by IEC 61851-1, 6 A	A
$I^{\max}$	EVSE-rated maximum charging current, 32 A	A
$I_c^{\max}$	Maximum charging current of the on-board charger of EV c	A
$I_{c,s}^{\mathrm{hi}}$	Per-step upper current ceiling of EV c, tighter of hardware limit and SOC headroom	A
$I_{c,k}^{\mathrm{set}}$	Outer-loop current setpoint of EV c, used as the inner-loop anchor	A
$B_c^{\max}$	Maximum usable battery capacity of EV c	kWh
B_c^{arr}, B_c^{dep}	Battery energy of EV c at arrival and requested at departure	kWh
$\Delta B_c^{\mathrm{req}}$	Requested charging energy of EV c over the session	kWh
E_c^{init}	Battery energy of EV c at the start of its first active interval	kWh
$\mathrm{SOC}_{\mathrm{dep}}^{\max}$	Upper state-of-charge limit at departure, 90%	
η_c	Charging efficiency of EV c	
V_{LN}	Nominal phase-to-neutral voltage assumed in the optimizer, 230 V	V
I_ϕ^{cap}	Main installation fuse rating on phase ϕ	A
I_b^{branch}	Branch fuse rating of branch b	A
$I_{\phi,k}^{\mathrm{head}}$	Available current headroom on phase ϕ at interval k	A
$I_{\phi,k}^{\mathrm{load}}$	Office installation load current on phase ϕ	A
$I_{\phi,k}^{\mathrm{PV}}$	PV generation current on phase ϕ	A
$I_{\phi,k}^{\mathrm{ph}}$	Aggregate EV charging current on phase ϕ	A
$I_{\phi,k}^{\mathrm{grid}}$	Signed net grid current on phase ϕ; negative under PV export	A
$I_{\phi,k}^{\mathrm{imp}}$	Grid import current, $\max(0, I_{\phi,k}^{\mathrm{grid}})$	A
$I_{\phi,s}^{\mathrm{fix}}$	Current of non-participating vehicles on phase ϕ at short step s	A
λ_k	Day-ahead electricity price at interval k	€/MWh
Δ_s^{req}	Requested common-mode load adjustment at short step s	A
f_δ	Dead-zone function suppressing inner-loop response below a 2 A deviation	

Network and Load-Flow Quantities

$\mathbf{Z}_{ik}$	Three-phase branch impedance matrix of segment ik	Ω
Z_s, Z_m	Phase-conductor self-impedance and inter-phase mutual impedance	Ω/km
Z_0, Z_1	Zero-sequence and positive-sequence impedance	Ω/km
$R_{\phi\phi'}$, $X_{\phi\phi'}$	Resistance and reactance elements of $\mathbf{Z}_{ik}$	Ω
ℓ_{ik}	Length of branch segment ik	m
Z_{base}	Per-unit system base impedance	Ω

$V_{n,\phi}$	Complex bus voltage at bus n, phase ϕ	pu
$V_{n,\phi}^{d}$, $V_{n,\phi}^{q}$	Real and imaginary parts of $V_{n,\phi}$	pu
V_0, θ_ϕ	Slack bus voltage magnitude and phase angles	pu, deg
$I_{k,\phi}$	Nodal injection current at bus k, phase ϕ	A
$\mathbf{I}_{ik}$	Three-phase branch current vector of segment ik	A
$P_{n,\phi}$, $Q_{n,\phi}$	Active and reactive power injection at bus n, phase ϕ	kW, kvar
V_n^+, V_n^-	Positive- and negative-sequence voltage components at bus n	pu
V^0	Zero-sequence voltage component	pu
$\mathbf{A}$	Fortescue symmetrical-component transformation matrix	
$\varepsilon_{\mathrm{FBS}}$	Forward-backward sweep convergence tolerance, 10^{-6} pu	

Objective and Performance Metrics

C_{cost}	Charging cost objective term	€
$C_{\mathrm{imbalance}}$	Load-imbalance objective term	A
C_{penalty}	Unmet-energy penalty term	kWh
$C_{\mathrm{smoothing}}$	Inner-loop smoothing objective term	A
$C_{\mathrm{participation}}$	Inner-loop plan-anchor objective term	A
α, β, γ	Objective weighting coefficients for cost, imbalance and unmet energy	
ε	Inner-loop plan-anchor weighting coefficient, 10^{-3}	
VUF_n	Voltage unbalance factor at bus n	%
$\Delta V_{n,\phi,t}$	Voltage ramp over a fixed 20-second window	V
W	Voltage-ramp evaluation window length, 20 samples	
E_c	Total energy delivered to EV c	kWh
Ξ	Fleet-level relative energy shortfall	%
$\Delta C_{\%}$	Cost premium relative to the cost-only baseline	%

Photovoltaic Model Parameters

G^{inc}	Total module incident irradiance	W/m^2
G^{stc}	Irradiance at standard test conditions, 1000 W/m^2	W/m^2
G^{noct}	Irradiance at nominal operating cell temperature, 800 W/m^2	W/m^2
θ_M, Az_M	Module tilt and azimuth angle, $30°$ and $180°$	deg
θ_s, Az_s	Solar elevation and azimuth angle	deg
T_M, T_a	Module and ambient temperature	°C
T_{noct}	Nominal operating cell temperature, 44°C	°C
sw	Wind speed	m/s
a_f	Module absorptivity factor, 0.9	

A_M	Module area, 1.26 m^2	m^2
$P_{\mathrm{MPP}}^{\mathrm{stc}}$	Module maximum power point at standard test conditions, 245 W	W
$c_{P_{\mathrm{MPP}}}$	Temperature coefficient of maximum power, $-0.29\%/°$C	
η^{stc}	Module efficiency at standard test conditions, 0.1943	
η^{PV}	Real module efficiency	
η^{inv}	Inverter efficiency, 0.95	
FF	Module fill factor, 0.788	
P_{pv}	PV output power	kW

1 Introduction

This chapter establishes the motivation and scope of the thesis. Section 1.1 opens with the broader energy transition context, describing how the rapid uptake of low-carbon technologies in low-voltage distribution networks aggravates voltage unbalance on grid with inadequate mitigation. Section 1.2 narrows the focus to voltage unbalance as the specific power quality consequence of this asymmetry, defining the regulatory metric and the technical and economic consequences of exceeding it. Section 1.3 examines the existing literature on EV-based phase-balancing and identifies the gaps that this thesis attempts to address. Section 1.4 states the research objective and the three research questions pursued in response. Section 1.5 provides an overview of the thesis structure.

1.1. Context: Energy Transition and Low Voltage (LV) Grid Stress

The European energy transition is fundamentally restructuring electricity demand patterns at the residential level[11]. Driven by climate targets and national policy, the adoption of low-carbon technologies (LCTs), such as rooftop photovoltaic (PV) systems, electric heat pumps, and electric vehicles (EVs), is accelerating across the Netherlands and broader Europe [12][13]. In The Netherlands, despite the end of electric car purchase subsidy, electric car sales in the first quarter of 2025 were about 10% higher than during the same period[1]. While residential heat pump adoption is expected to grow substantially as gas boiler replacements accelerate globally [14]

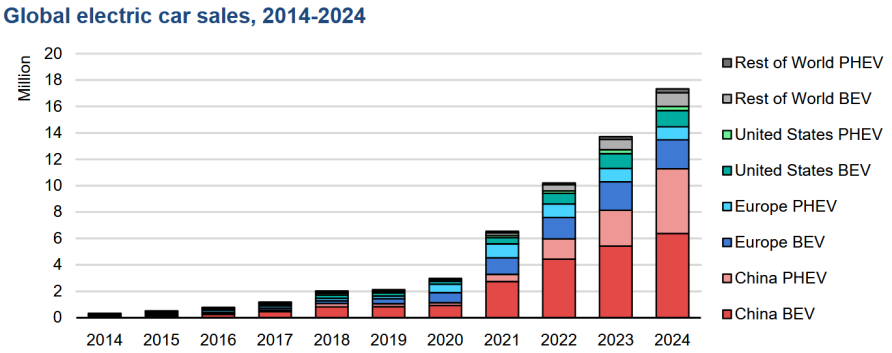

Figure 1.1: Global electric car sales rises rapidly in the last decade [1]

Even before the rapid growth of LCTs, electrical power grid infrastructures have been continuously facing challenges due to load variability. In The Netherlands, so called grid congestion resulted in 10,000 parties waiting for connections [15]. With the emergence of LCTs, loads became more responsive, variable, and in the case of PV, generative. This development is expected to enhance the preexisting issues of power grids.

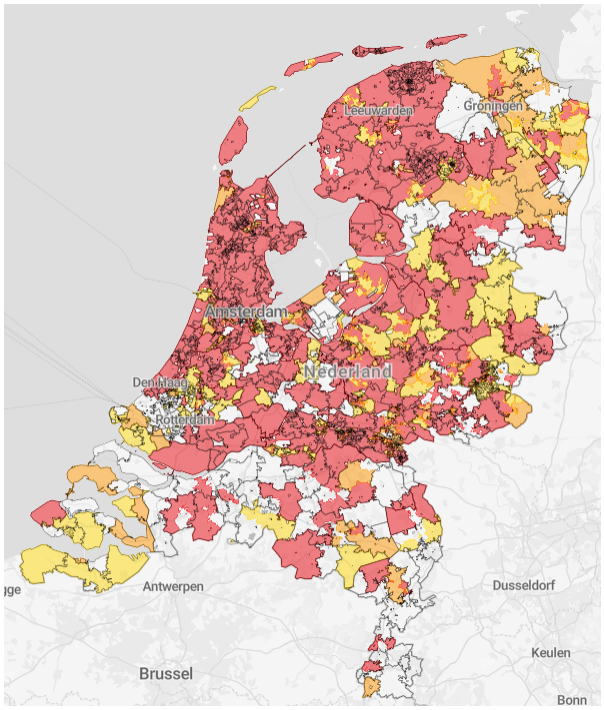

Figure 1.2: Middle-voltage grid capacity map shows the extent of grid congestion across The Netherlands. Yellow: limited transport capacity available without queue, Orange: Area under investigation with queue, Red: Shortage of transport capacity with queue [2]

More critically for this thesis, loads in low-voltage grids are structurally single-phase. In the Netherlands, as in most of Europe, residential connections are made single-phase to one of the three phases of the local LV feeder. Heat pumps below 3kW and residential PV systems below 10 kW are typically connected single-phase [11]. EV chargers at home and at smaller commercial installations are predominantly single-phase AC at 7.4kW. The result is an inherently asymmetric loading pattern that grows more pronounced as each of these technologies reaches higher penetration levels.

Crucially, these technologies do not act in isolation. [11] analyze six real Dutch LV distribution grids with different grid pattern, which are rural, suburban, and urban. They demonstrate that the combined integration of PV, EVs, and heat pumps can drive voltage unbalance well beyond regulatory limits, particularly during seasonal peaks [11]. The seasonal interaction is especially problematic: PV-heat pump combinations worsen unbalance in winter, while PV-EV combinations are most damaging in summer, when PV generation and evening charging demand both peak. The implication is that the unbalance problem is not attributable to any single technology but rather emerges from the compounding of several individually modest asymmetric loads onto feeders that already carry structurally uneven base loads.

This thesis is motivated by the intersection of two trends: the growing severity of phase unbalance in LV networks, and the simultaneous emergence of a controllable, flexible phase-switching EV charger, that can be turned from a contributor to the problem into part of the solution, if properly coordinated.

1.2. Problem: Voltage Unbalance and Its Consequences

Voltage unbalance in three-phase distribution networks arises when the three phase voltages differ in magnitude, phase angle, or both.

$$\text{VUF} = \frac{|V_{\text{neg}}|}{|V_{\text{pos}}|} \times 100\% \tag{1.1}$$

Under EN 50160, which governs power quality at the point of common coupling (PCC) in European LV networks, the VUF must not exceed 2% as a 10-minute mean value for 95% of one week [16]. When the VUF exceeds this threshold, a range of technical and economic consequences follow.

The direct consequences of unbalanced grid is inefficient grid infrastructure utilization [11]. From a technical standpoint, unbalanced voltages produce a negative-sequence current component in all

connected equipment. In distribution cables and transformers, this causes additional copper losses. The neutral conductor carries the residual unbalanced current, increasing the risk of thermal overload in networks with undersized neutral conductors. Per-phase voltage magnitudes deviate from their nominal values, with the most-loaded phase experiencing voltage depression and the least-loaded phase experiencing elevation, both of which can push sensitive equipment toward the boundaries of acceptable operating voltage [17]. Over time, sustained unbalance accelerates the aging of transformer windings and cable insulation through increased thermal stress [18].

From the perspective of the Distribution System Operator (DSO), sustained unbalance increases network losses, complicates transformer and cable sizing decisions, and in the worst case necessitates premature asset replacement or grid reinforcement. Existing technical remedies, such as static var compensators, D-STATCOMs, and passive load redistribution, are hardware-intensive, costly to deploy at the distribution level, and provide no demand response capability [18]. Software-defined control of existing flexible assets offers a more scalable and economically attractive path, provided the assets can respond with sufficient speed and precision. This is the central premise of this thesis.

1.3. Research Gap and Motivation

Coordinated phase-switching of EV chargers has been proposed as a software-defined alternative to hardware-intensive unbalance compensation methods, with simulation studies consistently reporting significant VUF reductions [3, 18, 19, 20]. However, a critical examination of this literature reveals three specific gaps that limit the reliability and practical relevance of the existing results.

The first gap is the limited inclusion of users' objectives and constraints in current existing studies, specifically public charging facility operators. With the advent of smart charging and open information of energy market, it is normal for these operators to optimize cost by maximizing installation utilization and maximizing charging power when electricity price is low[21]. In urban settings especially, where many chargers are located within public charging facilities such as apartment or office complex parking spaces, limitations and decisions of EV charger operator can make significant impact. Naturally, they do not have any inherent incentives to help maintain grid power quality. Thus, putting the localized constraints and objectives, such as installation capacity and cost optimization, will provide deeper insight on the potential of EV charging phase balancing with more realistic limitations.

The second gap concerns the fast voltage dynamics caused by rapid photovoltaic output fluctuations. In addition to VUF, rapid adoption of PV in LV grid produce irradiance ramp events on timescales of seconds, which propagate as voltage ramps at the facility bus and can cause visible or annoying light flicker for connected users [22]. The existing phase-balancing literature addresses VUF as a steady-state metric computed at 15-minute or longer resolution [3, 4], a timescale that is far too slow to respond to cloud transients. [22] demonstrate that EV load adjustment at 20-second resolution can effectively mitigate flicker, but their approach targets flicker in isolation without any phase-balance objective. No existing study addresses both VUF reduction and solar flicker mitigation simultaneously within a single coordinated EV charging framework. Whether a fast intra-timestep correction loop can be added to a phase-balancing outer loop without interfering with the phase-balance objective remains an open question. The third gap is the lack of hardware validation. All existing studies on this topic are validated through purely digital simulation, in which the EV charger is represented as an ideal controllable current source with a fixed or simplified power factor model. EV on-board charger currently remains as the bottleneck for faster grid services[8]. The power factor of an on-board charger varies with state-of-charge, the charger's response to a commanded current setpoint involves communication latency bounded by the IEC 61851-1 protocol, and phase switching introduces a transient that is absent from any simulation model[8]. These deviations from the ideal model are not merely inconvenient second-order effects. On the contrary, they directly affect constraint feasibility, control loop timing, and the achievable VUF reduction in a real implementation. No existing study has subjected a phase-balancing algorithm to a physical phase-switching charger under real grid voltages and currents.

Power Hardware-in-the-Loop (PHIL) simulation provides the methodological bridge between digital validation and field deployment. In a PHIL setup, the distribution grid is modelled in a real-time simulator (RTS), while the actual charger hardware is exercised with real voltages and currents through a power interface. This allows the charger's true dynamic behavior — its power factor response, switching transients, and communication characteristics — to be captured within an otherwise controlled and repeatable experimental environment [6, 21]. Despite the maturity of PHIL methodology for EV charging

research, no existing study has deployed it specifically for phase-balancing algorithm validation with a commercial phase-switching charger.

This thesis addresses the three gaps in sequence. The first is bridged by formulating the optimization around a public charging facility, so that installation capacity, branch protection and operator cost objectives constrain the balancing action, which yields a more realistic assessment of the available flexibility in urban settings. The second is addressed by adding a fast corrective layer that acts within the scheduling window, allowing the framework to respond to irradiance transients on a timescale the phase-balancing literature does not reach, without displacing the balance objective itself. The third is addressed by integrating a commercial phase-switching charger, specifically the Zaptec Pro, into a Power Hardware-in-the-Loop setup, so that the schedules produced by the optimizer are executed on physical hardware rather than on an idealized charger model. Together these steps constitute a more complete validation pathway than any existing work in the phase-balancing literature, combining operator-level objectives, fast flicker response, and hardware validation.

1.4. Research Objective and Questions

The objective of this thesis is to develop a dynamic phase assignment algorithm for fleets of phase-switching EV chargers, formulated to minimize voltage unbalance in radial LV distribution networks subject to a cost constraint, and to validate the technical feasibility of implementing this algorithm on physical charger hardware using a Power Hardware-in-the-Loop setup.

This objective is pursued through three research questions:

RQ1: *How can charging cost minimization, phase-balance, and solar-flicker mitigation be jointly formulated within a single coordinated EV charging framework, subject to public charging facility installation constraints?*

RQ1 addresses the design challenge of integrating three objectives that the existing literature has only ever pursued in isolation. Combining them naively is not guaranteed to work: a phase switch driven purely by cost could reintroduce the imbalance the framework aims to reduce, and a fast flicker-response correction operating on a 20-second timescale could interfere with the load-balance. Resolving this requires a formulation in which the three objectives are coupled through a shared decision structure. The formulation and its constraints are developed in Chapter 3.

RQ2: *To what extent do charging cost minimization, phase-balance improvement, and solar flicker mitigation trade off against one another when pursued simultaneously, and what cost premium does grid-friendly charging impose on the operator?*

RQ2 addresses the empirical consequence of the formulation developed under RQ1. Chapter 4 discusses the metrics and setup to evaluate the optimizer and the results are presented in Chapter 5.

RQ3: *How can the coordinated phase-balancing algorithm and a commercial phase-switching EV charger be jointly tested in a Power Hardware-in-the-Loop environment, and to what extent does the resulting setup demonstrate technical feasibility?*

RQ3 is deliberately posed in two parts. The first part concerns construction: no established testbed exists that couples a centralized phase-balancing optimizer to a commercial phase-switching charger, so the setup itself is a contribution rather than a given. This part is answered in Section 6.1, which documents the architecture, the components, and the substitutions made relative to the demonstrator of [6] and [21].

The second part concerns scope. Two objects are under test simultaneously: the optimizer, whose schedules must be physically executable, and the device under test (DUT), whose response must conform to the assumptions embedded in the optimization model. The question is therefore not how much better hardware performance is than simulation, but which feasibility claims the realized setup can and cannot support. Where the realized setup falls short of the intended architecture, the gap is reported as a bound on the claim rather than as a result, and as a specification for the work required to remove it.

1.5. Thesis Outline

The remainder of this thesis is structured as follows.

Chapter 2 reviews the literature on voltage unbalance compensation in LV distribution networks, with particular attention to EV-based approaches. It surveys the two methodological camps that have emerged in the field and examines the control architectures, optimization algorithms, and hardware validation work relevant to this thesis. The chapter closes with a comparative overview that identifies the gaps this thesis addresses.

Chapter 3 develops the mathematical model and solution method. The control framework is organized as a two-loop architecture operating at different timescales: an outer loop running at 15-minute intervals that determines phase assignments and base charging current setpoints through a Mixed-Integer Linear Program (MILP), and an inner loop running at 20-second intervals that adjusts charging current in response to instantaneous PV fluctuations and load ramps through a Linear Program (LP). The chapter presents the EV charging model, installation constraints, objective function, and the three-phase Forward-Backward Sweep load flow model used to evaluate VUF.

Chapter 4 defines the metrics, simulation setup, and case studies. It describes the public charging facility configuration, the LV grid, and the input data including PV generation profiles [23], NEDU residential load profiles[24], EV fleet schedules derived from the ElaadNL dataset[25], and the day-ahead electricity price profile from the ENTSO-E Transparency Platform[10]. Four simulation scenarios are defined, progressively adding layers of control from uncontrolled charging through cost-only optimization to full phase-balancing with inner-loop smoothing. The chapter also describes the sensitivity analysis conducted over the objective weighting coefficients.

Chapter 5 presents and discusses the simulation results across all four scenarios. It starts by establishing the baselines and subsequently analyzes the results of the proposed methods using the metrics defined by the chapter prior. The analysis of this result will answer RQ1 and RQ2.

Chapter 6 describes the Power Hardware-in-the-Loop setup and reports the hardware validation results addressing RQ3. The testbed uses an OPAL-RT OP5700 Digital Real-Time Simulator running the FBS load flow in MATLAB/Simulink, a Cinergia GE-40 three-phase power amplifier, a Zaptec Pro charger as the device under test, and a Delta Elektronika back-to-back power supply pair as the on-board charger emulator. The chapter characterizes the Zaptec Pro's phase-switching behavior, communication latency, and current modulation resolution in the context of the control framework's feasibility requirements.

Chapter 7 presents the conclusions, summarizing the findings per research question, stating the scientific contributions, acknowledging limitations, and recommending directions for future work.

2

Literature Review

This chapter reviews the literature on voltage unbalance compensation in LV distribution networks and other supporting studies. Section 2.1 defines VUF and establishes the categorization of strategies of VUF mitigation. Section 2.2 discusses EV potential for VUF mitigation while Section 2.3 elaborates on how other studies pairs VUF with other objectives, solving multiple problems under one model. Section 2.4 dives deeper into VUF minimization strategies by identifying voltage targeting strategy vs proxy solution and the components that builds the problem formulation. Section 2.5 discusses advantages and disadvantages of centralized vs decentralized system. Section 2.6 discusses one of the gap in this topic, that is the hardware validation or implementation. At the end Section 2.7 summarizes the gaps identified throughout the literature review, solidifying the contribution of the thesis.

2.1. Voltage Unbalance Compensation: Overview of Strategies

As established in Chapter 1, VUF is defined as $|V^-|/|V^+|$ following the IEC symmetrical component convention [16]. Recognizing the significant impact of zero sequence current in three-phase four-wire systems, which are common for LV distribution system, [17] proposed to include zero sequence voltage as part of the VUF formula, as seen on (2.1). Others, motivated by computational tractability, exclude the positive-sequence denominator altogether and work directly with $|V^-|$ as a proxy, which avoids the bilinear term $|V^-|/|V^+|$ in the constraint formulation. Since then, studies on VUF have been using different VUF equation, sometimes even combining the two principles, including zero-sequence voltage while excluding positive-sequence denominator. This thesis retains the standard IEC definition to maintain regulatory relevance and comparability with EN 50160, acknowledging that the zero-sequence exclusion introduces a modest conservatism in four-wire networks.

$$u = \sqrt{(V^-)^2 + (V^0)^2} \tag{2.1}$$

[18] summarized different VUF compensation approaches in their comprehensive review, organized into three categories [18]. The first is passive and power-electronics-based hardware. The passive approach are effective at times but require high capital and dedicated infrastructure, and offer no demand response capability, which is crucial in the context of highly volatile grids. Meanwhile, the active approaches such as load balancing or phase swapping require carefully maintained operation due to the potential of aggravating the unbalance. The second category uses three-phase DERs, primarily PV inverters with per-phase reactive power controllability, which can compensate unbalance dynamically but depend on available inverter headroom and introduce reactive power flows that may not align with network operator preferences. The third category uses single-phase DERs as controllable assets for phase-selective power redistribution. Voltage sensitivity analysis proves that this approach can achieve VUF targets comparable to hardware-based solutions[18].

The majority of studies that control EV charging load fall under the third category. Knezovic and Marinelli [26], Weckx and Driesen [19], [3] [3], [20] [20], and [27] all treat EVs as controllable loads with power factor control and/or power injections. To replicate this degree of freedom on hardware simulation, the onboard chargers need to be controllable. However, since this thesis aims to exclusively test phase-changing AC charger, this research will focus on methods in the first category. One

such example is study by [4] which demonstrated load phase switching, focusing entirely on dynamic phase assignments although it includes general load instead of just EV load. Another example by [5] demonstrates droop control using EV charger loads[5]. In this experiment, three EVs are each connected to one phase. However, the effectiveness of this method is contingent on having a chargeable EV present on each phase simultaneously, as the droop controller can only reduce the load on the phase it is connected to and has no mechanism to act on unoccupied phases. This dependency on coincidental fleet distribution limits its applicability in real parking scenarios where phase occupancy is inherently uneven. The approach taken in this thesis addresses this by operating at the level of the available current per phase across the entire charging park, steering allocation across the fleet through DPM regardless of how EVs happen to be distributed among phases. Together, [4] and [5] establish the two pillars of the method adopted here: dynamic phase-selective load assignment and per-phase load coordination without reactive power control, both within the first category of taxonomy from [18].

2.2. EV Charging as a Phase-Balancing Resource

Among residential loads, single-phase EV charging possesses properties that make it uniquely tractable for active phase balancing. Energy is delivered to the vehicle battery at the commanded rate regardless of which phase supplies the current, making phase switching entirely unnoticeable to the end user. Moreover, unlike general load, AC EV charging infrastructure does not contain motor loads such as heatpumps [4], making EV charging interruption during phase switch safer and less strenuous for the load.

The practical realization of EV-based phase balancing requires charger hardware capable of selecting and switching between phase connections under software control. Conventional single-phase AC chargers are wired permanently to one phase. Three-phase chargers distribute load across all three phases simultaneously. Phase-switching chargers resolve this constraint by using internal relay switching to connect to any of the three available phases on command, effectively providing the phase-selectivity of single-phase operation with the dynamic flexibility needed for optimization-based control.

[8] provides an experimental characterization of phase-switching charger behavior using commercial hardware including a Zaptec Pro, testing both power modulation and phase switching across three EV models [8]. Their results demonstrate that phase switching can be completed while maintaining compliance with IEC 61851-1, that the bottleneck for providing fast grid services is the EV on-board charger rather than the EVSE, and that switching transients are sufficiently brief that they do not constitute violations of the charging protocol. These experimental findings are directly relevant to the PHIL feasibility assessment in Chapter 6 of this thesis. The Zaptec Pro charger exposes a cloud-accessible REST API that enables algorithmic control of charging current and phase selection from a remote controller [28], providing the communication layer through which the MILP optimizer dispatches commands in the experimental setup.

2.3. Extended Objectives and Constraints in EV Phase Balancing

2.3.1. Cost Optimization

In addition to the voltage unbalance, some studies also embed other objectives and/or constraints to the problem formulation to more closely resemble real low-voltage grid, the problems it is facing due to low-carbon technology, and adjacent trends in EV charging algorithm. One of the common objective is minimize cost. With real-time pricing (RTP), consumers are incentivized to alter the charging schedule and many studies have created algorithm to achieve that goal, both within and outside the grid unbalance topic [21, 27, 29]. Beside benefit for consumer, the consumer's behavior in response to dynamic electricity price could also significantly mitigate grid congestion[30, 31]. However, utilizing EV charging as grid services such as phase balancing could rise the total charging cost relative to the optimized cost, since the optimizer needs to deviate from the lowest price charging window to satisfy the other objectives. Thus, comparing fully cost optimizing algorithm and one that include load balancing can give insight on the relative cost penalty cause by this service.

2.3.2. PV Integration

Rapid irradiance fluctuations caused by cloud transients produce short-term voltage fluctuations and subsequently light flicker[32]. [22] document this phenomenon using measurements from a real LV feeder in Utrecht, the Netherlands, demonstrating that PV output fluctuations at 2-second resolution produce voltage fluctuations that can cause visible and annoying light flicker for a significant portion of the day in 2050 scenarios [22]. In this thesis, these voltages fluctuation is referred as solar flicker. They propose mitigating these fluctuations by altering EV charging processes — modulating charging power to absorb ramp events — and demonstrate effective flicker reduction at the expense of increased charging costs. Their model operates on 20-second timesteps, driven by Dutch automatic frequency restoration reserve (aFRR) price signals, and uses linear programming to schedule charging power. This model is comparable to [19] because both are targeting load balance across the grid, only with different goals. VUF improvement studies using optimization acknowledge the impact of PV supply fluctuation and the increasing trend of PV penetration. However, cloud transient and the subsequent solar-flicker is not usually mitigated. [3] for example, have a 15-minute interval for the optimization which is far too slow to respond to solar flicker [3]. Modification in problem formulation and decreasing this interval to 20 seconds could potentially solve solar flicker as well. However, computational load will rise exponentially and new set of data needs to be provided. To avoid this problem, this thesis attempts to address solar flicker using additional optimizer with 20-seconds resolution.

2.4. Optimization-Based Phase-Balancing Strategies

2.4.1. VUF Targeting vs Proxy Solution

Optimization-based approaches to EV phase balancing can be divided into two methodological camps depending on whether the controller acts directly on VUF or on a simpler proxy. [19] is an example of proxy solution by load balancing[19]. Their objective problem is the minimization of inter-phase load variance. Since variance is a quadratic term, their problem formulation can be classified as quadratic programming(QP). Their result shows that although load balancing is simpler, it has better impact on grid voltages and losses compared to other strategies such as PV inverter power control. This is because despite significant voltages, PV inverter may increase component loading and consequently the power losses. Load balancing problem is a relatively simple to solve convex optimization problem, which is advantageous for faster wide spread adoption due to lower computational load.

However, [20] provides theoretical justification for targeting VUF directly instead of proxy by arguing *"Due to asymmetric structure of secondary distribution networks, active power balancing does not necessarily lead to operational quantities balancing; so considering network and power flow constraints is necessary"* [20]. The viability of active power balancing is indeed dependent on the grid structure. Thus, more advanced methods in the topic usually targets VUF directly and includes load flow model into the constraints of the problem formulation. [3] coordinate PEV charging power and phase switching using a genetic algorithm with load flow evaluation at each generation, demonstrating meaningful VUF reduction over a 24-hour horizon on a real unbalanced LV network [3]. [20] extend this to simultaneous unbalance and harmonic mitigation via iteratively linearized stochastic nonlinear programming with three-phase harmonic power flow constraints [20]. Franco et al. formulate the EV charging coordination problem in unbalanced distribution systems as a MILP, embedding a full three-phase current injection model as equality constraints [29]. Unfortunately, this improvement comes with exponential increase of computational load. This is important to consider since real low voltage grid can contain hundreds or thousands of buses.

Considering inclusion of public charging facility and solar flicker response in this thesis, local load balancing will be implemented. This lower computational load is critical considering the response time that is required to respond to local load ramps or switching and solar flicker.

2.4.2. Optimization Algorithm and Load Flow

In the previous subsection, it is established that studies using mathematical programming whose objectives directly address buses' voltage or VUF needs to embed load flow into the problem formulation unlike studies with proxy solution where load flow becomes validator of the decision created by the optimizer. [19] for example use backward-forward sweep (FBS) with 15- minute time resolution. Load flow simulation is more complicated to apply for voltage/VUF targeting optimization because load flow formulation is inherently nonlinear non-convex problem. Nonlinear non-convex optimization can some-

times be stuck in local optima. Thus, many studies that choose to use mathematical programming often linearize and use linear programming (LP) or mixed integer linear programming (MILP) as optimizer [20, 29].

Other studies employ metaheuristic algorithm which essentially decouple load flow from the main problem formulation. This allows them to use preexisting load flow simulation, allowing the true nonlinear power flow to be used without approximation, which means the load flow result requires no separate validation for linearization error. [3] who used genetic algorithm for example, implemented backward-forward sweep (FBS) method to calculate load flow and bus voltage [3]. As seen on figure 2.1 the genetic algorithm process the data population into decision variables which is then tested by the subsequent FBS load flow. Metaheuristics algorithm solution is near-optimality with varying degree depending on the algorithm[33]. Thus, the result might not be as optimal as mathematical programming. However, metaheuristics has a lower chance of infeasibility because mathematical programming is brittle especially as the number of variable increases.

In terms of solve time, mathematical programming especially LP tends to perform better[33] compared to metaheuristics. However, the computational load of metaheuristics optimization is actually more scalable with trade-offs such as increasing error threshold or decreasing iteration number. [33] provides a comprehensive review of metaheuristics algorithms that summarized the solution quality and time performance. [4] further decreases the search space by running sensitivity analysis to determine candidate buses for phase switching. The selection of load flow also affect optimization time. [34] conclude that FBS achieve the lowest computational time, relative to first-order Taylor(FOT) series approximation and fixed-point linearization, for modelling radial distribution grids with voltage unbalance.

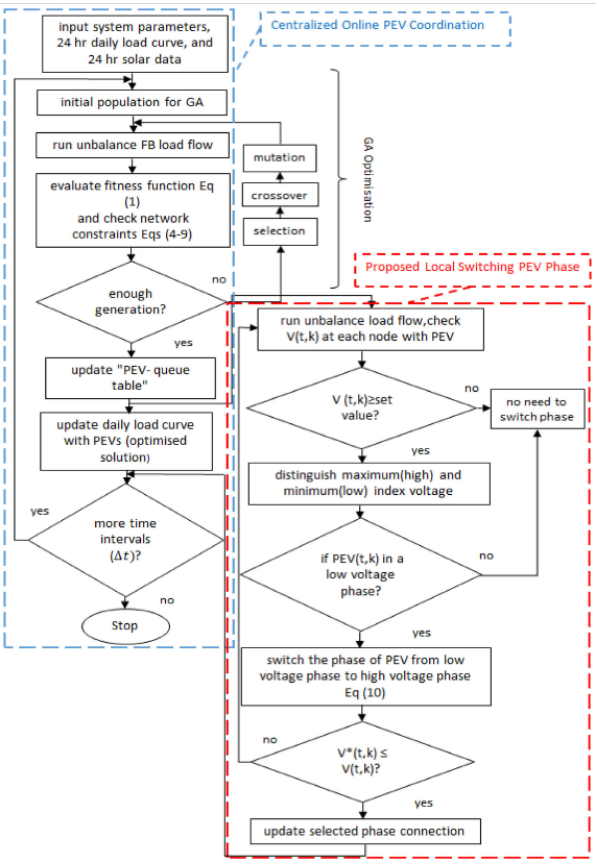

Figure 2.1: Proposed switch phase flowchart by [3]

Considering the computational load increase due to inclusion of solar flicker response, this thesis studies optimization on load level instead of grid or voltage level. Specifically, a modified version of load balancing problem formulation in [19] is adopted. The modification opt for linear version of balance minimization considering the addition of binary variable due to phase-switching significantly increase

solve time. Because load-balancing is used instead of voltage targeting strategy, a separate load flow is necessary to verify the VUF.

2.5. Control Architecture

The literature on EV-based phase balancing employs centralized and decentralized control architectures. Wu et al. provide an overview of hierarchical EV charging station operation in smart grid applications, covering the trade-offs between these architectures in terms of scalability, optimality, and communication requirements [35]. The systems used in these studies are LV grids. Thus, centralized control scheme in this topic typically refers to a system where one actor collectively makes decisions for one LV grid, as seen in example by [4] on figure 2.2.

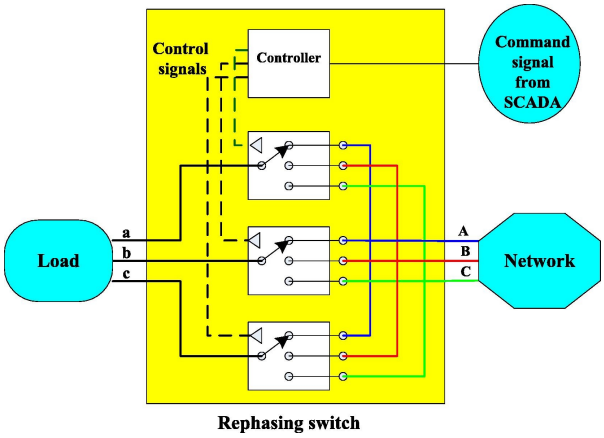

Figure 2.2: Centralized/ remote control re-phasing illustration [4]

Voltage-targeting optimizations uniformly require centralized control: computing VUF at every bus requires network-wide voltage state information, which cannot be obtained from local measurements alone. [3, 4, 20] all adopt centralized architectures for this reason. [19] which implemented grid level load balancing naturally also uses centralized control. Furthermore, centralized control tend to have lower cost and it is easier to be extended into hierarchical scheme [26]. However, this control scheme relies on reliable communication infrastructure which adds complication to the system.

On the other hand, decentralized control can be utilized for local load-targeting approaches or power droop, since per-phase voltage at the connection point is a locally observable quantity. [5] demonstrate this with a per-phase voltage droop controller implemented independently at each EVSE, where the charging current is adjusted proportionally to the locally measured phase-to-neutral voltage with no communication between chargers [5]. Each EV reacts only to its own phase voltage, making the approach entirely infrastructure-free and trivially scalable. However, decentralized control instantaneous response to measurement could lead to instability in certain cases. [36]. In phase-switching control specifically, if all chargers independently observe the same phase imbalance signal and respond simultaneously, they risk all switching to the same phase in unison, inverting rather than correcting the imbalance. [18] confirm this structural limitation, noting that decentralized controllers operating without inter-phase coordination may cause over-compensation due to the absence of coordination between distributed assets [18]. This simultaneity risk is a structural argument for centralized coordination when phase switching is involved.

Considering the limitation of controllable EVs that are located in one place, the optimizer built in this thesis is more akin with decentralized architecture if observed from grid perspective, but multiple chargers are coordinated with one optimizer and no other EV charger in the grid is controllable. Thus, the instability problem is no longer a risk despite being a decentralized approach.

2.6. Hardware Validation in EV Charging Research

The translation from simulation to physical hardware is a critical and largely unaddressed step in the EV phase-balancing literature due to the high cost of hardware implementation and experimental verification[37]. This fact creates gap for real-life implementation numerous findings in the topic of

EV as voltage-balancing agent which needs to be addressed because upon implementation a set of technical issues or regulatory limitations might surface. Identifying these issues will then create a more ready-to-apply solution on this topic.

[5] provide the closest existing work to hardware-validated phase-balancing, demonstrating EV capability for voltage unbalance compensation using real vehicles in a laboratory environment [5]. Figure 2.3 depicts their experimental setup which uses real physical hardware throughout — actual cables, a real wind turbine, and commercial EVSEs — with the droop controller implemented as software on each EVSE. Thus, the setup is incapable of simulating on grid level as it is limited to the cabling past the grid connection. This setup is enough to validate decentralized system to a certain degree, but scalability quickly becomes a problem as the validation requires the construction of real grid. Therefore, another method of validation is necessary to accommodate simulation in a bigger grid.

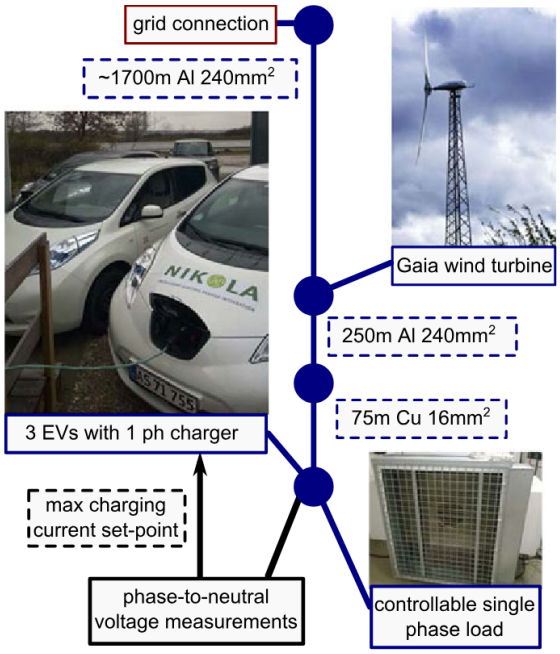

Figure 2.3: [5] experimental setup for the voltage unbalance testing[5]

Power Hardware-in-the-Loop (PHIL) simulation solves this issue by incorporating software simulated grid and hardware connection. PHIL can be described as a hybrid testing technique which decouples the grid, which is simulated in a Real-Time Simulator (RTS), from the rest of the circuit in the form of actual device under test (DUT) [38]. This setup circumvent the prior limitation that necessitates complete real hardware setup, which allows research of impact of smart concepts on large-scale grids in real-time in a laboratory environment under a wide range of conditions and in a repeatable, safe, and economical manner. [6] developed the PHIL testbed in TU Delft, with the objective of observing the impact of electric vehicle charging on distribution grid. The LV distribution grid is represented by a Digital Real-Time Simulator (DRTS) whose outputs are then amplified to synthesize the three-phase node voltage, while the physical DUT is a commercial EVSE. Figure 2.4 illustrates three different power routing options. [6] setup is built using the third option considering the energy saving advantage of the setup.

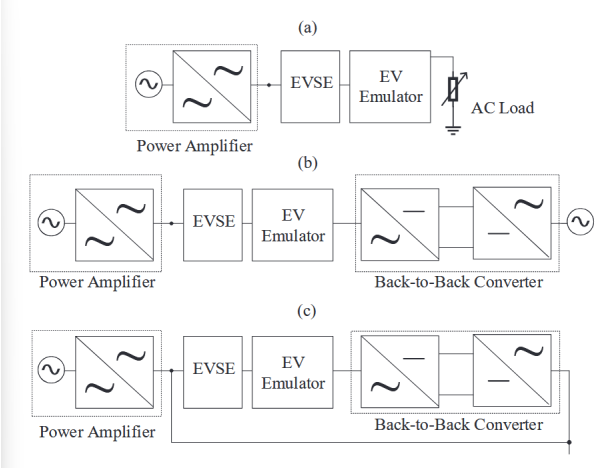

Figure 2.4: Available power routing options (a) AC Load (b) grid-connected (c) circulating power[6]

[21] is building upon [6] work and some improvements were made, which includes incorporating hierarchical mixed-integer programming EV smart charging algorithm[21]. This study evaluates the algorithm across eight scenarios including cases with uncertain solar and load predictions and engaged grid capacity restrictions. The algorithm targets cost and peak demand reduction but the PHIL methodology and testbed architecture are directly relevant and will be used as starting point for this Thesis.

2.7. Comparative Overview

Table 2.1 draws the reviewed studies together along the dimensions that matter for this thesis: the controlled quantity, the optimization method and its treatment of the load flow, the control architecture, and the extent of validation. Comparing the studies on these axes exposes a consistent set of trade-offs rather than a single dominant approach. The first axis is what the controller acts upon. One group targets voltage unbalance directly [3, 4, 20], while the other acts on a load-balancing proxy [19]. Direct VUF targeting yields a controller whose objective aligns with the regulatory metric directly, and thus promising more robust result. However, it is inseparable from a three-phase load flow and therefore incurs a substantially higher modelling and computational cost. Proxy load balancing sacrifices that direct alignment in exchange for a convex, low-cost formulation whose speed suits the fast response required for local ramps and solar flicker.

The second axis is how the load flow enters the problem, and it does not follow neatly from the first. Among the voltage-targeting studies, the mathematical formulations embed the power flow as explicit constraints [20, 29], guaranteeing constraint satisfaction but at the price of linearization error and brittleness as the number of variables grows. The metaheuristic formulations instead retain the true nonlinear power flow as an external evaluator [3, 4], trading guaranteed optimality for robustness and scalability. This distinction cuts across the voltage-versus-proxy split: a study can target voltage yet still evaluate the load flow externally.

The third axis, control architecture, is largely determined by the first. Voltage targeting requires network-wide state and is therefore uniformly centralized [3, 4, 20], whereas local load-balancing and droop schemes can be decentralized [5, 26] at the risk of over-compensation when uncoordinated. Meanwhile, [19] attempt at grid-wide balance requires also centralized control. The approach in this thesis occupies a middle position: decentralized from the grid's perspective, but internally coordinated through a single optimizer, so that risk does not arise.

The final and axis is validation where comparison reveals that the desirable properties appear only in isolation. The voltage-targeting studies are validated in simulation, some of them against standard test feeders [4, 20, 29]. Conversely, the hardware-validated studies either uses droop strategy [5] or do not address unbalance at all [21]. No reviewed study combines phase-changing and load balancing, a benchmark-validated load flow, and hardware validation. Closing this combination is the aim of the remainder of this thesis, Chapter 3 develops the optimization formulation and the load-flow model used to validate it, before the hardware evaluation is taken up in the subsequent chapters.

Table 2.1: Comparison of EV-based phase-balancing, VUF mitigation, or other related studies.

Study	Year	*EV*	*Voltage target*	*Flicker*	*Harmonics*	Mechanism	Optimizer	Architecture	*LF in formulation*	*Benchmark val.*	*HW validated*
Saleh & Latify [20]	2024	✓	✓	–	✓	Reactive/active injection + harmonic control	Iterative LP (stochastic)	Centralized	✓	✓	–
Jabalameli & Ghosh [3]	2021	✓	✓	–	–	Charging + var + phase switching	GA (FBS)	Centralized	–	–	–
Franco et al. [29]	2015	✓	–	–	–	EV charging schedule	MILP	Centralized	✓	✓	–
Soltani et al. [4]	2017	–	✓	–	–	Automatic re-phasing	MSFLA	Centralized	–	✓	–
Weckx & Driesen [19]	2015	✓	–	–	–	Per-phase active power balance	QP (FBS)	Centralized	–	–	–
Akhavan-Rezai et al. [27]	2017	✓	–	–	–	Demand response + PV	–	Centralized	–	–	–
Knezovic & Marinelli [26]	2016	✓	–	–	–	Phase-wise reactive droop	Droop	Decentralized	–	–	–
Martinenas et al. [5]	2017	✓	–	–	–	Local voltage droop	Droop	Decentralized	–	–	✓
Brinkel et al. [22]	2020	✓	–	✓	–	Smart charging ramp limiting	LP	Centralized	–	–	–
Yu et al. [21]	2024	✓	–	–	–	Cost/peak reduction	MILP	Centralized	–	–	✓
This thesis	**2026**	✓	✓	✓	–	**Phase switching + DPM**	**MILP**	**Centralized**	–	–	✓

Abbreviations: GA = Genetic Algorithm; MSFLA = Modified Shuffled Frog Leaping Algorithm; QP = Quadratic Programming; LP = Linear Programming; MILP = Mixed-Integer Linear Programming; FBS = Backward-Forward Sweep (load flow used as an external evaluator).
Notes: ✓ = addressed; – = not addressed. “LF in formulation” distinguishes mathematical methods (power flow as explicit constraints) from metaheuristic and proxy methods (power flow as external evaluator or absent); any method annotated “(FBS)” therefore evaluates the load flow externally. Benchmark validation refers to testing on a standard test feeder: [29] use the IEEE 34-bus feeder; [4] use the IEEE 37-bus feeder; [20] use the standard 39-bus test system; [3] use a real 74-node Western Australian feeder (not a standard benchmark).

3

Mathematical Model and Solution Method

This chapter develops the mathematical model and solution method. It starts with Section 3.1 which paints the overview of the proposed framework, consisting of outer and inner loop. These loops address different problems and are detailed by Section 3.2 and Section 3.3. Finally, Section 3.4 proposes validation method for the optimization result using forward-backward sweep.

3.1. Control Framework and Scheduling Architecture

As indicated in Section 1.3, this thesis aims to build a model to mitigate voltage unbalance using public EV charging facility. The charging facility consists of EV chargers, PV, and small office installation. Considering the computational load, proxy solution using load balancing is adopted with the main switch box as point of measurement. Cost optimization goal is also embedded into objective function to analyze the effect of load balancing towards overall charging cost. It is clear that the magnitude of effect depends on the particular condition of the scenario, but this thesis aims to show the correlations qualitatively. As the public charging facility is equipped with PV, the system also considers solar flickers and aims to minimize voltage problems due to the flicker. Consequently, higher resolution data and quicker response or adjustment from the system is needed [22]. Phase-shifting will also be a feature for load balancing, creating more degree of freedom which is expected to result in better balance. However, the process can be quite disruptive and cannot be executed too often. Moreover, the binary variable introduced by phase assignment significantly increase the load of the optimizer.

Considering the different response time required for the load balancing and solar flicker response, a two-loop system is adopted. Outer loop with 15-minute interval mainly address load imbalance, while inner loop with 20-second interval mainly address solar flicker and other changes in load.

The two loops from the previous section are now arranged into a flowchart Figure 3.1. The first optimizer timestep index is called long-period index (k), with interval of 15 minutes($\Delta t_1 = 15\,\mathrm{min}$), while the second optimizer timestep index is called short-period index (s), with interval of 20 seconds($\Delta t_2 = 20\ \mathrm{s}$). Phase assignment is enabled on the outer loop which adds binary variables. Thus, the outer loop problem becomes MILP while the inner loop remains simpler and faster LP. The other parts of the flowchart is part of the PHIL validation, which is detailed in Chapter 6.

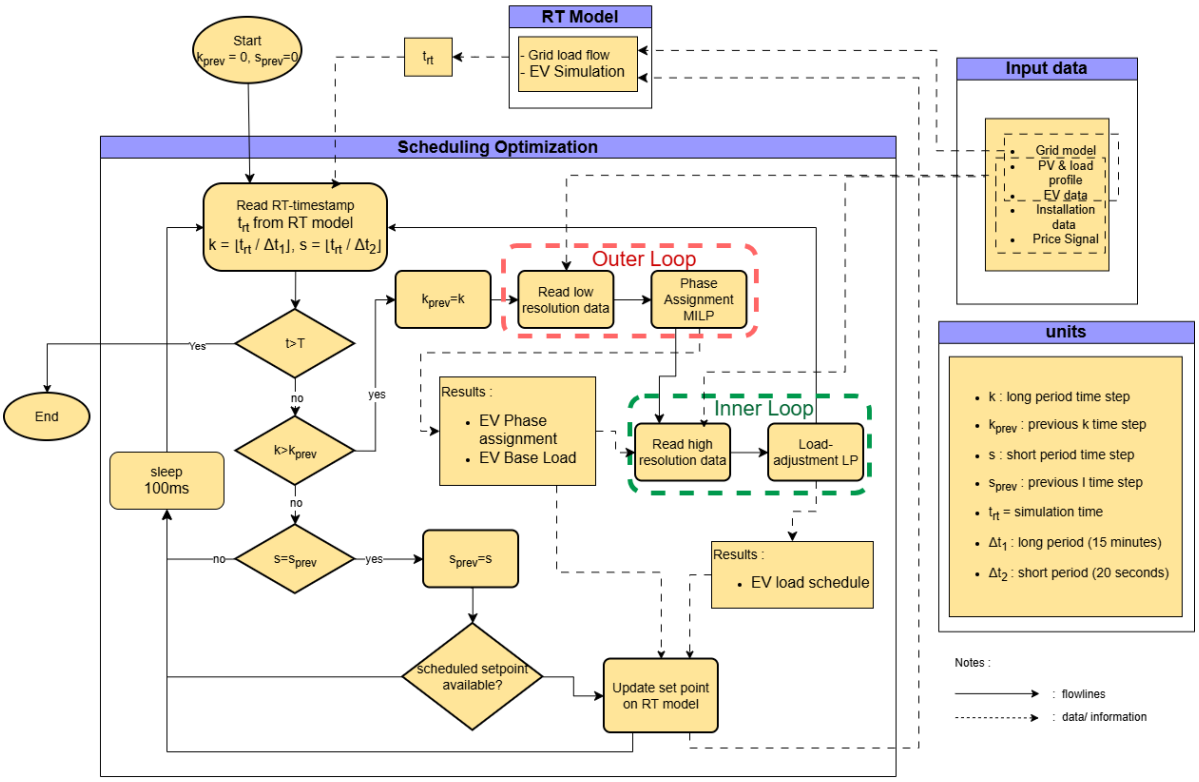

Figure 3.1: Proposed model flowchart consisting two loops

At initialization the previous long-period index k_{prev} and the previous short-period index s_{prev} are both set to zero. On every iteration of the control loop the optimizer reads t_{rt} and derives two integer time indices:

$$k = \left\lfloor \frac{t_{\text{rt}}}{\Delta t_1} \right\rfloor, \qquad s = \left\lfloor \frac{t_{\text{rt}}}{\Delta t_2} \right\rfloor, \tag{3.1}$$

If t_{rt} has exceeded the total simulation horizon T the loop terminates, otherwise the control logic checks which tier of the hierarchy has advanced. When neither index has changed, the loop sleeps for $100\ \text{ms}$ before re-polling, preventing busy-waiting without introducing meaningful latency.

When $k > k_{\text{prev}}$ a new long-period interval has elapsed. The optimizer sets $k_{\text{prev}} \leftarrow k$, reads the low-resolution input data (15 minutes interval), and solves the outer loop optimization. This problem determines two outputs for each connected EV over the upcoming scheduling horizon:

1. the phase to which each charger is assigned (a binary decision variable, fixed for the entire scheduling horizon);

2. the base load, which is the nominal charging current setpoint at each 15-minute interval.

The problem solved at every k is a receding horizon starting from the current k until K, the last index. Therefore, the horizon becomes smaller for every rerun.

When $s > s_{\text{prev}}$ but k has not changed, a new short-period interval has elapsed within the same 15-minute window. The optimizer sets $s_{\text{prev}} \leftarrow s$, reads the updated available current from the RT model (high-resolution data), and solves the Load-Adjustment LP.

3.2. Outer Loop

The outer loop consists of three objectives with adjusted weighting coefficients (α, β, and γ) as seen on (3.2), which are decided empirically and can be adjusted depending on the prioritization of the system. The first part is cost optimization objective (C_{cost}) solved with simply multiplying the total energy imported by the facility by the day-ahead energy cost (λ_k). The formula uses grid import current $I_{\phi,k}^{\text{imp}} =$

$\max(0, I^{\text{grid}}_{\phi,k})$, the non-negative part of the signed net grid current defined in (3.4). Exported surplus PV energy is not penalized, as no feed-in tariff is modelled. Energy cost from PV is assumed to be 0€/kWh, pushing the algorithm to maximize local PV load utilization. Voltage (V_{LN}) is assumed at fixed 230V maintaining simplicity and avoiding grid physics inclusion. The second component is the load balancing optimization ($C_{imbalance}$) which minimizes the deviation of the signed net grid current $I^{\text{grid}}_{\phi,k}$ on each phase from the three-phase mean. The signed quantity is used rather than import alone because a phase exporting surplus PV is equally unbalanced as one importing heavily. The formulation was derived from [19] that minimizes variance which is a quadratic programming. However, the inclusion of binary variables due to phase switching turns the problem into MIQP which raise the computation time considerably. Thus, the current formulation is adopted with the trade-off of inability to penalize unbalance spikes and only suppress the sum across timesteps. The third part is unmet charging power penalty ($C_{penalty}$). This penalty is chosen instead of using hard constraint to avoid rigid problem that results in infeasibility.

$$\min \quad \underbrace{\alpha \sum_{k\in\mathcal{K}} \sum_{\phi\in\mathcal{F}} \lambda_k \cdot \max\left(0, I^{\text{grid}}_{\phi,k}\right) \cdot V_{\text{LN}} \cdot \Delta t_1 \cdot 10^{-3}}_{C_{\text{cost}}} + \underbrace{\beta \sum_{k\in\mathcal{K}} \sum_{\phi\in\mathcal{F}} \left| I^{\text{grid}}_{\phi,k} - \frac{1}{3} \sum_{\phi'\in\mathcal{F}} I^{\text{grid}}_{\phi',k} \right|}_{C_{\text{imbalance}}} + \underbrace{\gamma \sum_{c\in\mathcal{C}} \xi_c}_{C_{\text{penalty}}} \tag{3.2}$$

, where the shortfall ξ_c is the energy still missing at departure relative to the requested target,

$$\xi_c = B^{\text{dep}}_c - E_{c,\,k^{\text{dep}}_c-1}, \qquad B^{\text{dep}}_c = B^{\text{arr}}_c + \Delta B^{\text{req}}_c, \qquad \xi_c \geq 0 \qquad \forall\, c \in \mathcal{C} \tag{3.3}$$

and is enforced through the departure constraint (3.10). Here $E_{c,\,k^{\text{dep}}_c-1}$ is the battery energy at the last complete interval during which EV c is present, and B^{dep}_c is the energy requested for the session. A vehicle that reaches its target therefore contributes nothing to C_{penalty}, while one that departs short contributes in proportion to the deficit. Because γ is set well above α and β, the term acts as a soft constraint. It keeps the problem feasible when a session cannot be served in full, rather than forcing the solver to declare infeasibility, and

$$I^{\text{grid}}_{\phi,k} = I^{\text{load}}_{\phi,k} + I^{\text{ph}}_{\phi,k} - I^{\text{PV}}_{\phi,k} \qquad \forall\, \phi \in \mathcal{F},\ k \in \mathcal{K} \tag{3.4}$$

The signed net grid current $I^{\text{grid}}_{\phi,k}$ is given by (3.4). It can be negative when PV generation exceeds local load and EV demand. The cost term uses only the import portion $\max(0, I^{\text{grid}}_{\phi,k})$, while the imbalance term uses the full signed value so that net-export phases are penalized equally with net-import phases. It is important to note that $I^{\text{PV}}_{\phi,k}$ in basic form is the minimum measurement in 15 minutes to avoid negative value headroom.

3.2.1. EV Charging Model

This section defines the mathematical model of the EV and EVSE as seen by the outer loop optimizer. The framework adheres to IEC 61851-1 Mode 3 AC charging, in which the EVSE communicates a maximum current permission to the vehicle via a PWM control pilot signal[39]. Additional constraints and limitations are put in place to anticipate hardware validation which is part of the research question. 3-phase Charging mode is disabled due to DUT's inability to switch from three-phase to single-phase seamlessly especially when choosing phase L2 or L3 during the switch[8]. V2G operation is also out of scope so all chargers are modelled as unidirectional loads.

Decision Variables

Let $\mathcal{C}$ denote the set of EVs present at the installation during the scheduling horizon $\mathcal{K} = \{1, \ldots, K\}$ of K long-period intervals, each of duration $\Delta t_1 = 0.25\,\text{h}$. The following decision variables are defined for each EV $c \in \mathcal{C}$ and interval $k \in \mathcal{K}$:

- $I_{c,k} \geq 0$ charging current setpoint [A];
- $s_{c,k} \in \{0, 1\}$ on/off binary: $s_{c,k} = 1$ if EV c is actively charging at interval k;
- $E_{c,k} \geq 0$ cumulative battery energy at end of interval k [kWh];

- $z_{c,\phi,k} \in \{0,1\}$ phase-assignment binary: $z_{c,\phi,k} = 1$ if EV c is connected to phase $\phi \in \mathcal{F} = \{1,2,3\}$ at interval k;
- $\xi_c \geq 0$ unmet SOC shortfall at departure [kWh].

It is important to note that the currents in the decision variable refers to the actual charging current instead of the EVSE current setpoint. However, the model is designed to only control the EVSE and not the OBC directly. As seen on Figure 3.2 there is a discrepancy between EVSE's PWM and charging current, which is referred as accuracy [7, 8]. Thus, a correction to the EVSE setpoint must be applied to compensate for this offset especially upon entering hardware test. The calculation will be elaborated in Chapter 6.

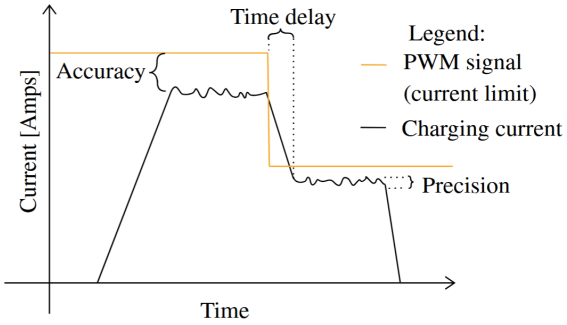

Figure 3.2: Illustration of EV charging KPI adapted from [7], as reproduced in [8].

Charging Current Limits

IEC 61851-1 specifies that a charger either delivers at least $I^{\min} = 6\,\mathrm{A}$ or is fully off. Partial charging below the minimum is not permitted [39]. $I^{\max}$ is the EVSE-rated maximum current and is capped at 32A.

$$I^{\max} s_{c,k} \geq I_{c,k} \geq I^{\min} s_{c,k} \tag{3.5}$$

Phase Assignment

The algorithm only simulate single-phase charging and each EV charging on phase ϕ means the current $I_{c,k}$ flows on that phase only. Therefore, the phase assignment binary must satisfy:

$$\sum_{\phi \in \mathcal{F}} z_{c,\phi,k} = s_{c,k} \qquad \forall\, c \in \mathcal{C},\ k \in \mathcal{K} \tag{3.6}$$

Battery Energy Dynamics

Battery energy evolves linearly with charging current. At the first active interval:

$$E_{c,k_c^{\mathrm{arr}}} = E_c^{\mathrm{init}} + V_{\mathrm{LN}} \cdot \eta_c \cdot \Delta t_1 \cdot 10^{-3} \cdot I_{c,k_c^{\mathrm{arr}}} \tag{3.7}$$

and for subsequent intervals within the parking window:

$$E_{c,k} = E_{c,k-1} + V_{\mathrm{LN}} \cdot \eta_c \cdot \Delta t_1 \cdot 10^{-3} \cdot I_{c,k} \qquad \forall\, k > k_c^{\mathrm{arr}} \tag{3.8}$$

where $V_{\mathrm{LN}} = 230\,\mathrm{V}$ is the nominal phase-to-neutral voltage and η_c is the charger efficiency. Battery energy is bounded by:

$$0 \leq E_{c,k} \leq B_c^{\max} \qquad \forall\, c,\ k \tag{3.9}$$

Departure SOC Constraint

Both literature and real EV have a wide range of design considerations including in determining the SOC cut-off. Constant-current-constant-voltage (CC-CV) method is the dominant strategy for AC charging to determine the actual charging current and the maximum SOC [40]. Due to the non-linear current decrease in the CV region, as seen on Figure 3.3, charging usually stops as on the AC side the current hits the minimum charging earlier defined in (3.5). Thus, charging process usually stops before hitting 100% using this strategy without direct control or command to stop charging with varying

ceiling between EVs. To simplify the algorithm, (3.11) assumed 90% as upper limit of SOC in place of CV charging behavior. Meanwhile, the lower limit is kept as a soft penalty which feed the objective function. ξ_c refers to the unmet energy upon departure. The penalty targets $B_c^{\text{dep}} = B_c^{\text{arr}} + \Delta B_c^{\text{req}}$, the requested energy from the session data. Since EVs may arrive and depart at times that do not align exactly with interval boundaries, the constraint is applied at $k^{dep} - 1$, the last complete interval during which the EV is present.

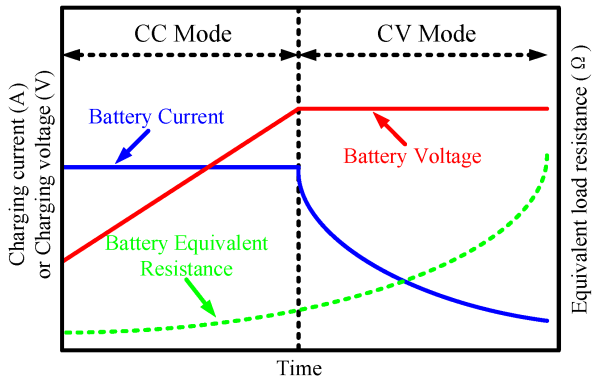

Figure 3.3: CC/CV charging profile [9]

The shortfall variable ξ_c is not computed explicitly but is instead determined by (3.10) acting together with the objective. Since $\xi_c \geq 0$ enters the objective with a positive weight γ, the optimizer drives it to the smallest value satisfying the constraint, so at the optimum it settles at $\xi_c = \max\left(0,\ B_c^{\text{dep}} - E_{c,k_c^{\text{dep}}-1}\right)$. It is therefore zero whenever the requested energy is delivered and equal to the missing energy otherwise, which is what makes the departure target a soft rather than a hard requirement.

$$E_{c,k_c^{\text{dep}}-1} + \xi_c = B_c^{\text{dep}} \qquad \forall\, c \in \mathcal{C} \tag{3.10}$$

$$E_{c,k_c^{\text{dep}}-1} \leq \text{SOC}_{\text{dep}}^{\max} \cdot B_c^{\max} \qquad \forall\, c \in \mathcal{C} \tag{3.11}$$

3.2.2. Installation Constraints

The constraints in this section bound the aggregate EV load at the installation level, protecting the electrical infrastructure and linking phase assignment to the grid current variables used in the objective function. The installation layout is detailed on Section 4.2, branching into 3 branches with three to four phase-changing AC chargers.

Per-Phase Aggregate and Grid Current

The total EV current on phase ϕ at interval k is:

$$I_{\phi,k}^{\text{ph}} = \sum_{c \in \mathcal{C}} I_{c,k} \cdot z_{c,\phi,k} \qquad \forall\, \phi \in \mathcal{F},\ k \in \mathcal{K} \tag{3.12}$$

Installation Capacity Constraint

Because the EV charger is sharing an installation with PV and the building installation, the sum of EV current is not bounded by the installation capacity (I_{cap}), but rather the headroom as expressed on (3.13) and (3.14):

$$\sum_{c \in \mathcal{C}} I_{c,k} \cdot z_{c,\phi,k} \leq I_{\phi,k}^{\text{head}} \qquad \forall\, \phi \in \mathcal{F},\ k \in \mathcal{K} \tag{3.13}$$

$$I_{\phi,k}^{\text{head}} = I_{\phi}^{\text{cap}} - I_{\phi,k}^{\text{load}} + I_{\phi,k}^{\text{PV}} \tag{3.14}$$

Those equations ensure that the total load on each phase never exceeds the main switch rating. It is assumed that this capacity is purely on the coupling of installation and the grid and not the current capability of the installation itself. Thus, the headroom may exceed I^{cap} when PV generation is present, consistent with the grid current definition in (3.4).

Branch Fuse Constraint

The installation is divided into branches b, each protected by a branch fuse I_b^{branch}. For each branch and phase:

$$\sum_{c \in \mathcal{C}_b} I_{c,k} \cdot z_{c,\phi,k} \leq I_b^{\text{branch}} \qquad \forall\, b,\ \phi \in \mathcal{F},\ k \in \mathcal{K} \tag{3.15}$$

where $\mathcal{C}_b \subseteq \mathcal{C}$ is the subset of EVSEs on branch b. This prevents a sub-circuit from being overloaded even when the main fuse has remaining headroom.

To summarize, the outer loop minimizes the weighted sum of charging cost, phase imbalance and unmet energy given in (3.2), subject to the charging current limits (3.5), the phase assignment rule (3.6), the battery energy dynamics (3.7)–(3.9), the departure state-of-charge conditions (3.10)–(3.11), and the installation limits (3.13)–(3.15). The decision variables are the charging current $I_{c,k}$, the on/off binary $s_{c,k}$, the phase-assignment binary $z_{c,\phi,k}$, the battery energy $E_{c,k}$, and the shortfall ξ_c. Only the first interval of each solution is implemented; the remaining intervals are discarded and the problem is re-solved at the next long-period index over the shortened horizon, initialized from the measured battery energy.

3.3. Inner Loop

The inner loop is an LP that acts as a fast corrective layer under the outer loop that aims to stabilize voltage by minimizing power fluctuations. The LP is chosen because of the ability to coordinate chargers on different states and bounds as opposed to a droop which acts in a decentralized manner. When solar supply or installation load fluctuates, the inner loop responds by adjusting available charging load. Only EVs whose measured battery energy is below a participation threshold of 75% of capacity at the start of the window, denoted $\mathcal{P}_k \subseteq \mathcal{C}$, take part in this modulation. Vehicles above the threshold simply execute their outer setpoints, which prevents vehicles close to their energy ceiling from being interrupted by the modulation near the end of their session. The LP works for a single step rather than a receding horizon in the outer loop because single step is adequate to suppress the aggregate difference. Adding forecast is not expected add meaningful improvement towards the result and will unnecessarily increase computational load.

The magnitude of corrections are coupled across three phases. This is important because the participation room can be different between the phases which risk unbalance at the expense of voltage correction. Furthermore, the PV system is a three-phase balanced inverter, so a solar fluctuation disturbs all phases equally. The inner loop is therefore formulated as a single optimization over all three phases simultaneously, coupled through one shared decision variable: the common-mode shift Δ_s, the amount in amperes by which every phase's aggregate EV current deviates from its outer plan. The objective consists of two components. The first component $C_{smoothing}$ closes the gap of load adjustment(Δ_s) and requested adjustment(Δ_s^{req}). The second component $C_{participation}$ acts as a proportional plan-anchor that keeps vehicle's current as close as possible to its outer setpoint, which then distributes the shift to the participating EVs. ϵ is set at a low value of 0.001 to avoid trading with the first objective. The objective function then becomes:

$$\min \quad \underbrace{\left|\Delta_s - \Delta_s^{\text{req}}\right|}_{C_{smoothing}} + \underbrace{\varepsilon \sum_{c \in \mathcal{P}_k} \left|i_{c,s} - I_{c,k}^{\text{set}}\right|}_{C_{participation}} \tag{3.16}$$

subject to the common-mode constraint

$$\sum_{c \in \mathcal{P}_k :\, \phi_c = \phi} i_{c,s} \;-\; \Delta_s \;=\; \sum_{c \in \mathcal{P}_k :\, \phi_c = \phi} I_{c,k}^{\text{set}} \qquad \forall\, \phi \in \mathcal{F} \tag{3.17}$$

where $i_{c,s}$ is the charging current of EV c at step s and $I_{c,k}^{\text{set}}$ is its setpoint from the most recent outer solution.

Δ_s^{req} is defined in (3.18) as the difference between the 20 second instantaneous currents and the average which is used in the outer loop.

$$\Delta_s^{\text{req}} = \frac{1}{3} \sum_{\phi \in \mathcal{F}} f_\delta\left(I_{\phi,s}^{\text{PV}} - \bar{I}_{\phi,k}^{\text{PV}}\right) \tag{3.18}$$

where $\bar{I}^{\mathrm{PV}}_{\phi,k}$ is the 15-minute averages underlying the outer plan. f_δ is a dead-zone function to limit the responsiveness of the inner loop to 2A deviation and above. The dead zone dampens the load adjustment preventing overreacting towards small voltage ramps with the trade-off of lower voltage ramp suppression.

A set of constraints are formulated following Section 3.3. The decision variables are the charging currents $i_{c,s}$ of the participating vehicles $c \in \mathcal{P}_k$ together with the shared shift Δ_s. Non-participating vehicles execute their outer setpoint directly, clipped to its remaining SOC ceiling, and its current is subtracted from the capacity limits available to the LP as described below. Phase shifting is disabled in the inner loop.

Each vehicle has a per-step upper ceiling $I^{\mathrm{hi}}_{c,s}$, defined as the tighter of the EVSE hardware limit and the remaining SOC headroom at step s. Its purpose is to throttle the charger near the energy ceiling, preventing the modulation from charging a vehicle past the maximum SOC of (3.11):

$$I^{\mathrm{hi}}_{c,s} = \min\left(I^{\max}_c,\ \frac{\mathrm{SOC}^{\max}_{\mathrm{dep}} \cdot B^{\max}_c - E_{c,s}}{\eta_c \cdot V_{\mathrm{LN}} \cdot \Delta t_2 \cdot 10^{-3}}\right) \tag{3.19}$$

where $E_{c,s}$ is the measured battery energy at step s and $\Delta t_2 = 20\,\mathrm{s}$ is the inner loop interval. The charging current of every participating vehicle is then bounded by:

$$I^{\min} \le i_{c,s} \le I^{\mathrm{hi}}_{c,s} \qquad \forall\, c \in \mathcal{P}_k \tag{3.20}$$

The lower bound enforces the IEC 61851 minimum current rule of (3.5) directly: a vehicle that the outer plan has switched on is never modulated below $6\,\mathrm{A}$, so the inner loop cannot toggle chargers on and off within a window. In the rare case that $I^{\mathrm{hi}}_{c,s} < I^{\min}$, the bounds would be inconsistent and the vehicle is instead fixed to zero for the step. The participation threshold of Section 3.3 makes this rare, since a vehicle below 75% SOC at the window start cannot reach the 90% ceiling within a single 15-minute window.

Phase headroom is limited by (3.21) at 20-second resolution, using the instantaneous PV and load measurements rather than the window averages.

$$\sum_{c \in \mathcal{P}_k :\, \phi_c = \phi} i_{c,s} \le I^{\mathrm{head}}_{\phi,s} - I^{\mathrm{fix}}_{\phi,s} \qquad \forall\, \phi \in \mathcal{F} \tag{3.21}$$

where $I^{\mathrm{fix}}_{\phi,s}$ is the total current of the non-participating vehicles on phase ϕ, which occupy part of the headroom but are not adjustable by the LP.

Similarly, the branch headroom is dictated by (3.22) at 20-second resolution, likewise reduced by the non-participating current on the branch.

$$\sum_{c \in \mathcal{P}_k \cap \mathcal{C}_b} i_{c,s} \le I^{\mathrm{branch}}_b - I^{\mathrm{fix}}_{b,s} \qquad \forall\, b \tag{3.22}$$

Because the headroom and fuse limits enter the same optimization that selects Δ_s, the shift is chosen small enough that every phase can deliver it simultaneously. A limit binding on one phase reduces the correction on all three rather than being violated or absorbed unevenly. Should a sudden 20-second load excursion nevertheless leave the constraint set with no feasible common shift, the dispatch for that step falls back to executing the outer setpoints under proportional capacity scaling, identical to the capacity-only inner loop of the non-smoothing scenarios.

3.4. Three-Phase Load Flow Model

3.4.1. Network Representation

The LV distribution network is modelled as a three-phase radial feeder in the phase frame. Let $\mathcal{N}$ denote the set of buses and $\mathcal{L}$ the set of branch segments. Each branch $ik \in \mathcal{L}$ connecting upstream bus i to downstream bus k is characterised by a symmetric 3×3 complex impedance matrix $\mathbf{Z}_{ik}$ in the phase frame $\mathcal{F} = \{1, 2, 3\}$:

$$\mathbf{Z}_{ik} = \begin{bmatrix} Z_s & Z_m & Z_m \\ Z_m & Z_s & Z_m \\ Z_m & Z_m & Z_s \end{bmatrix}_{ik}, \tag{3.23}$$

where Z_s is the phase-conductor self-impedance and Z_m is the inter-phase mutual impedance. The matrix is symmetric with equal off-diagonal entries because the cable cross-sections are identical for all three phase conductors. The entries of (3.23) are derived based on modified Carson's method using kron's reduction[41, 42].

$$Z_s = \frac{Z_0 + 2Z_1}{3}, \tag{3.24}$$

$$Z_m = \frac{Z_0 - Z_1}{3}. \tag{3.25}$$

Equations (3.24)–(3.25) follow directly from inverting the symmetrical-component transformation $\mathbf{Z}_{012} = \mathbf{A}^{-1}\mathbf{Z}_{abc}\mathbf{A}$, where $\mathbf{A}$ is the standard Fortescue transformation matrix. The self-impedance Z_s accounts for the phase conductor resistance and self-inductance together with the neutral return contribution, while the mutual impedance Z_m represents the coupling between phase conductors through the common neutral. The per-branch matrix $\mathbf{Z}_{ik}$ is obtained by multiplying Z_s and Z_m by the branch length ℓ_{ik} and normalizing to the per-unit system base $Z_{\text{base}} = V_{\text{base}}^2/S_{\text{base}}$.

Loads are modelled as constant-power injections at each bus n and phase ϕ: $P_{n,\phi}$ (active) and $Q_{n,\phi}$ (reactive), with positive values denoting consumption. PV generation enters as a negative active-power injection on the connected phase, consistent with the net-injection convention. EV chargers add load on the phase determined by the phase assignment $z_{c,\phi,k}$ from the MILP optimizer. The slack bus is fixed at the secondary terminal of the distribution transformer at nominal voltage:

$$V_{0,\phi} = V_0 \angle \theta_\phi, \quad \phi \in \mathcal{F}, \tag{3.26}$$

with $V_0 = 1.0\,\text{pu}$, $\theta_1 = 0°$, $\theta_2 = -120°$, $\theta_3 = +120°$.

3.4.2. Forward-Backward Sweep Algorithm

The load flow is solved using the Forward-Backward Sweep (FBS). FBS is chosen over other techniques such as Newton-Raphson (NR) and the Fixed-Point (FP) alternatives because its significantly faster computation time. FBS achieves this by exploiting the radial nature of distribution grids. large radial networks and directly suitable for real-time execution in the RT simulation model upon hardware validation [34]. Let the complex bus voltage at bus n and phase ϕ be:

$$V_{n,\phi} = V_{n,\phi}^d + j\, V_{n,\phi}^q, \tag{3.27}$$

where superscripts d and q denote the real and imaginary parts (direct and quadrature components) in the rectangular frame. The algorithm iterates between two sweeps until convergence.

The first sweep is backward sweep using Kirchhoff Current Law (KCL). Starting from the leaf nodes and working toward the slack bus, the branch current $\mathbf{I}_{ik}$ of each branch ik is computed by accumulating the nodal injection currents of all downstream buses. At node k the local injection current is:

$$I_{k,\phi} = \left(\frac{P_{k,\phi} - j\, Q_{k,\phi}}{V_{k,\phi}^{(m)}} \right)^*, \tag{3.28}$$

where $V_{k,\phi}^{(m)}$ is the voltage from the previous iteration m (the linearization point). The branch current in the three-phase vector form is then given by KCL:

$$\mathbf{I}_{ik} = \mathbf{I}_k + \sum_{kl \in \mathcal{L}} \mathbf{I}_{kl}, \tag{3.29}$$

where the summation runs over all branches kl that originate at k (i.e. all children of k).

The second sweep, forward sweep, then follows using Kirchhoff Voltage Law (KVL). Starting from the slack bus and working toward the leaf nodes, the voltage at downstream bus k is updated by applying the voltage drop across branch ik [34, Eq. 15]:

$$\mathbf{V}_k = \mathbf{V}_i - \mathbf{Z}_{ik}\, \mathbf{I}_{ik}, \tag{3.30}$$

or in the rectangular real/imaginary decomposition [34, Eqs. 16a–16b]:

$$V_{k,\phi}^{d} = V_{i,\phi}^{d} - \sum_{\phi' \in \mathcal{F}} \left(R_{\phi\phi'}\, I_{ik,\phi'}^{d} - X_{\phi\phi'}\, I_{ik,\phi'}^{q} \right), \tag{3.31}$$

$$V_{k,\phi}^{q} = V_{i,\phi}^{q} - \sum_{\phi' \in \mathcal{F}} \left(R_{\phi\phi'}\, I_{ik,\phi'}^{q} + X_{\phi\phi'}\, I_{ik,\phi'}^{d} \right), \tag{3.32}$$

where $R_{\phi\phi'} = \mathrm{Re}\{Z_{ik,\phi\phi'}\}$ and $X_{\phi\phi'} = \mathrm{Im}\{Z_{ik,\phi\phi'}\}$ are the resistance and reactance elements of $\mathbf{Z}_{ik}$, and $I_{ik,\phi'}^{d}$, $I_{ik,\phi'}^{q}$ are the real and imaginary parts of the branch current on phase ϕ'.

After each forward sweep the slack bus voltage is re-enforced to (3.26). The algorithm converges when the maximum voltage change across all buses and phases falls below the tolerance $\varepsilon_{\mathrm{FBS}}$:

$$\max_{n \in \mathcal{N},\, \phi \in \mathcal{F}} \left(\left| V_{n,\phi}^{d,(m+1)} - V_{n,\phi}^{d,(m)} \right| + \left| V_{n,\phi}^{q,(m+1)} - V_{n,\phi}^{q,(m)} \right| \right) < \varepsilon_{\mathrm{FBS}}, \tag{3.33}$$

with $\varepsilon_{\mathrm{FBS}} = 10^{-6}\,\mathrm{pu}$.

3.4.3. Voltage Unbalance Factor Computation

Once the FBS has converged, the VUF is computed post-hoc at each bus n from the converged phasor voltages $V_{n,1}$, $V_{n,2}$, $V_{n,3}$ using the Fortescue transformation:

$$V_{n}^{+} = \tfrac{1}{3}\left(V_{n,1} + \alpha\, V_{n,2} + \alpha^{2} V_{n,3} \right), \tag{3.34}$$

$$V_{n}^{-} = \tfrac{1}{3}\left(V_{n,1} + \alpha^{2} V_{n,2} + \alpha\, V_{n,3} \right), \tag{3.35}$$

where $\alpha = e^{j2\pi/3}$ is the Fortescue rotation operator. The VUF at bus n is then:

$$\mathrm{VUF}_{n}\ [\%] = \frac{|V_{n}^{-}|}{|V_{n}^{+}|} \times 100, \tag{3.36}$$

per the IEC definition adopted in this thesis. The VUF limit of $2\,\%$ set by EN 50160 is evaluated against the worst-case bus in each simulation scenario [16].

4

Simulation Setup and Case Studies

The fourth chapter details the metrics, the case studies, and the setup of the experiment. Section 4.1 provides metrics used to analyze the result of the optimization and load flow validation. Section 4.2 and Section 4.3 present the model/representation of the public charging facility installation and grid layout respectively. Subsequently, Section 4.4 provides the time series data for the installation and the grid. Section 4.5 assembles scenarios that combines all the data provided and combinations of objectives to analyze effectiveness optimizer. Furthermore, Section 4.6 introduces sensitivity analysis and objective weighting to the scenarios, testing the optimizer in different conditions. It is important to note that the case study is illustrative rather than representative. It demonstrates the mechanism and the direction of the effects, while the magnitudes are specific to the modeled network.

4.1. Performance Metrics

The following metrics are computed from the optimizer and load flow output. They are chosen to capture performance along three axes: power quality, energy delivery, and economic cost.

4.1.1. Load Imbalance

Load imbalance ($C_{imbalance}$) from the objective (3.2) is the main performance matrix whose correlation with VUF is being studied. It is evaluated at 20-second resolution and creating instantaneous imbalance value ($C_{\text{imbalance},s}$) as expressed by(4.1)

$$C_{\text{imbalance},k,s} = \frac{1}{3} \sum_{\phi \in \mathcal{F}} \left| I^{\text{grid}}_{\phi,k,s} - \bar{I}^{\text{grid}}_{k,s} \right| \qquad \forall\, k \in \mathcal{K},\ s \in S \tag{4.1}$$

where

$$\bar{I}^{\text{grid}}_{k,s} = \frac{1}{3} \sum_{\phi \in \mathcal{F}} I^{\text{grid}}_{\phi,k,s} \qquad \forall\, k \in \mathcal{K},\ s \in \mathcal{S} \tag{4.2}$$

I^{grid} is defined in (4.3) as the short-period counterpart of (3.4):

$$I^{\text{grid}}_{\phi,k,s} = I^{\text{load}}_{\phi,k,s} + I^{\text{ph}}_{\phi,k,s} - I^{\text{PV}}_{\phi,k,s} \qquad \forall\, \phi \in \mathcal{F}, k \in \mathcal{K},\ s \in \mathcal{S} \tag{4.3}$$

The horizon average and the worst-case instantaneous value are then expressed by (4.4):

$$\overline{C}_{\text{imbalance}} = \frac{1}{K \cdot S} \sum_{k \in \mathcal{K}} \sum_{s \in \mathcal{S}} C_{\text{imbalance},k,s}, \qquad C^{\max}_{\text{imbalance}} = \max_{k \in \mathcal{K},\ s \in \mathcal{S}} C_{\text{imbalance},k,s} \tag{4.4}$$

4.1.2. VUF

The VUF, as defined in (3.36), is the primary power quality metric evaluated at the grid level. It is computed at every bus $n \in \mathcal{N}$ from the three-phase voltages returned by the FBS load flow at each

short-period step (k, s), and summarized per bus by (4.5). However, only VUF at the tested bus will be presented

$$\overline{\mathrm{VUF}}_n = \frac{1}{K \cdot S} \sum_{k \in \mathcal{K}} \sum_{s \in \mathcal{S}} \mathrm{VUF}_{n,k,s}, \qquad \mathrm{VUF}_n^{\max} = \max_{k \in \mathcal{K},\, s \in \mathcal{S}} \mathrm{VUF}_{n,k,s} \tag{4.5}$$

4.1.3. Per-Phase Voltage Magnitude

Per-phase voltage magnitudes $|V_a|, |V_b|, |V_c|$ at the end nodes of the feeder are reported to complement the VUF metric. While VUF captures the global degree of asymmetry through sequence components, per-phase magnitudes reveal which phases experience overvoltage or undervoltage and to what extent. The operational band of $[0.90\, U_{\mathrm{nom}},\ 1.10\, U_{\mathrm{nom}}]$ is used as the reference for compliance assessment. Phase voltages are computed from the FBS load flow at each outer timestep. (4.6) then computes the highest and lowest voltage during the simulation. The metric is reported as the per-phase voltage profile over the 24-hour horizon.

$$V_{n,\phi}^{\min} = \min_{k \in \mathcal{K},\, s \in \mathcal{S}} \left|V_{n,\phi,k,s}\right|, \qquad V_{n,\phi}^{\max} = \max_{k \in \mathcal{K},\, s \in \mathcal{S}} \left|V_{n,\phi,k,s}\right| \tag{4.6}$$

4.1.4. EV Energy Delivery and SOC Fulfillment

SOC fulfillment is an important metric to analyze, to make sure that the load balancing does not interfere with the the one thing that is important of EV users, that their EVs are fully charged upon departure. Total unmet energy $\sum_{n \in \mathcal{C}} \xi_n$ is reported per scenario. A shortfall of zero across all EVs indicates full demand satisfaction, whereas any positive value represents energy not delivered within the parking window. This is the most important metric for the end user that are the EV owners which expects charging target to be filled upon departure.

$$E_c = \sum_{k \in \mathcal{K}} \sum_{s \in \mathcal{S}} V_{\mathrm{LN}} \cdot \eta_c \cdot i_{c,k,s} \cdot \Delta t_2 \cdot 10^{-3} \quad [\mathrm{kWh}] \tag{4.7}$$

The shortfall at departure is:

$$\xi_c = \max\left(0,\ B_c^{\mathrm{dep}} - E_{c,\, k_c^{\mathrm{dep}} - 1}\right) \quad [\mathrm{kWh}] \tag{4.8}$$

The fleet total and the per-vehicle relative shortfall are:

$$\Xi = \frac{\sum \xi_c}{\sum (B_c^{\mathrm{dep}} - B_c^{\mathrm{arr}})} \times 100 \quad [\%] \tag{4.9}$$

4.1.5. Charging Cost and Cost Premium

Total charging cost is computed as value of C_{cost} from the objective function (3.2) . Cost is normalized to €/kWh delivered to allow fair comparison across scenarios with different total energy import. The *cost premium* is defined as the percentage increase in cost per kWh relative to Scenario 2 (cost-only optimization):

$$\Delta C_{\%} = \frac{C_{\mathrm{Sc}x}/E_{\mathrm{Sc}x} - C_{\mathrm{Sc2}}/E_{\mathrm{Sc2}}}{C_{\mathrm{Sc2}}/E_{\mathrm{Sc2}}} \times 100\,\% \tag{4.10}$$

This metric directly quantifies the economic cost of introducing phase balancing, enabling DSOs to gauge incentive required to entice public charging operators to participate in VUF reduction efforts.

4.1.6. Voltage Ramps/Solar Flicker

Rapid voltage variations caused by fluctuating PV output and stepping charger setpoints are quantified through the voltage ramp between consecutive short-period steps, following the flicker assessment of [22]. Aside from damages to electrical appliances connected to the grid, this phenomenon causes light flicker that can cause annoyance and health problem to those who are exposed[43, 44]. Two severity thresholds from [22] are applied: ramps exceeding $1.6\,\mathrm{V}$ (0.7% of nominal) correspond to visible light flicker, and ramps exceeding $3.6\,\mathrm{V}$ (1.6% of nominal) to annoying flicker. The voltage ramp is evaluated

every second across a fixed 20-second window ($W = 20$ samples at the $1\,\mathrm{s}$ resolution of the load-flow data):

$$\Delta V_{n,\phi,t} = \Big|\,|V_{n,\phi,t}| - |V_{n,\phi,t-W}|\,\Big|, \qquad W = 20, \quad \forall\, n \in \mathcal{N},\ \phi \in \mathcal{F},\ t \in \{W+1, \dots, T\}. \tag{4.11}$$

The occurrence is then the share of the per-second evaluations whose ramp exceeds each threshold over 24-hour horizon. The flicker metric is reported at the public charging facility bus. These thresholds are perception-based indicators rather than regulatory limits: as noted in [22], EN 50160 flicker compliance cannot be assessed at a 20-second resolution, which requires sub-second measurement. The occurrence values are therefore interpreted as relative indicators.

4.2. Public Charging Facility Installation

The public charging facility represents the point of common coupling at which the EV chargers, the rooftop PV system, and a small management office are jointly connected to the LV grid through a single main fuse box, as illustrated in figure 4.1. This shared connection point means that the aggregate current drawn from or injected into the grid at this bus reflects the net of EV charging, office consumption, and PV generation. Consequently, the available current headroom for EV charging at each timestep is not fixed, but fluctuates with the office load and the instantaneous PV output.

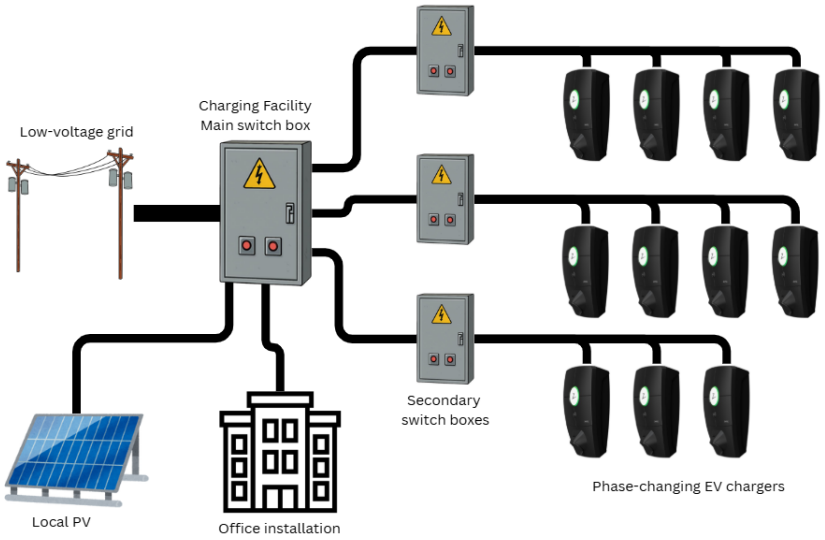

Figure 4.1: Illustration of public charging facility

The facility is connected to the grid at a single bus within the test network (see section 4.3) via one three-phase installation fuse rated at I_{cap}. Downstream of this fuse, the installation branches into three sub-circuits: the office building supply, the PV inverter, and the EV charging spots. Each EV charging spot is equipped with a smart charger capable of continuously modulating its current setpoint between $I_{\mathrm{min}} = 6\,\mathrm{A}$ (the IEC 61851 minimum) and I_{max}, depending on the on-board charger of the connected vehicle. Further detail of the installation can be found in Appendix A.1

4.3. Test Network

Given that the load-flow has been validated in IEEE networks, a test with real grid can be done instead of using standardized benchmark grid. In that case, a real grid is preferred to allow validation under a realistic operating condition, particularly with a Dutch LV network. Thus, a real anonymized low-voltage distribution grid from the Dutch DSO Stedin is used in this thesis. Particularly, a grid with no PV penetration is chosen, isolating the PV supply to the public charging facility because in the decentralized system the load adjustment can only respond locally. Thus, this isolation allows more direct observation of optimizer response towards PV supply fluctuations.

Table 4.1: Low-voltage network characteristics. Extended detail in Appendix A.2

Parameter	Value
Grid topology	Fully radial
Transformer rating	400 kVA
Transformer voltage level	10.75/0.42 kV
Slack bus location	LV terminal of transformer
Slack bus voltage	230 V (phase-to-neutral)
Number of LV nodes	213

The 214 load connections are a mix of single-phase and three-phase consumers: 52 are connected to L1 only, 61 to L2 only, 49 to L3 only, and 52 are three-phase connections. This asymmetric phase distribution is a direct source of structural voltage unbalance in the base case, independent of EV charging. Each connection is assigned a NEDU demand profile class (E1A, E1B, E1C, E2A, or E2B). The network data also contains the location, type, and phase of each load connection and the distributed PV generation, as illustrated in Figure 4.2. The subsequent section will then complete the simulation setup with the external input profiles.

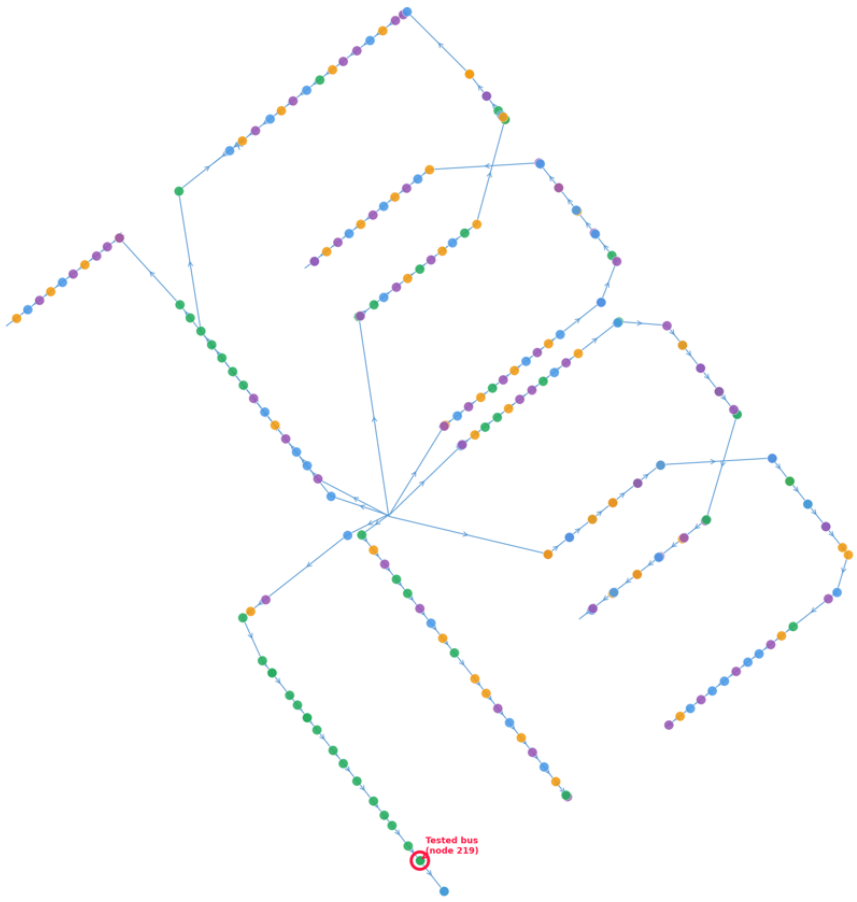

Figure 4.2: Radial low-voltage test grid (BU06370601, Stedin). Node colors indicate phase connection: blue = L1, purple = L2, orange = L3, green = three-phase. The transformer substation is shown at the root.

4.4. Input Data

Once the test network is set up, the time-series data in the form of load/generation profile is injected into the grid. The simulation uses three categories of external time-series input: distributed photovoltaic generation, residential base load, and day-ahead electricity prices. Data resolutions are divided into 15 minutes and 20 seconds matching the outer and inner optimization loops. The simulation is run on

a random summer weekday considering that VUF tends to be higher in the summer [11]. The following subsections describe the source and characteristics of each input.

4.4.1. PV Supply

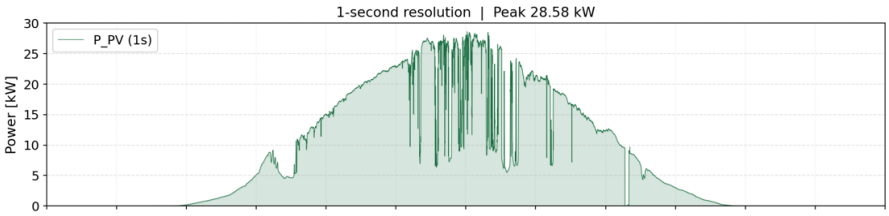

Figure 4.3: Power output profile of a three-phase 50 kWp rooftop PV system on a representative summer day.

This thesis replicates the PV generation model used in [11], which is developed based on [45] and [46]. A 50 kWp three-phase PV system is installed on the public charging facility, as shown in Figure 4.3. The output of the PV is assumed to be uncontrolled with grid following inverter which do not contribute to voltage regulation. The model computes instantaneous PV power output from weather inputs (irradiance, ambient temperature, wind speed) obtained from [47], which provides high resolution data up to 1 second. The data is computed thought (4.12)-(4.19)

The total module incident irradiance G^{inc} at a given angle of incidence (AOI) is determined via the Isotropic Sky Model:

$$\cos(AOI(t)) = \cos(\theta_M)\sin(\theta_s(t)) + \sin(\theta_M)\cos(\theta_s(t))\cos(Az_M - Az_s(t)) \tag{4.12}$$

where θ_M and Az_M are the module tilt and azimuth set at 30 and 180 degrees respectively. θ_s and Az_s are the solar elevation and azimuth angles. The dynamic module temperature T_M is computed using the Duffie–Beckman model:

$$T_M(t) = T_a(t) + \frac{G^{\text{inc}}(t)}{G_{\text{noct}}}(T_{noct} - 20)\frac{9.5}{5.7 + 3.8s_w}\left(1 - \frac{\eta^{\text{stc}}}{a_f}\right) \tag{4.13}$$

where T_a is ambient temperature, s_w wind speed, $G_{\text{noct}} = 800\,\text{W/m}^2$, $T_{\text{noct}} = 44°\text{C}$, and $a_f = 0.9$ the absorptivity factor. The temperature effect on MPP power and efficiency is then:

$$P_{\text{MPP}}^{T_M, G^{\text{stc}}}(t) = P_{\text{MPP}}^{\text{stc}} + c_{P_{\text{MPP}}}(T_M(t) - T_a) \tag{4.14}$$

$$\eta^{T_M, G^{\text{stc}}}(t) = \frac{P_{\text{MPP}}^{T_M, G^{\text{stc}}}(t)}{A_M G^{\text{stc}}} \tag{4.15}$$

The irradiance effect on MPP power and efficiency under actual incident irradiance is:

$$P_{\text{MPP}}^{25°\text{C}, G^{\text{inc}}}(t) = FF \cdot V_{\text{OC}}^{25°\text{C}, G^{\text{inc}}}(t) \cdot I_{\text{SC}}^{25°\text{C}, G^{\text{inc}}}(t) \tag{4.16}$$

$$\eta^{25°\text{C}, G^{\text{inc}}}(t) = \frac{P_{\text{MPP}}^{25°\text{C}, G^{\text{inc}}}(t)}{A_M G^{\text{inc}}} \tag{4.17}$$

Finally, the real module efficiency and PV output power are:

$$\eta^{PV}(t) = \eta^{25°\text{C}, G^{\text{inc}}}(t)\left[1 + \frac{c_\eta}{\eta^{\text{stc}}}(T_M(t) - 25°\text{C})\right] \tag{4.18}$$

$$P_{pv}(t) = \eta^{PV}(t)\, A_M\, G^{\text{inc}}\, \eta^{\text{inv}} \tag{4.19}$$

The model parameters are $G^{\text{stc}} = 1000\,\text{W/m}^2$, $P_{\text{MPP}}^{\text{stc}} = 245\,\text{W}$, $\eta^{\text{stc}} = 0.1943$, $A_M = 1.26\,\text{m}^2$, $c_{P_{\text{MPP}}} = -0.29\%/°\text{C}$, $\eta^{\text{inv}} = 0.95$, and $FF = 0.788$ [11].

4.4.2. Residential Load Profiles

Each load node in the grid is assigned an individual residential profile drawn from the same source database, scaled by the annual consumption class of the connection. This per-node assignment is derived directly from the grid data, which encodes each connection's consumer phase assignment and annual energy consumption. Subsequently, load profiles data from the NEDU standardized electricity consumption profiles ($P_d(t)$) published by the Dutch Platform Verbruiksprofielen [24] are multiplied with the annual consumption (E_i) creating the per-minute load profile ($P_l^i(t)$)

$$P_l(t) = E_i \cdot P_d(t) \tag{4.20}$$

Figure 4.4 illustrates the total load of the grid excluding the load from public charging facility that will be inserted into post optimization. An additional load is added to represent small EV charging facility management office. This building represents the administrative and control infrastructure for the chargers. As illustrated on Figure 4.1, the office, PV, and EV chargers are connected through on main fuse box. Thus, the metering is combined and the office load profile affects current headroom of the installation.

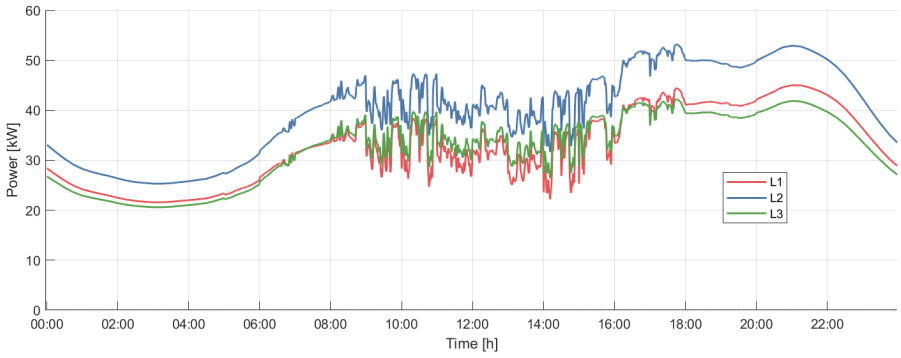

Figure 4.4: Total load of the simulated grid per phase without public charging facility load

4.4.3. EV Fleet Schedules and Parameters

The EV fleet parameters are derived from the ElaadNL open dataset [25], which provides real-world charging session records from Dutch public and semi-public charging infrastructure. The dataset contains session-level information including plug-in time (t^{arr}), plug-out time (t^{dep}), charged energy volume, charger type, and phase configuration.

For this study, the EV schedule is pre-processed from the ElaadNL data into the format consumed by the optimizer, with each EV record providing the following parameters:

- t_c^{arr}: arrival time of EV c
- t_c^{dep}: departure time of EV c
- B_c^{arr}: battery energy at arrival, expressed in kWh
- B_c^{max}: maximum usable battery capacity in kWh
- I_c^{max}: maximum charging current in amperes (16 A or 32 A), determined by the on-board charger
- $\phi_c^{1/3}$: charging phase configuration (single-phase or three-phase)
- ϕ_c^{SE}: the supply equipment phase to which the EV is connected

Some adjustments are done on the generated data from Elaad and information gaps are filled to assume perfect information that are required in the optimization. The original data has a strict charging duration, which does not differentiate end of charging session and EV departure. Thus, time departure extension is required to give room for optimization and allow load management. Because the charging location is a assumed as an office building or a third place charging is more active during the day, filling up all charging spots. The parking schedule is detailed in Appendix A

4.4.4. Electricity Price Profile

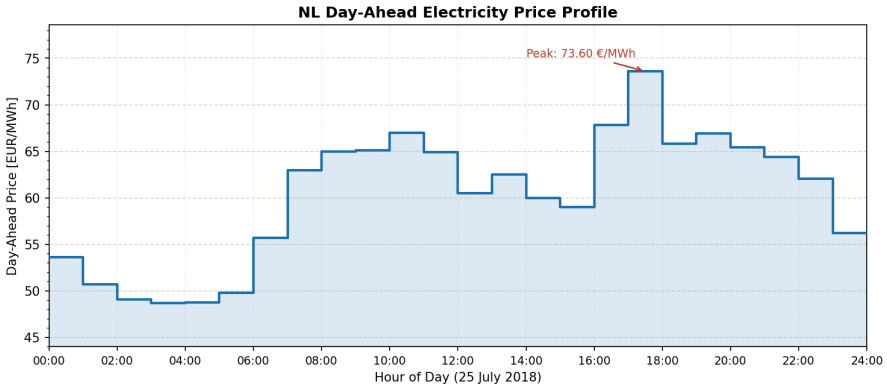

Figure 4.5: NL day-ahead electricity price profile for 24 July 2018, retrieved from the ENTSO-E Transparency Platform [10]. Prices are in EUR/MWh at hourly resolution.

The electricity price signal used in the cost objective of the optimizer is the Dutch day-ahead market price retrieved from the ENTSO-E Transparency Platform [10]. The profile shown in Figure 4.5 corresponds to 24 July 2018, a representative summer weekday, and exhibits the characteristic double-peak structure of European spot markets: a morning ramp as demand rises after 06:00, a midday dip coinciding with high solar generation, and an evening peak between 17:00 and 18:00 reaching EUR73.60/MWh before declining toward midnight. The overnight valley falls below EUR50/MWh, providing an economic incentive for demand shifting.

The price vector $\lambda(k)$ is provided as a fixed exogenous input to the outer MILP at each receding-horizon step, indexed to the same 15-minute intervals as the load and PV profiles. For intervals within a single hourly price block, the same price is applied to all four sub-intervals.

4.5. Scenario Design and Objective Weight Tuning

Different scenarios and are deployed to analyze the impact of the load balancing towards the defined performance metrics. The scenarios are arranged with adding layers of objective functions from uncontrolled scenario to one with all objective functions and loops. The scenarios are summarized in Table 4.2.

Scenario 1 — Uncontrolled charging (baseline). No optimization is applied. EVs begin charging at the maximum available current immediately upon arrival. However, a charging queue is enabled to prevent simultaneous overloading of the connection: when the aggregate current at a phase would exceed the installation's grid capacity I_{cap}, arriving EVs are held in queue and admitted in order of arrival. Phase switching is disabled and each parking spot has pre-assigned phase for the entire simulation. This scenario represents the baseline.

Scenario 2 — Cost minimization only. The outer MILP minimizes total charging cost (C_{cost}) subject to capacity and SOC constraints. The inner LP enforces only the per-phase capacity headroom constraint. The optimization allows charging load adjustments responding to price signal and availability of owned PV supply. This is the most user-focused scenario which becomes the base for estimation of cost premium that is incurred due to load imbalance optimization. Similar with scenario 1, phase-shift are not included in this scenario and the charger phases are pre-assigned.

Scenario 3 — Cost minimization with phase balancing. The outer MILP adds a linearized phase-imbalance term ($C_{imbalance}$) to the Scenario 2 objective, penalizing the deviation of per-phase current from the mean across phases. The chargers are now assumed as phase changing chargers The inner LP retains the capacity-only enforcement of Scenario 2. This scenario reveals the effectiveness of load balancing towards VUF and trade-off between cost minimization and phase imbalance, quantifying the cost premium.

Scenario 4 — Cost minimization, phase balancing, and smooth inner loop. Almost identical to Scenario 3 in the outer MILP, but also employ complete load-smoothing inner LP, that minimizes the total variation of per-phase EV current across 20-second steps within each 15-minute outer interval,

as described in Section 3.3. This targets the higher-frequency current fluctuations not captured by the outer optimizer, particularly relevant for reducing solar flicker and improving power quality at the charger terminals.

Table 4.2: Overview of the four simulation scenarios.

Scenario	Outer objective	Inner mode	Phase switching
Sc1	None (max current, queued)	Capacity clip	No
Sc2	Cost minimisation	Capacity clip	No
Sc3	Cost + phase imbalance	Capacity clip	Yes
Sc4	Cost + phase imbalance	Load smoothing	Yes

4.6. Objective Weighting Tuning and Sensitivity Analysis

The tuning examines how the performance metrics respond to variations in the two key objective coefficients: α (charging cost weight) and β (load imbalance weight). These are the primary tunable parameters in Scenarios 3 and 4, and their relative magnitude controls the trade-off between economic efficiency and load balance quality. Meanwhile,γ (unmet energy penalty) is kept at significantly higher number, 1000. This high empirical coefficient is in place as a soft constraint. Thus, tuning of this coefficient is less substantial.

The sweep covers the following parameter ranges, consistent with the ranges coded into the optimizer:

- $\alpha \in \{0.01, 0.1, 1, 10, 100\}$
- $\beta \in \{0.01, 0.1, 1, 10, 100\}$

For each (α, β) combination, both Sc3 and Sc4 are executed over the full 24-hour horizon, producing a grid 50 optimization runs in total. The outputs are summarized as heatmaps over the (α, β) grid for the key performance indicators: average charging cost per kWh, total unmet energy (SoC shortfall), average phase-current variance, and local PV utilization. This allows the sensitivity of each metric to the cost-balance trade-off to be visualized directly.

In addition to the objective weight sweep, the sensitivity of the results to the scale of the public charging facility itself is examined. The base case scenarios assume a fixed facility configuration. However, the value of solution achievable through optimization is expected to depend on the number of simultaneously connected EVs and the installed PV capacity, since these determine the amount of flexible current available to the outer MILP for phase rebalancing.

To test this, the facility configuration is varied along two independent axes:

- **EV fleet size**: the number of occupied charging spots is varied from 9 to 13 vehicles, spanning the range from partial to near-full utilization of the available parking spots.
- **PV capacity**: the installed PV system size is varied from 30 kWp to 70 kWp, representing a range from the current installation size (Figure 4.3) to a larger future deployment consistent with increasing rooftop PV adoption in the Netherlands.

For each combination of fleet size and PV capacity, all relevant scenarios are re-executed over the full 24-hour horizon under the same EV arrival/departure schedule, price profile, and objective weights ($\alpha = 1$, $\beta = 1$) used in the base case, isolating the effect of facility scale from the effect of objective tuning. The resulting grid of runs is summarized as heatmaps of the metrics over the (fleet size, PV capacity) plane. This allows an assessment of whether the balancing benefit of the optimizer scales with the flexibility available at the facility, and whether higher PV penetration increases or decreases the achievable VUF reduction under phase rebalancing.

5

Results and Discussion

The fifth chapter presents and discusses the outcome of the optimization and load flow across the scenarios assembled in Chapter 4. Section 5.1 establishes the two baselines, the uncontrolled charging (Scenario 1) and the cost-only optimization (Scenario 2), against which the balancing scenarios are measured. Section 5.2 then evaluates the two balancing scenarios, Scenario 3 and Scenario 4, metric by metric, covering load imbalance and VUF, per-phase voltage magnitude, voltage ramps, energy delivery, and charging cost. Subsequently, Section 5.3 relaxes the fixed objective weighting and examines the cost-balance trade-off as the weights are swept, delimiting the range over which the tuning is stable. Finally, Section 5.4 tests whether these findings persist as the operating conditions vary, sweeping the EV count and the installed PV capacity.

It is important to note that the case study is illustrative rather than representative, as established in 4. The scenarios are compared on a single network and day, so the results should be read for the mechanisms they expose and the direction of the primary effects, not as transferable magnitudes. A further distinction applies within the comparison itself. Some differences between scenarios are direct consequences of the control objective, while others are incidental correlations of this particular case that do not necessarily reproduce on other networks or load patterns. Whether an effect is coincidental or caused by the control action, it is identified in Section 5.2, and Section 5.4 serves to separate the effects that persist across different sizing from those that do not.

5.1. Baseline Scenarios

5.1.1. Scenario 1 : Uncontrolled Charging

Scenario 1 establishes the reference behavior of the charging facility with no control, only constrained by the current limits of branches and main switch. The result is that all plugged EVs are drawing maximum current allowed and the charging is compressed into a window of roughly seven hours, after which the last vehicle reaches its energy target and the facility falls silent for the remainder of the parking duration. Figure 5.1 shows the per-phase charging current over the 24-hour horizon. The most heavily loaded phase reaches 125 A surpassing installation limit of 80 A per phase with the help of local PV that expands the capacity of installation. In this scenario the average load imbalance sits at 5.27 A, with the worst instantaneous value of 23.44 A occurring when consecutive arrivals happen to stack on the same phase.

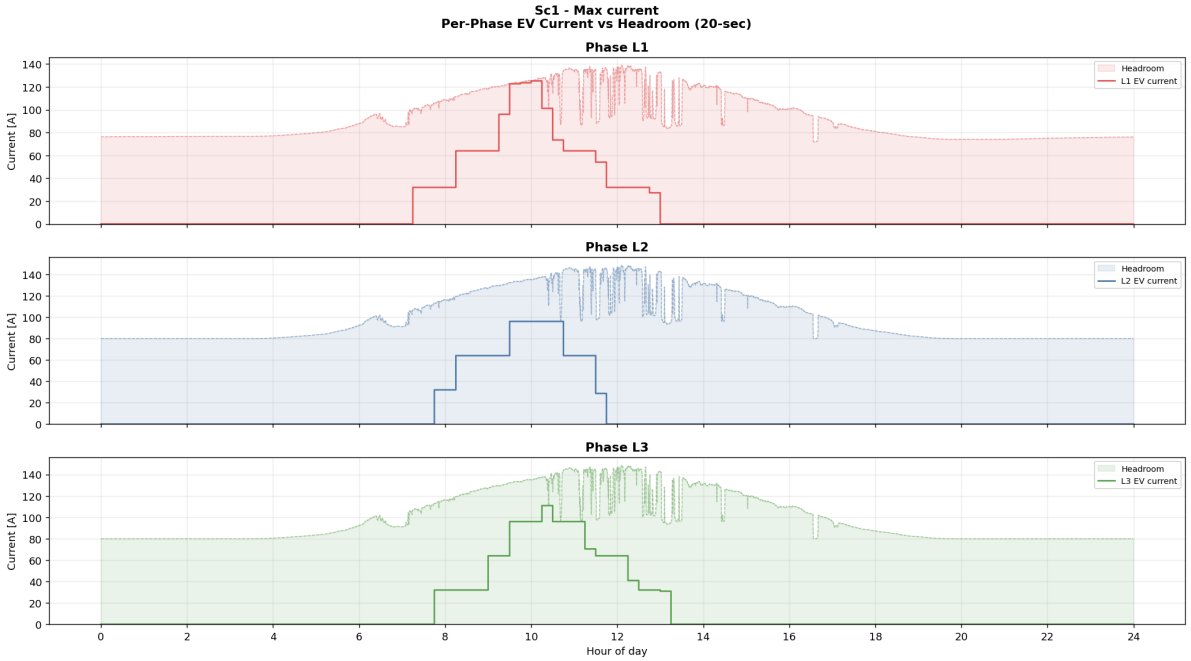

Figure 5.1: Per-phase current of scenario 1, indicative of uncontrolled charging behavior. The headroom is expanded during the day due to PV generation.

The EV charging imposes dominant influence towards the grid in this particular setup and the load imbalance reflects on the VUF. This is confirmed by the load flow result in figure 5.2b, which shows that the VUF is highly elevated during the charging window although it stays within regulatory limit of 2%. This is an expected behavior which is echoed by other studies which earlier forms the motivation of this thesis. Besides having no control, the current energy delivery is unfavorable for phase balance because each parking spot carries a preassigned phase and vehicles occupy spots in order of arrival. The worsening of load imbalance and VUF in this scenario clearly sets up the potential of improvement by the optimizer.

Besides VUF, the load flow result depicted on Figure 5.2 also shows that aggregated EV charging and PV supply worsen voltage magnitudes and solar flicker. The lowest voltage magnitude is recorded at 0.962 pu(figure 5.2a. Over the simulated day the bus shows a visible-flicker occurrence of 0.915% and never reaches the annoying-flicker threshold. This occurrence separates into two mechanisms by its timing (figure 5.3). 402 flicker events are recorded at the 15-minute schedule boundaries, where the charging setpoints step, and are attributed predominantly to EV charging transients. The remaining 1968 events fall within the scheduling windows, where the charging setpoint is held constant, and are attributed predominantly to the continuous variation of PV generation. The separation is approximate rather than exact. EV arrivals and departures occurring mid-window contribute some within-window transients, and the headroom-following of the charging current is itself PV-correlated, but the split still identifies the dominant source in each category. These two components frame the power-quality analysis in the remainder of this chapter. The within-window PV-driven component is the direct target of the inner loop introduced in Scenario 4, which modulates the charging current to absorb PV fluctuations. The boundary component originates from the discrete 15-minute re-planning of the outer loop and is not part of the balancing objective. Thus, the optimizer does not attempt to affect this number, but may still indirectly influence it. Comparison with the other scenario will later uncover this possibility.

5.1.2. Scenario 2 : Cost-Only Optimization

Scenario 2 replaces the uncontrolled dispatch with the outer MILP minimizing charging cost alone, while retaining the fixed phase assignments. The effect on the temporal charging profile is immediate and visible in Figure 5.4: rather than charging at maximum current on arrival, the optimizer spreads the same energy across the low-price windows and the local PV generation. The resulting charging cost of 0.80 cent/kWh, against 2.46 cent/kWh under Sc1, defines the economic reference point from which the cost premium of the balancing scenarios is measured in the subsequent sections.

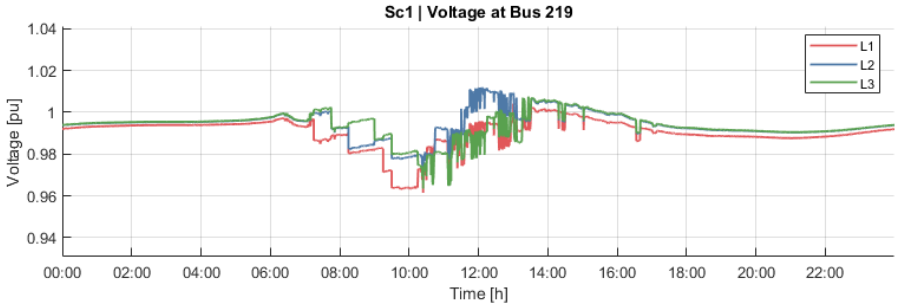

(a) Voltage profile at tested bus: Scenario 1

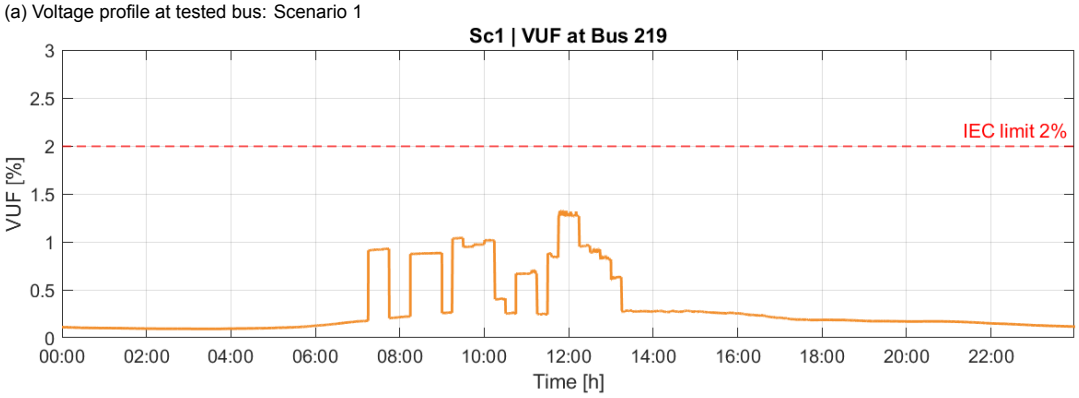

(b) VUF at tested bus: Scenario 1

Figure 5.2: Scenario 1 results: voltage profile and voltage unbalance factor.

This cost reduction, however, comes at the expense of phase balance. The full-horizon average load imbalance decreases from 5.27 A (Sc1) to 4.50 A, but this improvement is incidental. The metric that matters for power quality moves in the opposite direction as the worst-case instantaneous imbalance rises by 80.3%, from 23.44 A to 42.27 A, as cost minimization synchronizes the fleet into the same cheap windows and stacks current on whichever phases those vehicles occupy. It reflects the charging being spread over a longer window, which lowers the time-averaged deviation, rather than any deliberate balancing action.

The worsening of the load-balance is reflected at the grid level, resulting cost-only optimization drives the facility bus into a limit violation during the synchronized charging peak, as shown in Figure 5.5. The deepest voltage excursion at the facility bus drops slightly to 0.957 pu (against 0.962 pu under Sc1), remaining within the operational band. The VUF, however, tells the decisive story. Its horizon average at the facility bus decreases to 0.270% (from 0.311% under Sc1), mirroring the incidental improvement in average imbalance, but its maximum rises to 2.41%, against 1.32% under Sc1. This worst-case value exceeds the EN 50160 limit of 2%, whereas uncontrolled charging remained compliant at every

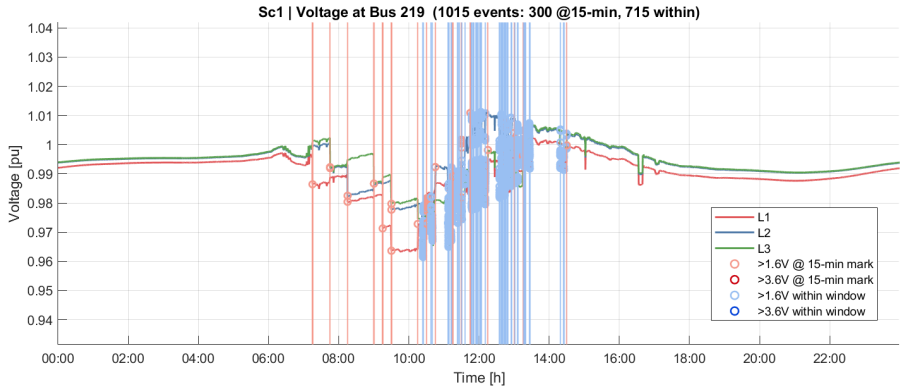

Figure 5.3: Example of detection of voltage flicker at boundaries(switching) and within it(solar flicker)

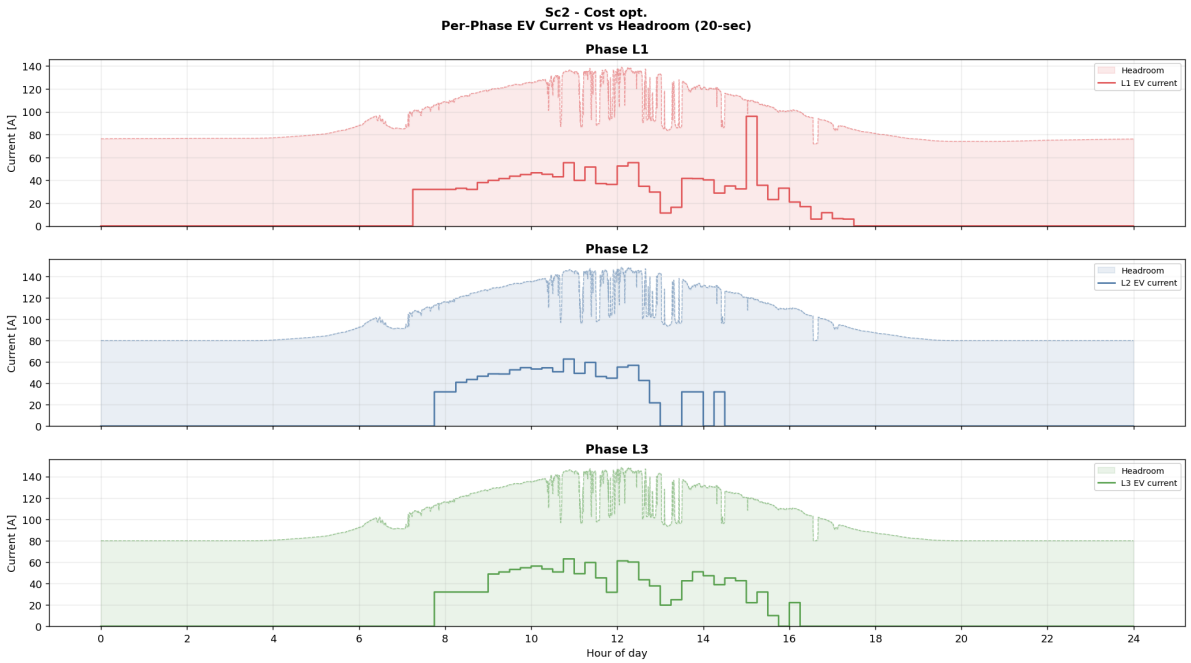

Figure 5.4: Per-phase current of Scenario 2, illustrating load concentration during low-price windows.

instant.

The cost synchronization introduces annoying flicker that uncontrolled charging did not produce as flicker picture worsens in a similar way . The visible-flicker occurrence at the facility bus is essentially unchanged at 0.900% (against 0.915% under Sc1), with 2333 ramp events above the 1.6 V/20 s threshold, of which 287 fall on the 15-minute schedule boundaries and 2046 within the windows, a distribution close to that of Sc1. The significant change is at the severe threshold, where 56 ramps now exceed the 3.6 V/20 s annoying-flicker level (0.022% of the time) whereas Sc1 registered none.

The comparison between the two baselines yields the motivating conclusion for the remainder of this chapter, cost-driven smart charging is not a phase-balancing measure, and it actively aggravates worst-case unbalance precisely because it synchronizes the fleet. The average metrics improved incidentally, while the tail of the VUF distribution, which is what the EN 50160 assessment and the DSO ultimately care about, crosses the 2% limit, and the flicker acquires annoying-level events that were previously absent. Mitigating unbalance therefore requires the phase dimension to enter the objective function explicitly, which is the subject of Scenario 3.

5.2. VUF Optimized Results

5.2.1. Load Balance and VUF

The introduction of the balance term ($C_{imbalance}$) improves both load and voltage balance, as observed on Figure 5.6. Table 5.1 summarize the maximum and average load imbalance for 24-hour simulation. Over the 24-hour horizon the average imbalance of the net phase currents from their three-phase mean falls from 5.27 A (Sc1) and 4.50 A (Sc2) to 1.23 A under both Sc3 and Sc4, and the worst instantaneous deviation collapses from 23.44 A and 42.27 A in the baselines to 2.92 A under Sc3 and 2.93 A under Sc4. The imbalances in Sc3 and Sc4 only happens outside of charging sessions while the load imbalance during it is near zero because the single-phase office load can only be compensated while EVs are connected. Due to the near zero load imbalance in Sc3, Sc4's benefit is confined to the voltage ramp metrics of Section 5.2.5. However, this result also show that the inner loop does not cause instability and provide a result that is at least on par with Sc3.

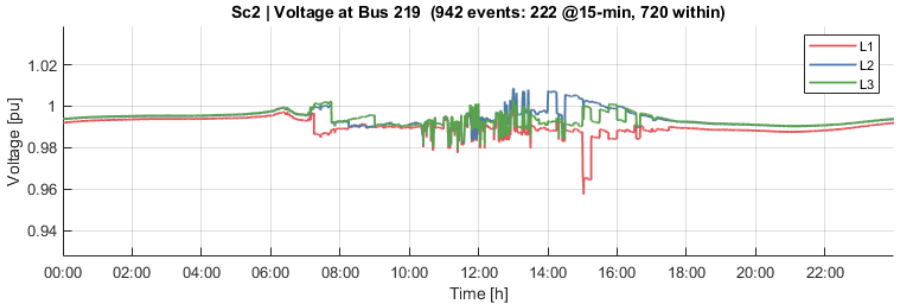

(a) Voltage profile at tested bus: Scenario 2

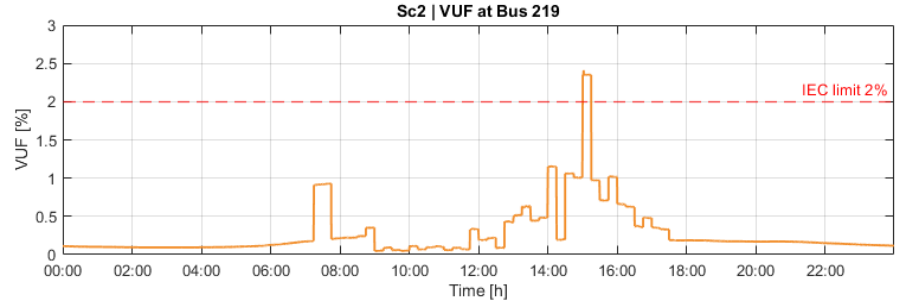

(b) VUF at tested bus: Scenario 2

Figure 5.5: Scenario 2 results: voltage profile and voltage unbalance factor.

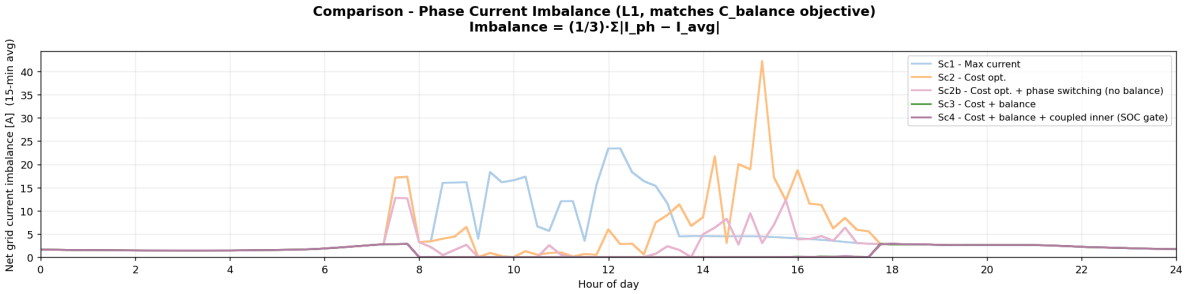

Figure 5.6: Load balance comparison of 4 scenarios. Scenario 3 and 4 trace almost entirely similar balance during the simulation.

Table 5.1: Load imbalance and facility-bus VUF across scenarios. The Sc2 maximum VUF exceeds the EN 50160 limit of 2%.

Metric	Sc1	Sc2	Sc3	Sc4
Average load imbalance [A]	5.27	4.50	1.23	1.23
Maximum load imbalance [A]	23.44	42.27	2.92	2.93
Average VUF [%]	0.311	0.270	0.100	0.100
Maximum VUF [%]	1.32	2.41	0.187	0.188

At the grid level, the improvement affects the VUF at the facility bus proportional to the load imbalance improvement. The mean VUF over the 24-hour horizon drops from 0.311% (Sc1) and 0.270% (Sc2) to 0.100% under both Sc3 and Sc4, and the maximum from 1.32% and 2.41% to 0.187% and 0.188% respectively, landing below even the baselines' average values. As seen on Figure 5.7, the daily VUF profile inverts its shape relative to the baselines: whereas in Sc1 and Sc2 the charging window is the period of elevated unbalance, in Sc3 and Sc4 it is the period of lowest unbalance at the facility. This pattern is similar with the load balance profile, which confirms that EV charger load balance correlate directly with VUF especially on a grid or branch with high EV participation.

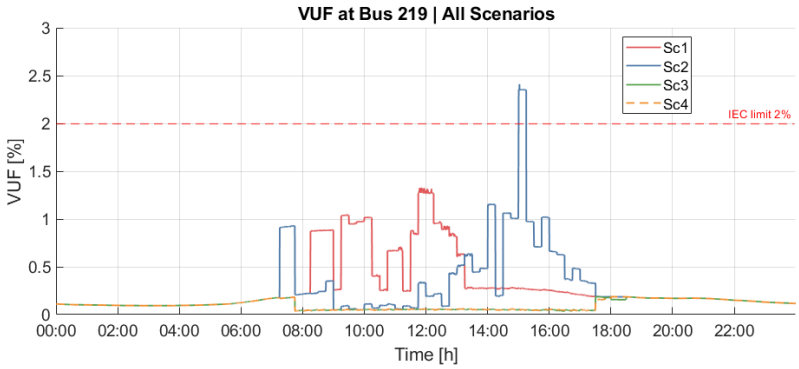

Figure 5.7: VUF at tested bus of all scenarios. Sc3 and Sc4 track very similar profile, improving VUF during charging window

5.2.2. Per-Phase Voltage Magnitude

All scenarios remain comfortably within the $[0.90, 1.10]$ pu operational band, so voltage magnitude is a matter of margin rather than compliance in this case study. Besides, the objective function only aim to balance load and not suppress it. Thus, voltage magnitude correction is not expected as direct consequence. Nevertheless, the deepest excursion at the facility bus is 0.962 pu under Sc1, 0.957 pu under Sc2, 0.977 pu under Sc3, and 0.980 pu under Sc4. Figure of voltage profile comparison of 4 scenarios can be found in Appendix B. The balancing scenarios improve the voltage margin relative to both baselines. The cost optimization with the phase changing ability pushes load to high PV supply window and the balancing term spread the load evenly to all phases reducing the voltage spike between phases. Thus, the voltage metric improves despite not being part of the objective. However, under extreme conditions the optimization can result in more pronounced voltage drop. For example, during a day with a wide difference between hourly prices, the cost optimization could push charging heavily on the low price window, causing high load and thus high voltage drop.

5.2.3. EV Energy Delivery and SOC Fulfillment

Energy delivery is unaffected by either balancing layer on this installation. All 11 vehicles reach their requested departure energy in both Sc3 and Sc4, delivering the same 213.1 kWh as the baselines with zero shortfall. The charging window lengthens slightly, ending at 18:29 under Sc3 and 17:59 under Sc4 against 14:16 under Sc1, as the optimizer exploits the full parking durations. The peak single-phase loading provides a complementary service-side observation: Sc4's maximum 15-minute phase current of 56.6 A, against 65.1 A under Sc3, both within the 80 A installation rating, indicating that the smoothing layer incidentally also softens the plan's reliance on PV-boosted headroom. The zero shortfall confirms that, at this occupancy level, the session slack comfortably absorbs the small within-window delivery deviations of the inner loop discussed in Section 3.3, so the drift-recovery mechanism of the receding horizon is sufficient without an explicit in-window energy correction.

5.2.4. Charging Cost and Cost Premium

The initial hypothesis was that the balancing scenarios would cost more than cost-only optimization, since they carry an additional objective. The result contradicts this. At the nominal operating point the charging cost is 0.72 cent/kWh under Sc3 and 0.64 cent/kWh under Sc4, against 0.80 cent/kWh under Sc2, so both balancing scenarios are cheaper than pure cost optimization. This apparent paradox arises because Sc3 and Sc4 differ from Sc2 in two ways at once, not one. They add the balance term, but they also enable phase switching, which enlarges the feasible set and can only lower the cost minimum.

To separate isolate the effect of load balancing from the advantages presented by phase switching, the cost-only scenario is re-run with phase switching enabled, denoted Sc2b. The comparison Sc2 to Sc2b isolates the effect of switching freedom, which reduces the cost from 0.80 to 0.67 cent/kWh. The comparison Sc2b to the balancing scenarios then isolates the true cost of the balance objective on a common footing, and it is small, a premium of $+7.5\%$ for Sc3 and -4.5% for Sc4, as summarized in Table 5.2. The balancing premium is therefore measured against Sc2b rather than Sc2, since only Sc2b shares the phase-switching freedom of the balancing scenarios. On this basis the balance objective is

close to cost-neutral, adding a few percent of either sign rather than a systematic penalty.

Sc4 is cheaper than both Sc3 and Sc2b because the inner loop ability to responds to the 20-second PV fluctuations that the 15-minute outer loop is blind to(depicted on Figure 5.8, with al phase included in Appendix B). During PV supply dips, Sc4 with the help of inner loop reduces load while other scenarios are unable to adjust, pushing them to import energy from the grid, sometimes at high price. Similarly, during PV supply spike the outer loop remains at constant while inner loop absorbs the extra supply. This responsiveness raises the PV utilization of Sc4 to 74.92%, the highest of all scenarios, against 74.05% for Sc3 and 74.61% for Sc2b, and the additional self-consumed generation displaces grid import at a cost saving. The effect is modest, under two percentage points of utilization, but it is consistent in direction with the inner loop's purpose and anticipates the flicker-mitigation mechanism examined in Section 5.2.5. The average total current is identical across all scenarios at roughly 40.6 A, confirming that these differences arise from redistributing and retiming the same energy rather than delivering a different amount of it.

Table 5.2: Charging cost, cost premium, and PV utilization at the nominal operating point. The cost premium of the balancing scenarios is measured against Sc2b, the cost-only scenario with phase switching enabled, so that phase-switching freedom is held constant and the premium reflects the balance objective alone.

Scenario	Cost [cent/kWh]	Premium vs Sc2b [%]	PV utilization [%]
Sc1	2.46	–	54.72
Sc2	0.80	–	73.07
Sc2b	0.67	0.0	74.61
Sc3	0.72	$+7.5$	74.05
Sc4	0.64	-4.5	74.92

5.2.5. Voltage Ramps

Phase balancing does not introduce severe flicker, and the inner loop of Sc4 yields both the lowest visible-flicker occurrence and no annoying flicker at all. The visible-flicker occurrence at the facility bus decreases monotonically across the scenarios, from 0.915% under Sc1 to 0.900% under Sc2, 0.847% under Sc3, and 0.523% under Sc4, as summarized in Table 5.3. Meanwhile, only Sc2 registers any annoying flicker (56 events, 0.022% of the time), produced by the synchronized current steps of price-driven charging while Sc1, Sc3, and Sc4 register none.

Further split between ramp events at 15-minute boundary and within-window(as introduced in 5.1.1), further confirms the effectiveness of the inner loop to suppress the solar flicker. Sc2 and Sc3 both decreases total flicker which is caused by the reduce of flicker at 15-minute boundary. The decrease at Sc2 and Sc3 is driven by the optimizer cost optimization slowly ramping the load instead of maximizing EV load at when the EV arrives and also phase switching allowing move even distributions of load. However, only Sc4 is effectively decreasing flicker within windows isolating the ability of solar flicker reduction to the inner loop native to Sc4.

Table 5.3: Voltage ramp and flicker statistics at the facility bus across scenarios. Event counts are decomposed by their timing relative to the 15-minute scheduling boundaries. Each phase is counted separately, so the counts are taken over phase-windows and are consistent with the occurrence percentages.

Metric	Sc1	Sc2	Sc3	Sc4
Visible flicker, > 1.6 V/20 s [%]	0.915	0.900	0.847	0.523
Annoying flicker, > 3.6 V/20 s [%]	0.000	0.022	0.000	0.000
Total ramp events (> 1.6 V)	2370	2333	2195	1356
at boundaries (switching)	402	287	142	119
within windows (PV flicker)	1968	2046	2053	1237

Figure 5.8 illustrates the result of inner loop by comparing Sc3 and Sc4. As opposed to Sc3's load profile, Sc4 responds to PV supply fluctuation which is represented by the current headroom. The responses can be seen as load dips that are proportional to the headroom. However, the response is not guaranteed as the adjustment is constrained by coupling term ensures symmetrical load adjustment, allowing load adjustment response only if all three phases have rooms to adjust, which ultimately prevents load imbalance. Further tests with different amount of EV participation and PV capacity needs to

be carried out to test the limit of the inner loop.

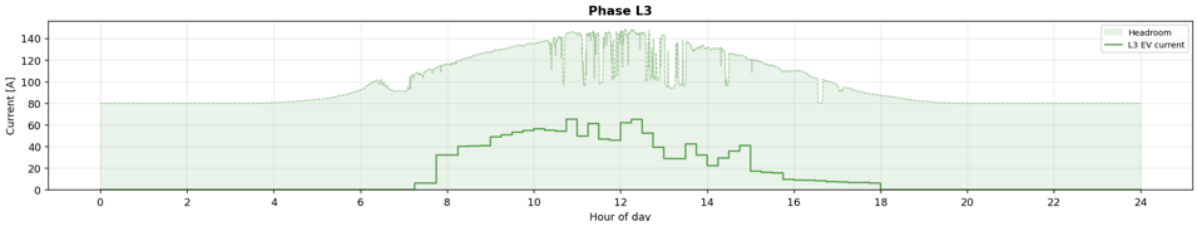

(a) Load profile of Sc3 phase L3

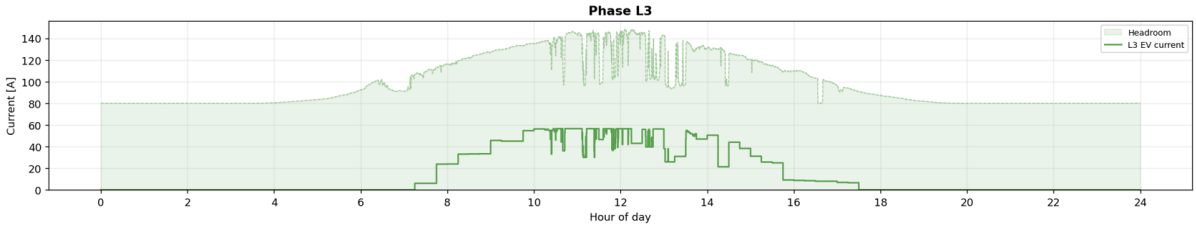

(b) Load profile of Sc4 phase L3

Figure 5.8: Aggregate charging load of Sc4 reacts to changes of headroom/PV supply.

5.2.6. Other Observations

The phase shifting pattern of Sc3 and Sc4(represented by Sc4 on Figure 5.9) reveals a flaw that needs to be addressed by future works. Currently, phase shifting is unconstrained and costless. Thus, EV chargers switch phases often during the simulation. This will be critical upon hardware implementation considering that the phase shifting process is not instantaneous. This limitation will be discussed further in Chapter 6. Furthermore, during phase-switching or charging mode switch power goes to zero before picking back up [8]. Thus, the real voltage transient in real setting is worse than the simulated result in Subsection 5.2.5.

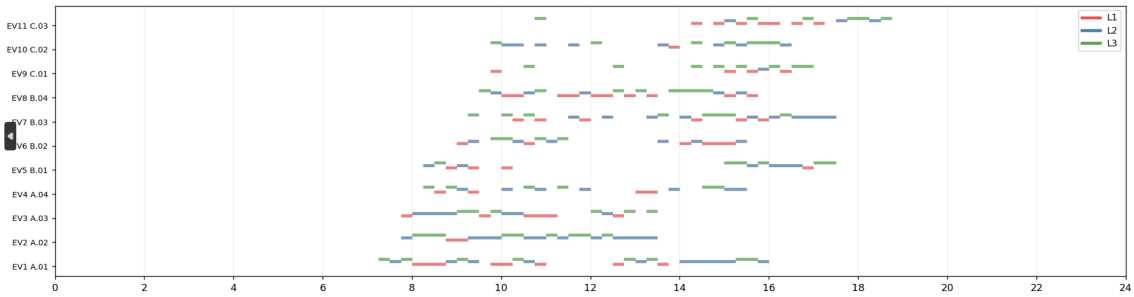

Figure 5.9: Phase assignment of EVs in Sc4

5.3. Objective Weighting Tuning

The weighting coefficients are the only part of the formulation that a user is expected to change after deployment, which makes their behavior worth characterizing in its own right. Two reasons motivate the sweep. First, the nominal weights of the preceding sections were set empirically for one operating point. A facility with a different fleet size, PV capacity, or price spread may not be best served by the same balance. Second, and more usefully, the weights are the interface through which prioritization can be expressed, a DSO seeking maximum unbalance mitigation and an operator seeking minimum cost occupy different points on the same axis. What matters for practical use is therefore not only which weighting performs best, but whether performance degrades gracefully or abruptly as the weighting moves away from it.

In terms of load imbalance, the result from the initial weighting remains true up to an α/β ratio of 100, where the average imbalance stays within 1.22–1.24 A and the worst case at 2.92 A for both scenarios. At α/β of 1000 and above, C_{cost} becomes too dominant and imbalance rises, with the average reaching

1.77 A and the worst case 12.18 A, as illustrated in Figure 5.10. A separate spike appears in Sc4 at the smallest weighting ($\alpha = \beta = 10^{-2}$), reaching a worst case of 9.78 A. This value does not follow the ratio trend and is a symptom of the poor solver conditioning noted above rather than a genuine optimum.

The opposite trend holds for the average energy cost, which decreases as the weighting ratio increases. Across the grid the cost ranges from about 1.06 cent/kWh at the lowest ratios to 0.53 cent/kWh at the higher ratios, a spread of roughly 0.5 cent/kWh. It can be concluded that the ratio produces a predictable and relatively small absolute change in the two main objectives for this particular condition, provided the solver has converged. In practice this provides safety in tuning within the converged region, with little risk of falling off a cliff, while the large jumps at the extreme corners are dominated by non-convergence rather than genuine trade-off behavior.

The unmet energy remains below 0.1% across Sc3 and across most of Sc4. In Sc4 a number of cells show higher unmet, up to 1.2%, concentrated where the cost weight α is small. Two mechanisms contribute to this. Where the solver reached optimality, the weak cost incentive combined with the inner loop's downward modulation leaves a small residual shortfall, which is a genuine result.

The key point of the weight tuning is that different setup (grid configuration, load pattern, etc) and prioritization can be advantaged by a certain weight combination. For this particular setup, the α/β ratio of 100 seems to produce the best result, pushing both the cost and imbalance to their minimum. However, repeating the analysis with different data input is necessary to come into more informed conclusions.

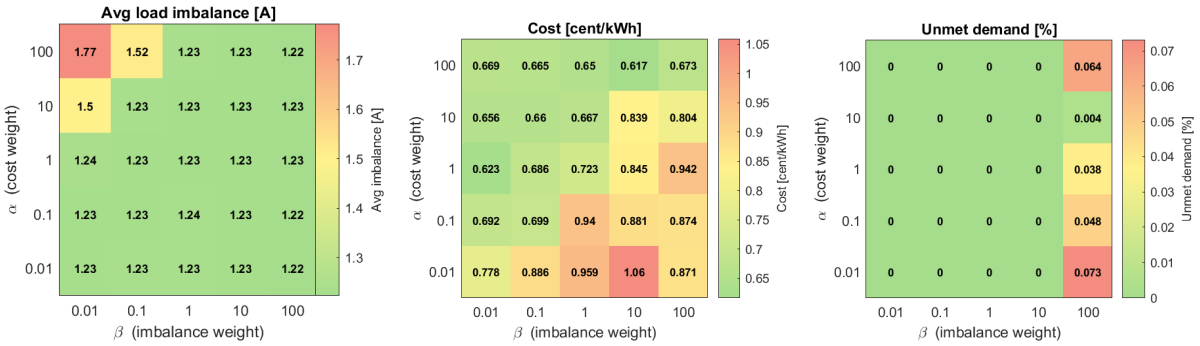

(a) Average load imbalance. (b) Cost. (c) Unmet demand.

Figure 5.10: Sensitivity of the main performance metrics to objective weighting in Scenario 3.

5.4. Sensitivity to Operating Conditions

5.4.1. Charging Cost

Charging cost tracks PV utilization across sizing combinations, and the two sweep axes therefore act on it in opposite directions. At low installed capacity the generation is fully absorbed whatever the fleet size, so utilization saturates near 90% and the cost rises with EV count because every additional vehicle must be served from the grid. Under Sc4 the cost at 30 kWp climbs from 1.93 to 3.05 cent/kWh as the fleet grows from 9 to 13 vehicles, while utilization moves by less than three percentage points. At high installed capacity the binding constraint reverses. Generation now exceeds what the fleet can absorb, utilization falls to roughly 55% at 70 kWp, and adding vehicles consequently raises utilization rather than cost, because each additional session consumes generation that would otherwise be exported. The fleet size that is most expensive at 30 kWp is therefore among the cheapest at 70 kWp.

That ordering does not hold cell by cell at the right-hand edge of the grid, and the reason is worth stating because the same behavior recurs in the load-imbalance results. At 60 and 70 kWp the differences between neighboring fleet sizes fall to a few hundredths of a cent, and the ranking of those cells no longer follows the fleet axis in any consistent way. This is an expected result. The optimizer does not minimize cost alone but the weighted sum of cost, imbalance and unmet energy, so once one term has been driven close to its floor the solver is free to trade it against the others at no penalty to the objective value. Any solution within that margin is equally good from the optimizer's point of view, and the reported metric consequently appears to swing back and forth. The practical implication is that the sweep should be read for its gradients rather than for the ordering of adjacent cells wherever the underlying quantity has saturated.

The cost advantage of Sc4 over Sc3 due to inner loop analyzed in Subsection 5.2.4 remains except for a region with low EV to PV ratio. Sc4 costs less in 21 of the 25 cases, by 0.096 cent/kWh on average. The four exceptions all lie in the same corner of the grid, where the ratio of connected EVs to installed PV is the lowest. A reduction of cost gap between Sc3 and Sc3 is expected as the ability to adjust to PV fluctuation diminishes when EV/PV ratio is lowered(comparison can be found in Appendix B). However, further analysis into total EV load profile shows that when Sc4 inner loop fails to compensate, its outer loop yield a flatter load profile relative to Sc3 outer loop. Which is not an expected behavior since the outer loop of both scenarios are identical. Further analysis is required to uncover the relations between inner and outer loop in this region.

Measured against the phase-switching baseline Sc2b, the premium of the balance objective remains small across the grid, averaging $+7.3\%$ for Sc3 and -3.6% for Sc4, with Sc4 cheaper than Sc2b in 18 of the 25 cases. For the most part the advantage and disadvantage of the two balancing scenarios follow the pattern already visible in their own cost maps. A number of combinations depart from it, and in those cells the departure is attributable to the baseline rather than to the balancing scenarios, since Sc3 and Sc4 deviate together under the same conditions while Sc2b returns a cost that is difficult to reconcile with its neighbors. Those cells sit where the phase-switching binary makes the problem most expensive to solve, consistent with the solver conditioning discussed in Section 5.3. The central tendency rather than the individual cells is therefore the reliable measure here, and on that basis the conclusion of Section 5.2.4 holds across the operating space. The balance objective is close to cost-neutral, and it is not the dominant influence on what the operator pays.

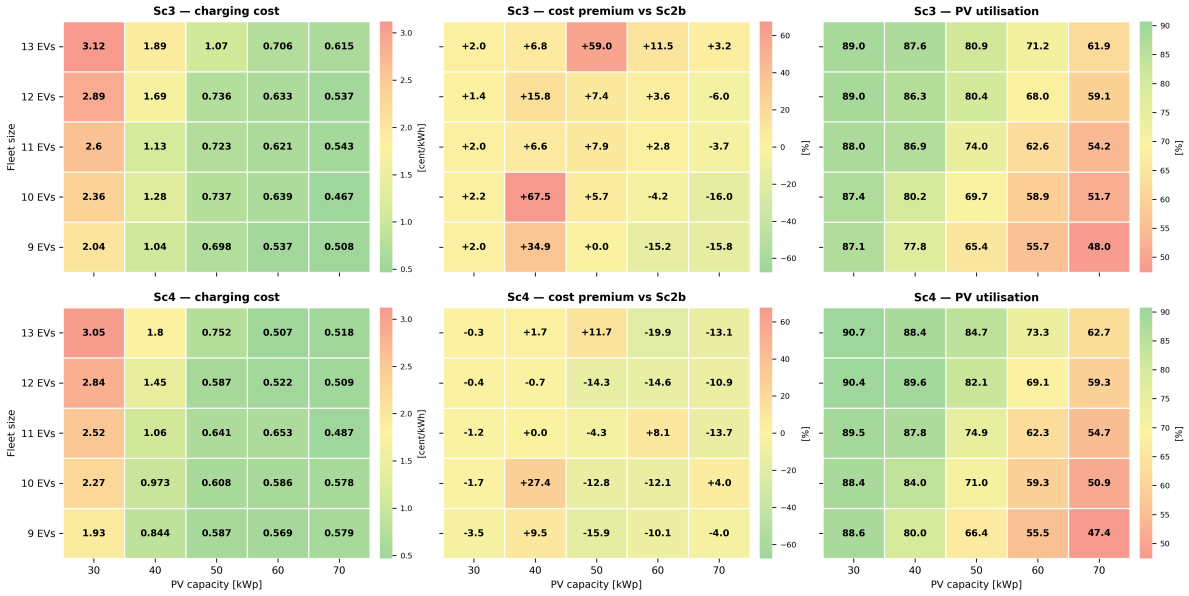

Figure 5.11: Charging cost, cost premium against Sc2b, and PV utilization across the 5×5 fleet-size $\times$ PV-capacity grid, for Sc3 (top) and Sc4 (bottom). Cost and PV utilization share a color scale within their column so the two scenarios are directly comparable, and the premium panels share a scale centered on zero, with green indicating a case cheaper than Sc2b.

5.4.2. Load Imbalance and Voltage Unbalance

The load imbalance and the VUF move together across the operating space, the first as the current-level cause and the second as its voltage-level consequence. Once phase balancing is active both are essentially invariant to the operating point. Across all 25 combinations the average imbalance under Sc3 and Sc4 stays within a 1.16–1.24 A band and the worst case within 2.84–3.20 A, a spread of under 0.1 A in the average despite the operating point varying over the full grid. The VUF follows directly, with the maximum at the facility bus held within 0.18–0.20% and the average within 0.097–0.102%, as shown in 5.4. Neither the EV count nor the PV capacity moves any of these values appreciably, so the balancing scenarios flatten both the current imbalance and the resulting voltage unbalance almost completely.

The baselines, by contrast, are strongly condition-dependent at both levels. Uncontrolled charging (Sc1) ranges from 4.92 to 6.64 A in average imbalance and up to 32.63 A at worst, and cost-only opti-

mization (Sc2) is worse at the tail, reaching 47.17 A as it synchronizes the fleet. This carries through to the VUF, where Sc1 rises to 1.82% and Sc2 exceeds the EN 50160 limit of 2% in 11 of the 25 cases, reaching 2.58%. The central conclusion of 5.2, that balancing suppresses unbalance, therefore holds across the full range of tested conditions at both the current and the voltage level, and is not an artifact of the nominal operating point.

Table 5.4: Range of facility-bus load imbalance and VUF across the 5×5 EV-count $\times$ PV-capacity grid.

Scenario	Avg. imbalance [A]	Max. imbalance [A]	$\mathrm{VUF}_{\mathrm{avg}}$ [%]	$\mathrm{VUF}_{\mathrm{max}}$ [%]
Sc1	4.92–6.64	18.40–32.63	0.30–0.39	1.04–1.82
Sc2	4.37–6.92	25.07–47.17	0.27–0.41	1.33–2.58
Sc3	1.16–1.24	2.84–3.18	0.097–0.102	0.18–0.20
Sc4	1.16–1.24	2.84–3.20	0.097–0.102	0.18–0.20

Going into the detail in the heatmap presented on Figure 5.12 reveals the stability of the current optimizer. The average VUF values of all 25 cases are at optimum. The noticeable improvement from 11 EVs to 12 EVs is caused by the arrival of EV number 12, which arrives at 07.15. Previously, only one EV was available that early, which cannot provide symmetry since at least two EVs are required to counterbalance the installation load on phase A. This improvement is not reflected by decreasing the EV number from 11 or increasing it from 12, since the same condition is not replicated by these changes. This finding highlights that although an increased EV number is likely to add more resource to achieve balance, the schedule greatly affects EV availability, which ultimately impacts the solution quality.

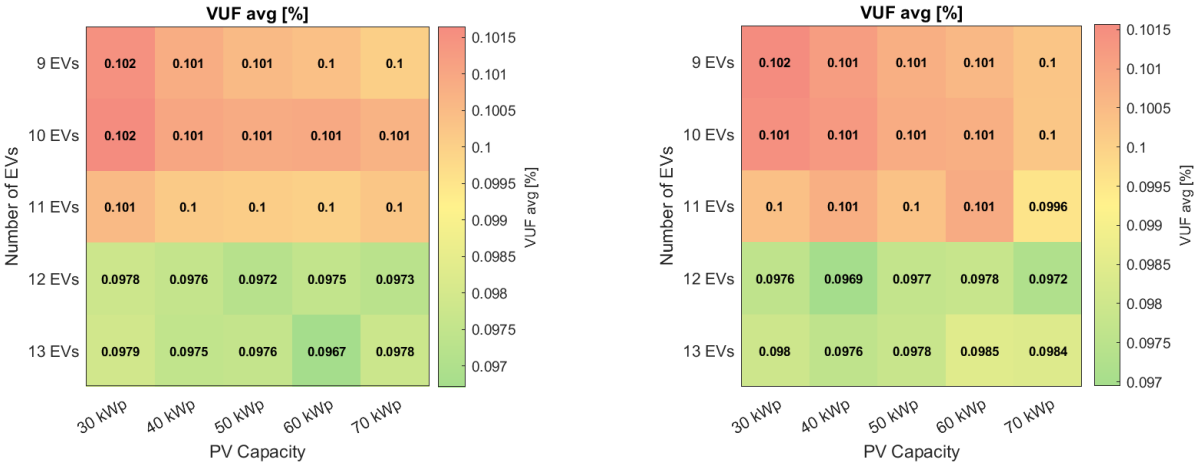

(a) Scenario 3

(b) Scenario 4

Figure 5.12: VUF sensitivity results for Scenarios 3 and 4.

5.4.3. Voltage Ramps

The flicker response, unlike the VUF, is governed by the PV capacity rather than the EV count. The visible-flicker occurrence rises steeply and monotonically with installed PV across all scenarios, while varying only weakly across the EV axis, as shown on Table 5.5. This is the physically expected behavior, since flicker originates in PV output fluctuation, and is consistent with the dependence on installed PV capacity reported by [22]. At every PV level, Sc4 achieves the lowest occurrence, which confirms that the inner loop's benefit is not specific to the nominal case. The sweep also exposes the operating limit of the inner loop. At low-to-moderate PV, Sc4 reduces visible flicker substantially below Sc3, for example from 0.86% to 0.56% at 50 kWp. The advantage narrows as PV capacity grows and the fluctuations begin to exceed the modulation authority of the connected EVs, and at 70 kWp the two scenarios converge to 1.43% and 1.26%.

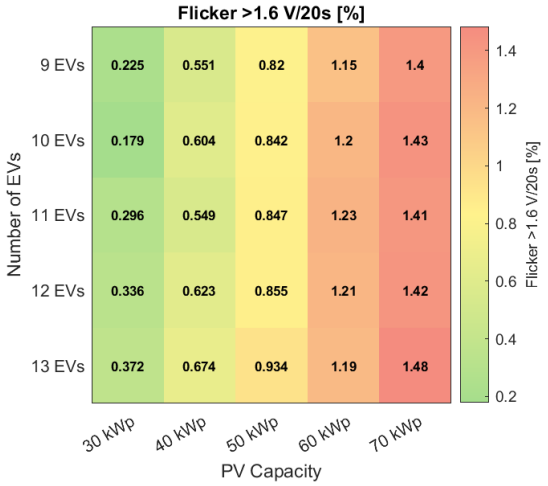

(a) Flicker 1.6: Sc3

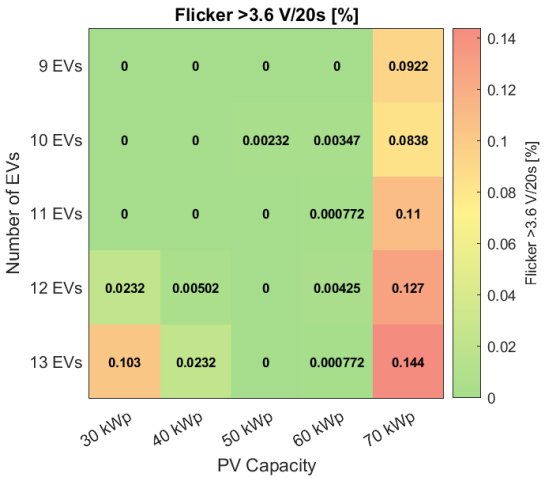

(b) Flicker 3.6: Sc3

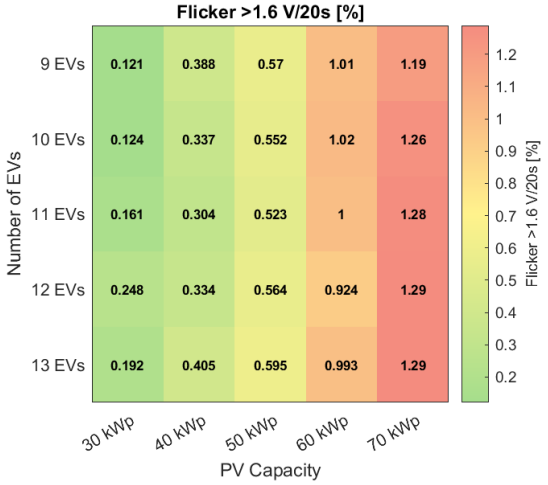

(c) Flicker 1.6: Sc4

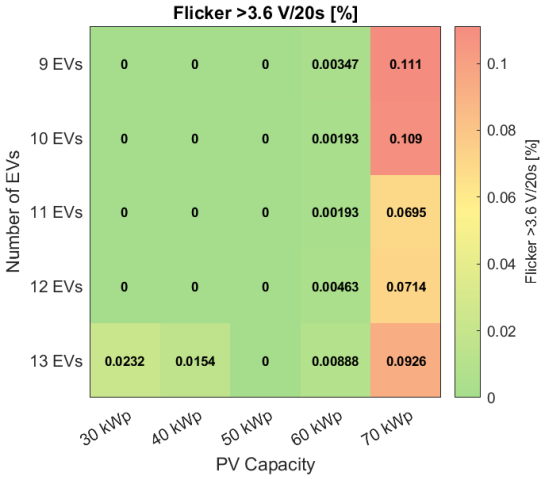

(d) Flicker 3.6: Sc4

Figure 5.13: Flicker sensitivity result heatmaps.

The annoying-flicker threshold shows the same limit more sharply. Sc4 holds annoying flicker at zero up to 50 kWp across all EV counts and near zero at 60 kWp, but at 70 kWp annoying flicker appears in every scenario, including Sc4, at up to 0.11%. In fact, in some conditions Sc4 seems to have worse annoying flickers than Sc3. However, this worsening cannot be attributed to failure of Sc4, because these annoying flickers are happening at the boundaries of 15 minutes, which means that it is the side effect of the outer loop that is blind to the smaller transient. Applying maximum power ramp could potentially reduce this type of flicker but it also has potential to raise another issue because the degree of freedom of the outer loop is now even more limited.

Table 5.5: Visible-flicker occurrence [% of time, facility bus] versus PV capacity, averaged over the EV-count axis.

Scenario	30 kWp	40 kWp	50 kWp	60 kWp	70 kWp
Sc1	0.25	0.61	0.91	1.27	1.46
Sc2	0.23	0.54	0.91	1.29	1.55
Sc3	0.28	0.60	0.86	1.20	1.43
Sc4	0.17	0.35	0.56	0.99	1.26

6

Power Hardware-in-the-Loop Validation

The sixth chapter presents the Power Hardware-in-the-Loop validation of the framework. Section 6.1 describes the testbed and its components, covering the setpoint scheduler that dispatches the precomputed optimization results, the real-time simulator hosting the grid model, the power amplifier forming the interface between simulation and hardware, the Zaptec Pro charger as the device under test, and the on-board charger emulator that terminates the power loop. Subsequently, Section 6.2 reports the feasibility results, beginning with the de-rated scenario required by the hardware limits and proceeding through the ability of the device under test to execute the commands the optimizer issues, the conformance between commanded and delivered power, the execution of the schedule over the full run, and the discrepancy in charged energy between the optimizer and the measurement. It is important to note that two objects are under test at once. The device under test is examined for whether it can execute what the framework requires, while the optimizer is examined for whether its assumptions survive contact with hardware.

6.1. PHIL Setup Description

The PHIL testbed used in this thesis builds on the demonstrator developed at TU Delft by [6] and subsequently extended by [21]. In both predecessors, a Digital Real-Time Simulator (DRTS) runs the LV distribution grid model while a power amplifier synthesizes the resulting three-phase node voltages at the terminals of a physical EV charger. The present work retains this architectural template but replaces the EVSE with a Zaptec Pro charger, substitutes the optimization, and replaces the load flow simulation with unbalanced FBS. The optimization is not run in real time and a setpoint scheduler is taking its place, running in parallel with the grid model extracting the simulation time ensuring synchronized timing. The logger then store the EV charger setpoint history and power measurement from the power supply for analysis. Table 6.1 summarizes the components used on the setup and the following subsections describe each component in turn. Figure 6.1 provides a schematic overview. The detail of the components can be found on Appendix C

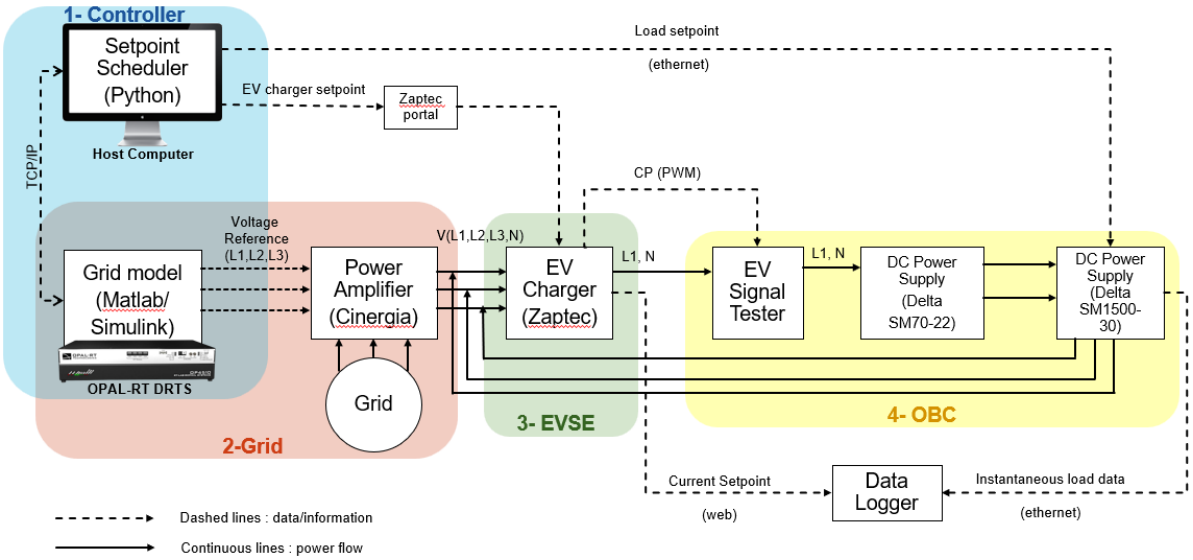

Figure 6.1: Overview of proposed PHIL testbed setup.

Table 6.1: Equipment used in the PHIL testbed.

Subsystem	Component	Model	Function
Controller	Host computer	–	Runs the setpoint scheduler
	Scheduler	Python	Dispatches current and phase commands
	Model host	MATLAB/Simulink, RT-LAB	Compiles and deploys the grid model
Grid	Real-time simulator	OPAL-RT OP4510	Solves the FBS load flow at 0.1 s
	Power amplifier	Cinergia GE-40	Synthesizes three-phase node voltages
EVSE	Device under test	Zaptec Pro	Phase-switching AC charger, 22 kW
	Connector	IEC 62196 Type 2	Control pilot and power interface
OBC	Signal tester	Walther-Werke EV tester	Provides the ready-to-charge CP signal
emulator	DC supply (load)	Delta Elektronika SM70-22	Draws single-phase AC, 1.5 kW limit
	DC supply (return)	Delta Elektronika SM1500-30	Returns power three-phase
Logging	Data logger	–	Records setpoints and instantaneous load

6.1.1. Setpoint Scheduler

The optimizer as described in Chapter 3 is run in advance and the results are fed into a setpoint scheduler whose output control EV charger and OBC load setpoints which is run on the host computer at sync with DRTS simulation time. The phase assignments and current setpoint is sent EV charger via the Zaptec portal and the load set point is sent via Ethernet. Ideally, the control pilot(CP) output of the charger is sniffed and processed to determine the load drawn by the OBC emulator. However, due to time and technical limitation encountered on the process, the setpoint scheduler uses EV charger setpoint and load schedule are used to determine the charging, trying to mimic real OBC response to CP.The behavior of OBC can greatly vary with the pairing of EV charger and the EV as tested by [8] who tested zaptec and ACDC chargers with three different cars. For example, the total time to process charging mode switch from three to single-phase vary between the two tested vehicles as depicted on figure 6.2

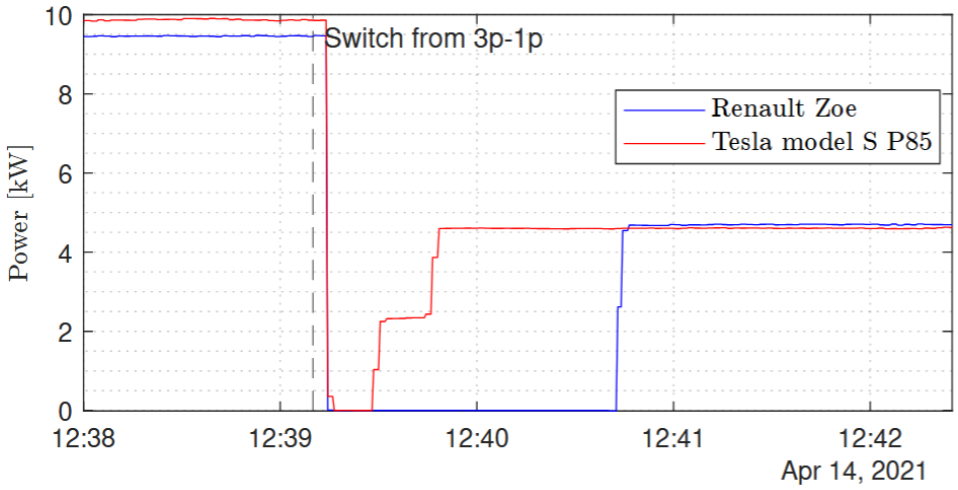

Figure 6.2: Changing charging automatically from three-phase to single-phase with Zaptec chargers[8]

Real EV on-board chargers do not respond to setpoint changes instantaneously. The IEC 61851-1 standard corroborates this, requiring the OBC to complete its response to a PWM change within 5 seconds [8]. As context, two EVs tested by [8] presented 4 seconds delay to PWM input. To comply with this behavior, the setpoint scheduler simulates ramps by sending 5 steps in 5 seconds whenever a new current setpoint is dispatched, both for ramp-up and ramp-down transitions.

Earlier in Subsection 5.2.6 it was established that phase switching requires a more conservative procedure. When the MILP assigns an EV to a different phase between two outer-loop timesteps, the scheduler ramps down the load to zero, executes phase change, and ramps back up to target load. The process takes 60 seconds which is consistent with the 30–90 second total switching time reported by [8] for Zaptec chargers paired with three-phase EVs.

The setpoint command path from the host computer to the charger involves two steps. An HTTPS REST call from the host to the Zaptec cloud server to issue the phase and current command and the Zaptec cloud forwarding the command to the charger over its cellular/wifi connection. This cloud communication results in longer delay, 300-400ms on average, compared to direct commands. This is negligible relative to the 15-minute outer-loop timestep and even relative to the 20-second inner-loop timestep, and does not represent a binding constraint on feasibility. The binding constraint is instead the OBC response time following a phase-switch command.

As discussed in Section 3.2.1 the optimizer aims to control the OBC with only controlling EV charger PWM setpoints. Thus, the scheduler anticipates this by increasing the setpoints relative to the intended load by 400W, following typical OBC accuracy [8]

6.1.2. Real-Time Simulator

The grid model described in Chaper 4 is implemented in MATLAB/Simulink on the host computer, compiled OPAL-RT OP4510 Digital Real-Time Simulator via RT-LAB. Once running on the target machine, the DRTS solves the three-phase load-flow equations within a fixed simulation timestep of 0.1s.

6.1.3. Power Amplifier and Interface Algorithm

The power amplifier converts the per-phase reference voltage computed by the DRTS into actual three-phase voltages at the charger connection terminals. The testbed uses Cinergia GE-40 which simultaneously translate DC analog signals to produce three-phase AC voltage instead of using power amplifier for each phase. It is also capable of bidirectional power flow which is important for this setup with the reason explained in Subsection 6.1.5.

6.1.4. Device Under Test: Zaptec Pro Charger

The Zaptec Pro is a single-outlet, wall-mounted AC charger rated at 22 kW (3-phase, 32 A per phase) with an IEC 62196 Type 2 socket. It is one of the few commercially available EVSE products with native phase-switching capability. The feature allows the charger to assign the EV connection to any

one of the three supply phases under remote command, or operate in three-phase mode when the connected vehicle supports it [28]. It is originally designed in conjunction to load sensing device which detects load usage of an installation to then allow zaptec pro to manage load, maximizing installation capacity. Without it, the setpoint is easily customizable through the its user portal. This makes it the natural candidate as the physical device under test for a framework whose central decision variable is the per-EV phase assignment.

The charger is connected directly to the output terminals of the power amplifiers, which supply the three-phase reference voltages derived from the DRTS grid model. The Zaptec Pro communicates the actual charging current to an internal certified meter, which logs data to the Zaptec cloud [8].

All AC charging in the Zaptec Pro follows the IEC 61851-1 Control Pilot (CP) protocol. The charger generates a 1 kHz PWM signal on the CP pin of the Type 2 cable Subsequently, the duty cycle encodes the maximum allowed charging current and is interpreted by the EV on-board charger (OBC). Zaptec Pro achieves a current modulation resolution of 0.1 A, substantially finer than the ± 0.6 A tolerance permitted by the standard, and that the PWM update latency from API command receipt to CP signal change is approximately 400 ms [8]. The IEC 61851-1 standard requires the charger to update the CP signal within 10 seconds of receiving an external command and the OBC must then respond within 5 seconds. Both limits are met by the Zaptec Pro and by all three-phase EVs tested against it in the literature [8].

An important hardware constraint is that single-phase charging through a three-phase OBC is only possible via the first phase (L1) of the vehicle's Type 2 cable. Furthermore, switching from single three-phase to single-phase mode is only possible to L1 because switching to other phases will cause an error [8]. This constraint contributes to the limitation of only using single-phase charging defined in Section 3.2.

6.1.5. On-Board Charger Emulator

Because the Zaptec Pro requires an active IEC 61851 CP handshake to enable charging, a physical EV or emulator must be connected to the Type 2 socket. As depicted on Figure 6.3, Walther-Werke EV signal tester providing the CP ready-to-charge signal, with a back-to-back Delta Elektronika power supply pair forming the power circulation unit.

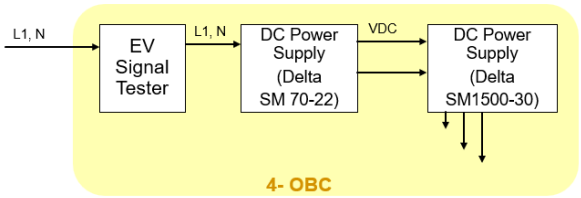

Figure 6.3: The setup of On-Board Charger Emulator, receiving single phase AC supply and feeding back the grid though three-phase connection

The absence of DC power supply that can work with both single and three phase input interchangeably ultimately excludes three-phase charging as described on Section 3.2.1. Thus, a single-phase power supply is used to interface with the output of the power, prioritizing single-phase phase-changing simulation. Due to the limited availability of single phase DC power supply, a single Delta SM70-22 is used as shown on figure 6.3. Consequently, the OBC can only handle 1500W maximum power and the scenarios run in Chapter 5 cannot be used. Therefore, another scenario needs to be run considering this limitation.

The second DC power supply is a three-phase power supply that feeds back the power through the a terminal post the power amplifier creating a circulating power routing. The setup is ideal for energy saving especially when emulating three-phase charging as all the power loops are entirely closed. However, this setup draws power from single phase while feeding back in three-phase. Thus, only a third of the power is being circulated back while the remaining is imported back to the grid which necessitates power amplifier that can handle bidirectional power flow. Replacing this power supply with another single-phase power supply will not solve the issue because the charger draws power from different phases due to the phase changing feature.

6.2. Feasibility Results

Before the individual criteria are examined, one precondition applies to all of them. The grid model executed on the OPAL-RT OP4510 met its 0.1 s timestep for the entire duration of every run reported in this chapter, with no computation overruns recorded by RT-LAB. This satisfies the requirement that the simulator resolve the system equations within the defined timestep at all times, since overruns may render a real-time simulation unstable [38]. The results that follow can therefore be attributed to the behavior of the hardware and the control path rather than to degradation of the simulation itself.

6.2.1. Optimizer Re-run for PHIL simulation

The limitation of the OBC necessitates a scenario where the tested OBC power does not surpass 1500W. Thus, for this particular scenario one EV current limit is capped at 6A. The schedule for this particular scenario can be found in Appendix A. The optimization results along with grid inputs are then cut into 1 hour simulation time from 11:30 to 12:30. This simulation time is chosen because it contains two phase changing events at 11:45 and 12:00.

6.2.2. DUT Command Execution Performance

The first feasibility criterion concerns whether the Zaptec Pro executes the two command types the optimizer relies on: current setpoint adjustment and phase reassignment. Both were executed for every command dispatched over the test run, with no rejected or dropped commands. Current setpoint changes propagated to the CP duty cycle well within the 20 s inner-loop interval, and phase reassignment completed within the 60 s window allocated by the scheduler. Against the criterion as stated, the answer is affirmative: the DUT is capable of executing the command set the phase-balancing framework requires, and neither the cloud communication path nor the charger itself constitutes a binding constraint at the timescales at which the optimizer operates.

This conclusion is nonetheless bounded by how the phase switch is currently sequenced. In the realized setup the scheduler drives the OBC emulator load directly, ramping it to zero before the phase-switch command is issued, so the charger performs the switch under a known no-load condition. In a real installation the ordering is not guaranteed by construction. The charger interrupts the CP signal and the OBC responds with its own latency, which means the outcome depends on whether the current has decayed before the contactor operates.

The quantity that matters here is therefore the interval between the setpoint change and the physical execution of the phase change by the charger. This interval bounds how quickly the ramp-down must complete: if the charger actuates its contactor before the current has reached zero, the switch occurs under load, which is both electrically undesirable and a hazard to the contactor. The converse delay, the time taken to re-establish the CP signal and restore charging on the new phase, does not carry the same weight, since the scheduler already allocates a conservative 60 s window on the assumption of OBC restart, and any re-establishment latency shorter than that is absorbed without consequence.

This interval is not available from the literature in the form required. [8] report a charger-side latency of approximately 400 ms from API command receipt to CP signal change, and a total three-phase to single-phase transition time of roughly 15 to 90 s across two tested vehicles (Figure 6.2), but that spread reflects OBC restart behavior rather than the charger. The contactor actuation time relative to the CP interruption is not separately reported. Repeated switching tests instrumented at the CP pin would isolate it and establish whether it is consistent enough to be treated as a fixed parameter, from which a minimum safe ramp-down duration could be specified.

Resolving this would also enable the more substantial improvement identified in Section 6.1.1, which is replacing the schedule-driven OBC emulator with one that decodes the CP duty cycle and sets the DC supply current accordingly. Such an emulator would react to the CP interruption as a real OBC does, rather than relying on the scheduler to impose the correct ordering in advance.

6.2.3. Setpoint versus Measured Power

The second criterion concerns conformance between the power the optimizer commands and the power actually delivered. Table 6.2 compares three measurement points at a single steady-state operating point: the DC output of the Delta supply, an independent AC measurement at the charger terminals, and the value reported by the Zaptec internal meter.

Table 6.2: Comparison of measured quantities at a single operating point.

	Voltage (V)	Current (A)	Power (kW)	PF
Delta	69 VDC	19.7	1.37	–
Measurement	228.5 VAC	6.7	1.50	0.98
Zaptec	–	6.7	1.52	–

The AC-side current and the value reported by the Zaptec internal meter agree, and the corresponding power values differ by approximately 1.3%. Therefore, the charger's own telemetry is an acceptable proxy for delivered power, which matters because it is the only measurement exposed through the Zaptec API and would constitute the feedback signal in a closed-loop implementation. Meanwhile, the offset between commanded and delivered current(accuracy), is not visible in this comparison because it is already compensated, as described in Section 6.1.1. The scheduler dispatches a setpoint approximately 1.7 A above the intended value, corresponding to the 400 W margin. This correction is applied ahead of the measurement and does not require revision.

The result requiring attention is instead the difference between the AC power drawn and the DC power delivered to the load. At this operating point 1.50 kW is drawn on the AC side against 1.37 kW delivered on the DC side, a loss of approximately 8.7%. This loss is not currently represented in the optimization model, which computes charging schedules on the basis of AC power drawn from the feeder and treats the energy delivered to the battery as equal to it. The consequence is that the departure SOC achieved in practice falls short of the scheduled value by this margin, and that the optimizer systematically underestimates the charging duration required to meet a given energy demand.

The absence of a closed feedback path compounds this. Because the measured charger current is logged rather than returned to the load flow, the loss identified above does not propagate into the grid model, and the simulated feeder sees the commanded rather than the delivered power. The PHIL results and the offline power flow therefore agree by construction rather than by validation. Closing this path would allow the divergence between them to be quantified directly, and is a prerequisite for verifying that any efficiency correction introduced into the optimizer produces the intended effect.

6.2.4. Schedule Execution over the Full Run

Figure 6.4 presents the complete test run, comparing two data from the two sections prior that is logged, confirming that the behavior observed at those points is representative rather than incidental. The executed phase assignment and the load measurement match the scheduled assignment at every timestep, and every phase-switch event in the schedule appears in the executed trace at the expected time. Combined with the absence of rejected commands noted in Subsection 6.2.2, this is the direct evidence for the executability claim that the DUT reproduced the optimizer's phase allocation over the full run without divergence.

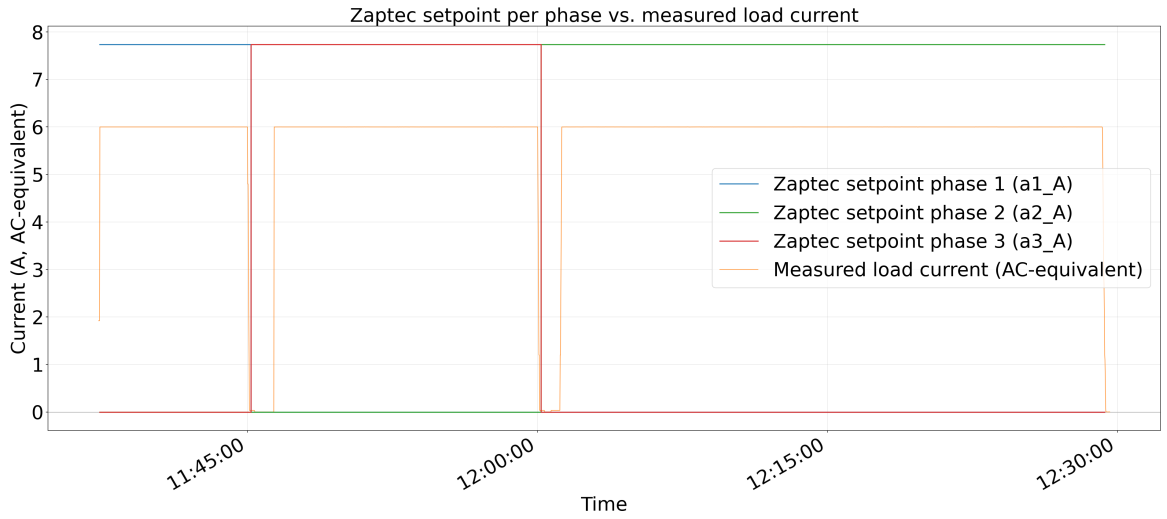

Figure 6.4: Charger setpoint and load measurement during hardware test

6.2.5. Charged Energy Discrepancy

The preceding subsections established that the schedule is executed and that a conversion loss is present at a single operating point. This subsection quantifies the cumulative effect on delivered energy, which is the quantity the optimizer's demand constraint is written in terms of.

The comparison is restricted to the window between 11:40 and 12:25 of simulation time. The first ten minutes of the hardware log are excluded because the components of the testbed are energized sequentially rather than simultaneously. The RT-LAB model is started first, and the amplifier only begins supplying the circulating loop once it receives a valid reference from the running model. Consequently, measurements taken before the setup is fully established do not correspond to the configuration under test. The final five minutes are excluded for the opposite reason, as the system is ramped down manually at the end of the run to avoid interrupting the supply while current is still circulating. The window retained contains both phase-switch events present in the schedule.

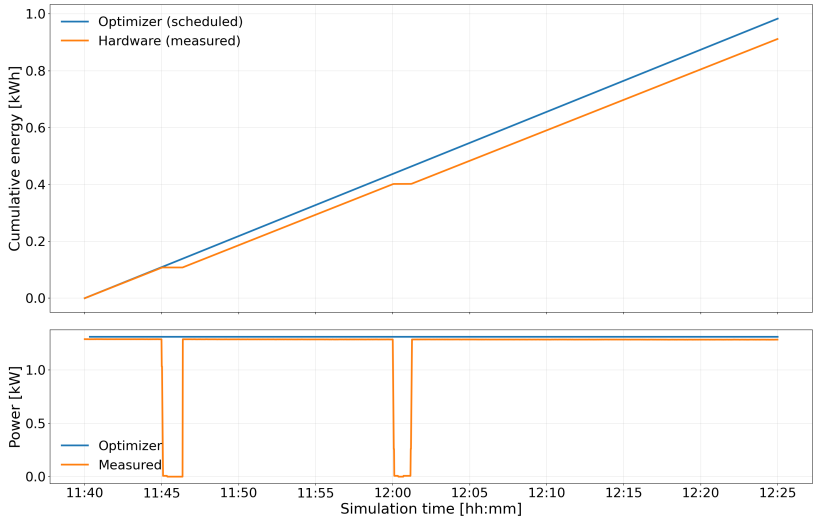

Figure 6.5: Cumulative charged energy over the comparison window: energy scheduled by the optimizer against energy measured at the DC load. The lower panel shows the corresponding instantaneous power.

Over the 45-minute window the optimizer schedules 0.983 kWh while 0.912 kWh is delivered, a shortfall of 0.072 kWh or 7.3%, which is attributed to two main reasons. The first is power difference due to voltage assumption used in the optimizer. Because the optimizer is load targeting, the voltage

at the bus is assumed and fixed, as opposed to voltage simulated in the PHIL setup which is dynamic, following the value of the load flow. The scheduled and measured power differ by considerably less than this figure suggests, the optimizer dispatches 1311W against 1283W measured, a difference of 2.2%. This difference cannot be solved other than by opting for voltage targeting algorithm which will considerably increase the computational load.

The second and larger part of the shortfall arises due to the phase-switch events which is not considered in the optimization. During the 45 minutes window, 100 second was spent with no current and 20 on ramping, corresponding to 0.050 kWh at the prevailing power, or 70% of the shortfall. Solving this difference is relatively simple. It can be done by adding constraints or switching penalty to the outer loop. This addition will also improve a known issue that has not been tested yet in this setup due to incomplete information loop, which is the power dips during the switching which can aggravate short term voltage unbalance and severely increase flickers.

7

Conclusions

This chapter closes the thesis. Section 7.1 answers each of the three research questions in turn. Section 7.2 states the scientific contributions. Section 7.3 gives recommendations for charging facility operators, distribution system operators, and for anyone extending the experimental setup. Section 7.4 delimits the claims, and Section 7.5 identifies the work required to remove those limits.

7.1. Summary of Findings

7.1.1. RQ1: Joint formulation of cost, phase balance and flicker mitigation

How can charging cost minimization, phase-balance, and solar-flicker mitigation be jointly formulated within a single coordinated EV charging framework, subject to public charging facility installation constraints?

The three objectives can be combined, and the formulation that achieves it is a decomposition across two timescales, joined by three couplings that prevent the fast layer from undoing the slow one. Consequently, the objectives prove not to be in fundamental conflict at this scale. The scenario carrying all three attains the same load imbalance as the scenario carrying only cost and balance, while at the same time reducing visible flicker occurrence.

The decomposition follows from the timescales on which the objectives act. The outer loop is a MILP solved every 15 minutes that decides the phase assignment and the base charging current of every connected vehicle, weighing charging cost against a linearized phase-imbalance term and a soft penalty on unmet energy at departure. The inner loop is then an LP solved every 20 seconds that modulates the aggregate charging current in response to the PV fluctuations the outer loop cannot resolve. Feasibility is preserved throughout by expressing the departure energy requirement as a soft penalty rather than a hard constraint. As a result, a scheduling conflict degrades service slowly instead of rendering the problem infeasible.

Decomposing the problem is not by itself sufficient, since two loops acting on the same chargers can work against one another. Three couplings therefore hold them together, each one removing a specific failure mode. First, the available per-phase current headroom is the single physical quantity that cost, balance and flicker all act on, so the installation limits bind every objective simultaneously rather than being enforced separately for each. Second, a proportional anchor term keeps each vehicle near its outer setpoint, and the fast loop consequently corrects around the balanced plan instead of replacing it. Third, and most importantly, the inner loop is restricted to a single common-mode shift applied identically to all three phases, which makes the flicker correction balance-neutral by construction. The price of that restriction is reduced authority, because a limit binding on one phase reduces the correction on all three.

The concern that a fast correction operating on a 20-second timescale might interfere with the balance objective therefore does not materialize, and the sensitivity analysis shows that this holds well beyond the nominal point. Scenario 4 matches Scenario 3 on balance and is moreover cheaper in 21 of the 25 operating conditions tested, despite the inner loop having no ability to revise the plan the outer loop has committed to. The within-window response instead allows the fleet to follow generation that Scenario 3, holding its setpoint flat for the full 15 minutes, cannot react to at all. The advantage does

narrow as PV capacity grows, and above roughly 60 kWp on this installation the connected fleet no longer has the authority to absorb the fluctuation. That limit bounds the applicability of the smoothing layer rather than the formulation itself.

7.1.2. RQ2: Trade-offs and the cost premium of grid-friendly charging

To what extent do charging cost minimization, phase-balance improvement, and solar flicker mitigation trade off against one another when pursued simultaneously, and what cost premium does grid-friendly charging impose on the operator?

The premium is small and centered near zero, provided it is measured against a baseline cost-only optimization with phase changing enabled. It amounts to $+7.5\%$ for Scenario 3 and -4.5% for Scenario 4 at the nominal point, and to $+7.3\%$ and -3.6% respectively on average across the operating grid. The trade-off that motivated the question is therefore far weaker than anticipated. The stronger finding in fact runs the other way, because cost optimization without a phase objective actively degrades power quality. Consequently, the relevant comparison is not what balancing costs the operator, but what its absence costs the grid.

A naive comparison of the balancing scenarios against cost-only optimization suggests that balancing is cheaper, by 10.7% and 18.0% on average across different installation sizing. This is an artifact because, the balancing scenarios differ from the cost-only baseline in two respects rather than one, since they add the balance term but also enable phase switching, which enlarges the feasible set and can only lower the cost minimum. Isolating the balance objective therefore requires a cost-only baseline that shares the phase-switching freedom, denoted Sc2b, against which switching alone is shown to reduce cost from 0.80 to 0.67 €ct/kWh. Measured on that common footing, Scenario 3 carries a median premium of $+2.8\%$ and Scenario 4 a median of -3.5%, being cheaper than Sc2b in 18 of 25 cases.

The fact that Scenario 4 is cheaper than pure cost optimization is explained by PV utilization rather than by any relaxation of the balance objective, and the inner loop contributes in both directions. During supply dips it reduces load, whereas the other scenarios hold their setpoint and must import the difference from the grid, at times when the price is high. During supply spikes it absorbs generation that the others export. Mean utilization consequently rises from 72.8% under Scenario 3 to 74.0% under Scenario 4, higher in 21 of the 25 cases, and the additional self-consumption displaces grid import at a saving of 0.096 €ct/kWh on average. The effect is modest but consistent in direction across the different sizing until EV to PV capacity ratio becomes to low, and it links the economic result directly to the flicker-mitigation mechanism, because the same modulation that smooths the voltage ramp is what captures the generation.

Cost-driven smart charging is not merely neutral with respect to power quality, but actively harmful. Relative to uncontrolled charging, cost-only optimization raises the worst-case load imbalance by 80.3%, from 23.44 A to 42.27 A, because it synchronizes the fleet into the same low-price windows and stacks current on whichever phases those vehicles happen to occupy. At the grid level the instantaneous VUF maximum consequently rises from 1.32% to 2.41%, crossing the 2% value of EN 50160 [16], and 56 annoying-level flicker events appear where uncontrolled charging produced none. Average metrics do improve incidentally, but it is the tail of the distribution that determines compliance.

Adding the balance term reverses this completely. Average imbalance falls from 5.27 A and 4.50 A in the baselines to 1.23 A, worst-case imbalance falls to below 3 A, and the VUF maximum at the facility bus falls to 0.187%, an order of magnitude below the cost-only baseline and below even the baselines' 24-hour averages. The daily VUF profile thus inverts. Whereas in the baselines the charging window is the period of highest unbalance, in the balancing scenarios it becomes the period of lowest, because the connected fleet is the only asset able to compensate the structurally single-phase office load. This result is moreover stable, holding to within 0.1 A of average imbalance and 0.02 percentage points of VUF across all 25 operating conditions tested. The baselines, by contrast, vary strongly with conditions, and the cost-only case exceeds the 2% value in 11 of 25 cases.

7.1.3. RQ3: Hardware-in-the-loop construction and the scope of its feasibility claims

How can the coordinated phase-balancing algorithm and a commercial phase-switching EV charger be jointly tested in a Power Hardware-in-the-Loop environment, and to what extent does the resulting setup demonstrate technical feasibility?

The testbed was built by putting a commercial phase-switching charger in place of the device under test in an existing PHIL setup. It shows that the charger executes the commands and the schedule the optimizer produces, but it does not yet show that the intended grid service is delivered, because the measurement loop was never closed. The answer is therefore a partial yes, and what is missing is a consequence of how the setup was realized rather than a shortcoming of the algorithm.

Construction. The testbed follows the architecture of [6] and [21]. The three-phase FBS load flow runs on an OPAL-RT OP4510 at a 0.1 s timestep, a Cinergia GE-40 amplifier turns the resulting node voltages into real voltages at the charger terminals, a Zaptec Pro [28] serves as the device under test, and a back-to-back Delta Elektronika supply pair with a Walther-Werke signal tester stands in for the on-board charger. A Python setpoint scheduler on the host sends current and phase commands through the Zaptec cloud REST API, keeping time with the real-time simulation clock. Three changes to the intended setup were necessary and are reported as such. First, the emulator follows the schedule directly instead of reading the control pilot. Second, the DC supply pairing caps the load at 1.5 kW, so the optimizer had to be re-run with one vehicle limited to 6 A over a one-hour window. Third, the loop draws power on one phase and returns it on three, so only a third of it circulates and a bidirectional amplifier becomes a requirement rather than a convenience.

What the setup shows. The run establishes four things. First, the real-time model kept its 0.1 s timestep throughout, with no computation overruns, so what follows reflects the behaviour of the hardware and the control path rather than a struggling simulation. Second, every current setpoint and every phase reassignment that was sent was accepted and carried out, with nothing rejected or dropped. The cloud path adds 300–400 ms, which is small against both the 20-second and the 15-minute intervals, so neither the communication route nor the charger limits the optimizer at the speeds it works at. Third, the executed phase trace matches the planned assignment at every timestep of the run, including both phase-switch events. Fourth, the charger's internal certified meter agrees with a separate AC measurement to within 1.3%. That last point matters because this meter is the only measurement the API exposes, so it would have to serve as the feedback signal in any closed-loop version of the setup.

What the setup does not show. Three limits go with those four claims. First, the measurement loop is open. The charger current is recorded but never fed back into the load flow, so the PHIL run and the offline power flow agree because they were built to, not because one confirmed the other, and no claim about VUF or voltage ramps can rest on the hardware. Second, the order of events during a phase switch was imposed rather than tested. The scheduler brings the emulator load to zero before sending the switch command, so the charger always switches with no current flowing. A real installation gives no such guarantee, because the charger interrupts the control pilot and the on-board charger then responds at its own pace. Whether the contactor moves before the current has died away depends on a delay that neither this work nor the literature, including the charger tests of [8], has measured. Third, the run covers one charger on one phase at 1.5 kW for 45 minutes, so nothing here carries over directly to a fleet.

Where the numbers diverge. Over the 45-minute comparison window the optimizer scheduled 0.983 kWh while 0.912 kWh was delivered, a shortfall of 7.3%, and two causes account for it. Both can be corrected in the model. The smaller one is the fixed 230 V assumed by the optimizer against the changing bus voltage of the simulation, which explains 1311 W dispatched against 1283 W measured and would only disappear if the formulation targeted voltage directly, at a considerable cost in computation. The larger one, roughly 70% of the shortfall, is the phase switches themselves. Charging stops for 15 to 90 s each time, depending on the vehicle and charger pairing, and the optimizer does not represent this at all. Both gaps mean the energy and flicker results of Chapter 5 are better than a real installation would achieve, and by a known margin.

7.2. Scientific Contributions

1. A coordinated EV charging formulation that combines dynamic phase assignment, a linear load-balancing proxy, and a fast PV-flicker correction layer under public charging facility installation

constraints. No reviewed study combines phase switching with load balancing and cost optimization. Moreover none addresses phase balance and solar flicker within a single framework.

2. The empirical result that cost-optimal smart charging degrades worst-case phase unbalance and introduces annoying-level flicker that uncontrolled charging does not produce. Price-responsive charging could be a power-quality liability unless the phase dimension enters the objective explicitly.

3. The first PHIL testbed coupling a coordinated phase-balancing optimizer to a commercial phase-switching charger, together with an explicit account of which feasibility claims such a setup can and cannot support, and a specification of the work required to close each gap.

7.3. Recommendations

7.3.1. For public charging facility operators

Cost optimization can be integrated to installation that already supports phase balancing and will provide improve cost, along with other metrics that indirectly improve power quality. The premium measured here is within a few percent of zero and is negative in the majority of tested conditions, because the phase-switching freedom that balancing requires also relaxes the cost problem. Operators evaluating grid-service participation should therefore evaluate the incremental cost of the balance objective against a baseline that already exploits phase switching, not against their current fixed-phase installation, or they will attribute the benefit of switching to the wrong cause.

Operators should also be aware that pure price-following without the phase-changing ability is not a neutral policy. In this case study it produced the worst power-quality outcome of any scenario, worse than no control at all on the metrics that determine compliance.

7.3.2. For distribution system operators and regulators

Because the cost premium of phase balancing is near zero, the incentive required to recruit charging facilities into unbalance mitigation is correspondingly small. The barrier is coordination and hardware capability rather than compensation. The more urgent regulatory observation is that the aggregate effect of uncoordinated price-responsive charging is to synchronize demand, which worsens exactly the tail behavior that EN 50160 [16] assessment is designed to capture. Tariff structures that reward energy shifting without any phase or ramp signal may therefore degrade LV power quality even as they relieve capacity constraints.

Facilities combining EV charging with local PV also warrant attention as flicker sources [22]. Above roughly a 60 kWp PV to eleven-vehicle ratio on this installation, no charging control strategy tested was able to suppress annoying-level ramps, which suggests that the PV-to-controllable-load ratio, rather than PV capacity alone, is the quantity worth monitoring.

7.3.3. For implementation and further experimental work

Three changes to the PHIL setup should precede any further hardware campaign, in this order. Replace the schedule-driven on-board charger emulator with one that decodes the control pilot duty cycle and sets the DC supply current accordingly, so that the emulator reacts to a control-pilot interruption as a real vehicle would rather than relying on the scheduler to impose the correct ordering. Close the measurement path by returning the measured charger current to the load flow, so that the divergence between simulated and hardware-driven grid state can be quantified rather than assumed away. Instrument the control pilot pin to measure the interval between control-pilot interruption and contactor actuation across repeated switching events, since this interval determines the minimum safe ramp-down duration and is not reported anywhere in the literature.

At the optimizer level, one most notable and cheap correction should be made regardless of further hardware work, which is a phase-switching penalty or minimum-dwell constraint in the outer loop.

7.4. Limitations

Scope of the case study. The simulation results rest on a single anonymized Dutch LV feeder, a single summer weekday, and a single charging facility configuration. They demonstrate a mechanism and the direction of its effects; the magnitudes are specific to this network and should not be transferred.

The facility bus is moreover a leaf node on a feeder with no other PV penetration, which isolates the optimizer's response but also maximizes the visibility of its action on the local VUF. On a feeder with distributed PV or with competing unbalanced loads elsewhere, the same current-level improvement would translate into a smaller voltage-level one.

Proxy rather than direct VUF targeting. The optimizer minimizes load imbalance at the facility connection, not VUF at any bus, and the load flow is an external evaluator rather than a constraint. The strong correlation observed here between the two is a property of this network, in which the controllable EV load dominates the local injection. The magnitude of the loss relative to a direct VUF-targeting formulation is not quantified in this work.

Perfect foresight. Arrival and departure times, energy requirements, PV output, office load and day-ahead prices are all treated as exactly known over the scheduling horizon. No forecast error is modelled, and the ElaadNL [25] sessions required departure-time extension to create scheduling slack. Under realistic uncertainty the achievable balance and the zero-shortfall energy delivery would both degrade.

Modelling simplifications. Charging is unidirectional and single-phase only; three-phase charging is excluded because the device under test cannot switch cleanly from three-phase to phases L2 or L3, and V2G is out of scope. The bus voltage is fixed at 230 V in the optimizer, the charger efficiency is constant, and reactive power is not controlled. Phase switching is treated as instantaneous and costless, which Chapter 6 shows it is not. The real transient, in which power falls to zero before recovering [8], will worsen both the short-term unbalance and the flicker figures reported in Chapter 5.

Metric interpretation. The VUF and flicker values reported are instantaneous or 20-second quantities, whereas EN 50160 [16] specifies a 10-minute mean over 95% of a week for VUF and requires sub-second measurement for flicker compliance. Statements that a scenario exceeds or satisfies the 2% value refer to an instantaneous comparison, which is a stricter test than the standard applies, and the flicker occurrences are relative indicators rather than P_{st} or P_{lt} assessments.

Solver conditioning. In 15 of the 25 weighting combinations per scenario, at least one receding-horizon step reached the per-step time limit, with the most affected cells terminating at gaps of 2.7–4.3%. These are concentrated at the small-β, high-ratio corner where the balance term is negligible and the MILP becomes degenerate. The conclusions drawn in Section 5.3 rest on the converged region and on the consistent direction of the trends. The extreme cells should be regarded as partly solver artifacts until recomputed without a time limit. Solve times are not systematically reported, which weakens the computational argument used to justify the proxy formulation.

Hardware validation. As stated under RQ3, the PHIL result covers a single charger at 1.5 kW on a single phase over 45 minutes, with an open feedback path and a schedule-driven rather than control-pilot-driven emulator. No hardware-based claim is made about VUF reduction, flicker, or fleet-level coordination.

7.5. Future Work

Model the switching transient. Adding a switching penalty or minimum-dwell constraint to the outer loop, together with an explicit representation of the zero-power interval, is the single highest-value modelling change. It removes the dominant component of the observed energy discrepancy and allows the flicker results to be re-evaluated against a transient that the current model omits entirely, and which acts in the direction of worsening the reported numbers.

Close the experimental loop. The three testbed changes recommended in Section 7.3 would convert the present executability result into a genuine closed-loop validation, at which point the divergence between simulated and measured grid state becomes measurable rather than structurally zero. Extending the setup to several chargers and to a three-phase-capable DC supply would allow the fleet-level coordination that the optimizer actually assumes to be tested.

Uncertainty and robustness. Replacing perfect foresight with forecast PV, load, arrival and energy-demand distributions, and comparing deterministic receding-horizon control against stochastic or robust formulations, would establish how much of the reported balance quality survives realistic information conditions.

Physical fidelity of the charging model. Incorporating a state-of-charge dependent power factor and efficiency, and the CC-CV taper that the present 90% ceiling approximates, would tighten the link between the scheduled and delivered energy that Chapter 6 shows to be currently loose.

Scaling and market design. Finally, the near-zero cost premium found here invites a study of incentive design: what mechanism, if any, is needed to recruit facilities whose participation costs them almost nothing, and how phase-balancing services would stack with the congestion and reserve products that the same flexible assets are already being asked to provide.

A

Appendix : Simulation Setup Data

This appendix provides the detailed installation, spot allocation, and EV session data underlying the case studies of Chapter 4. The 11-spot configuration (Table A.4) is used for the base-case scenarios (Sc1–Sc4) and the objective weight tuning, while the 13-spot configuration (Table A.3) is used for sensitivity analysis, where two additional EVs are added to Branch-D to test the optimizer.

A.1. Installation and Parking Data

Table A.1: EV Charging Installation Configuration – Site Alpha

Parameter	Value	Unit	Notes
Main Installation			
Main fuse / grid limit	80	A	I_{cap} per phase
Grid voltage (L-N)	230	V	Line-to-neutral RMS
PV system capacity	50	kWp	3-phase balanced inverter
Charging efficiency η	0.95	–	OBC efficiency
I_{min} (IEC 61851)	6	A	Minimum active charging current
I_{max} per EVSE	32	A	Type 2 single-phase limit
Branch / Busbar			
Branch-A	64	A	L1, L2, L3 — 4 spots
Branch-B	64	A	L1, L2, L3 — 4 spots
Branch-C	32	A	L1, L2, L3 — 3 spots
Branch-D*	32	A	L1, L2, L3 — 3 spots

* Branch-D present in the 12 and 13-spot scenario only.

Table A.1 lists the fixed electrical parameters of the charging installation used across all scenarios in Section 4.5, corresponding to the facility described in Section 4.2. Note that the PV capacity listed here (50 kWp) is the value used in the fleet/PV sensitivity sweep of Section 4.6. The branch fuse ratings are each sized below the sum of their downstream EVSE ratings, meaning simultaneous full-power charging at every spot on a branch is not physically permitted by the installation. This constraint is enforced explicitly by the outer MILP (branch constraint, Equation 3.2) and is what necessitates active current sharing among EVs on the same branch even before grid-level phase balancing is considered.

Table A.2 lists the home-phase assignment used for every spot when phase switching is disabled (Sc1 and Sc2, Section 4.5), and serves as the initial/fallback phase in Sc3 and Sc4 before the outer MILP reassigns a connected EV to a different phase. The phases are cyclically rotated across spots within each branch (L1–L2–L3–L1, …) by design, so that the installation is nominally balanced when all spots are occupied and drawing identical current; any imbalance observed in the results therefore originates from the EVs' differing arrival times, dwell durations, and charging power requirements rather than from a structural phase-allocation bias in the installation itself.

Table A.2: Parking Spot Directory – Site Alpha

Spot ID	Branch	Label	Phase	$I_{\max}$ (A)
1	A	A.01	L1	32
2	A	A.02	L2	32
3	A	A.03	L3	32
4	A	A.04	L1	32
5	B	B.01	L2	32
6	B	B.02	L3	32
7	B	B.03	L1	32
8	B	B.04	L3	32
9	C	C.01	L1	32
10	C	C.02	L2	32
11	C	C.03	L3	32
12*	D	D.01	L1	32
13*	D	D.02	L2	32

* Additional spots in the 13-spot scenario (`ev_13`).

Table A.3: EV Charging Session Schedule – 14-Spot Scenario (23 July 2018)

ID	Spot	Ph.	t_{arr} (HH:MM)	t_{dep} (HH:MM)	B_{arr} (kWh)	$B_{\max}$ (kWh)	ΔB_{req} (kWh)	Park (h)	$I_{\max}$ (A)	$\bar{P}_{\min}$ (kW)
1	A.01	L1	07:15	16:25	13.93	40	22.07	9.17	32	2.41
2	A.02	L2	07:45	13:48	9.51	40	24.05	6.05	32	3.98
3	A.03	L3	07:45	14:19	16.87	40	19.13	6.56	32	2.92
4	A.04	L1	08:15	15:47	20.18	40	15.82	7.53	32	2.10
5	B.01	L2	08:15	17:56	18.30	40	17.70	9.69	32	1.83
6	B.02	L3	09:00	15:59	20.13	40	15.87	6.98	32	2.27
7	B.03	L1	09:15	17:54	10.47	40	14.80	8.65	32	1.71
8	B.04	L3	09:30	16:13	10.05	40	25.95	6.72	32	3.86
9	C.01	L1	09:30	17:26	21.24	40	14.76	7.93	32	1.86
10	C.02	L2	09:30	16:58	20.45	40	15.55	7.47	32	2.08
11	C.03	L3	10:15	19:40	21.74	40	9.25	9.41	32	0.98
12	D.01	L1	07:15	16:25	13.93	40	22.07	9.17	32	2.41
13	D.02	L2	09:00	15:59	20.13	40	15.87	6.98	32	2.27

Table A.3 presents the ElaadNL-derived arrival/departure schedule (Section 4.4) after the departure-time extension and load-derived adjustments discussed in Section 4.4. The last column, $\bar{P}_{\min}$, is the minimum constant charging power required for that vehicle to reach its SOC target if it charged continuously and uninterrupted from arrival to departure; comparing this to the branch and installation fuse limits gives an immediate indication of which sessions are load-constrained rather than schedule-constrained. All combinations in the sensitivity analysis use the same table with differing numbers of EVs corresponding to the EV index. Thus, if the case is using 9 EVs, EV 1 to 9 are used in the simulation.

Table A.4 is a modification to accommodate hardware test in Chapter 6. EV 11 is notable for its uncharacteristically low $I_{\max} = 6\,\text{A}$, representing a vehicle limited to the IEC 61851 minimum current.This is because of the current limitation of the setup using power supply with the limit of 1.5kW for the OBC.

A.2. Test Network

The network originates from an anonymized GNF (Gaia Network File) export supplied by Stedin, the Dutch DSO responsible for the region in which the feeder is located. The GNF format is Stedin's internal exchange format for representing LV and MV network topology, containing node coordinates, cable/link segment records, switch states, and connected transformer and load metadata. Anonymization removes any information that could identify the specific street or customers served, retaining only the electrical topology, cable parameters, and aggregate connection data needed for simulation. Consequently, the feeder cannot be geographically located from the data used in this thesis, and node

Table A.4: EV Charging Session Schedule – 11-Spot Scenario

ID	Spot	Ph.	t_{arr} (HH:MM)	t_{dep} (HH:MM)	B_{arr} (kWh)	B_{max} (kWh)	ΔB_{req} (kWh)	Park (h)	I_{max} (A)	$\bar{P}_{\text{min}}$ (kW)
1	A.01	L1	07:15	16:25	13.93	40	22.07	9.17	32	2.41
2	A.02	L2	07:45	13:48	9.51	40	24.05	6.05	32	3.98
3	A.03	L3	07:45	14:19	16.87	40	19.13	6.56	32	2.92
4	A.04	L1	08:15	15:47	20.18	40	15.82	7.53	32	2.10
5	B.01	L2	08:15	17:56	18.30	40	17.70	9.69	32	1.83
6	B.02	L3	09:00	15:59	20.13	40	15.87	6.98	32	2.27
7	B.03	L1	09:15	17:54	10.47	40	14.80	8.65	32	1.71
8	B.04	L3	09:30	16:13	10.05	40	25.95	6.72	32	3.86
9	C.01	L1	09:30	17:26	21.24	40	14.76	7.93	32	1.86
10	C.02	L2	09:30	16:58	20.45	40	15.55	7.47	32	2.08
11	C.03	L3	10:15	19:40	21.74	40	9.25	9.41	6	0.98

identifiers are arbitrary integers assigned by the export process rather than meaningful addresses.

Because the raw GNF file encodes switch states rather than a pre-resolved radial topology, a parsing step is required to determine the active operating configuration of the feeder before it can be used in the load-flow model. Link segments represent switchable connections (fuses, disconnectors, or normally-open tie switches) that may or may not be in service at a given time. For this network, all link segments were found to be closed, and no normally-open tie switches were present, so no meshing or reconfiguration ambiguity needed to be resolved. This confirms the feeder is genuinely and permanently radial in its as-operated state, rather than radial only because open switches happen to isolate a subset of a more meshed underlying network — a distinction that matters for the applicability of the forward-backward sweep load-flow method, which assumes strict radiality.

The absence of any existing PV connections in this particular feeder, noted in Chapter 4 as the reason for its selection indicates that at the time the network snapshot was taken, none of the 214 connected customers on this feeder had registered rooftop PV with the DSO. Another version of the grid with 2030 forecast is also available for further study into the grid with high PV penetration.

Table A.5: Low-voltage network characteristics

Parameter	Value
Grid topology	Fully radial
Transformer rating	400 kVA
Transformer voltage level	10.75/0.42 kV
Transformer vector group	DYN5
Slack bus location	LV terminal of transformer
Slack bus voltage	230 V (phase-to-neutral)
Number of LV nodes	213
Number of branches	212
Minimum branch length	< 1 m
Maximum branch length	102 m
Median branch length	5.8 m
Dominant cable resistance	0.369 Ω/km
Dominant cable reactance	0.841 Ω/km
High-resistance segments	1.20 Ω/km (feeder ends)

B

Appendix B : Simulation Results Data

This appendix provides more detailed figures that might be required to provide more context of the results analyzed on Chapter 5.

B.1. Voltage comparison of main simulation

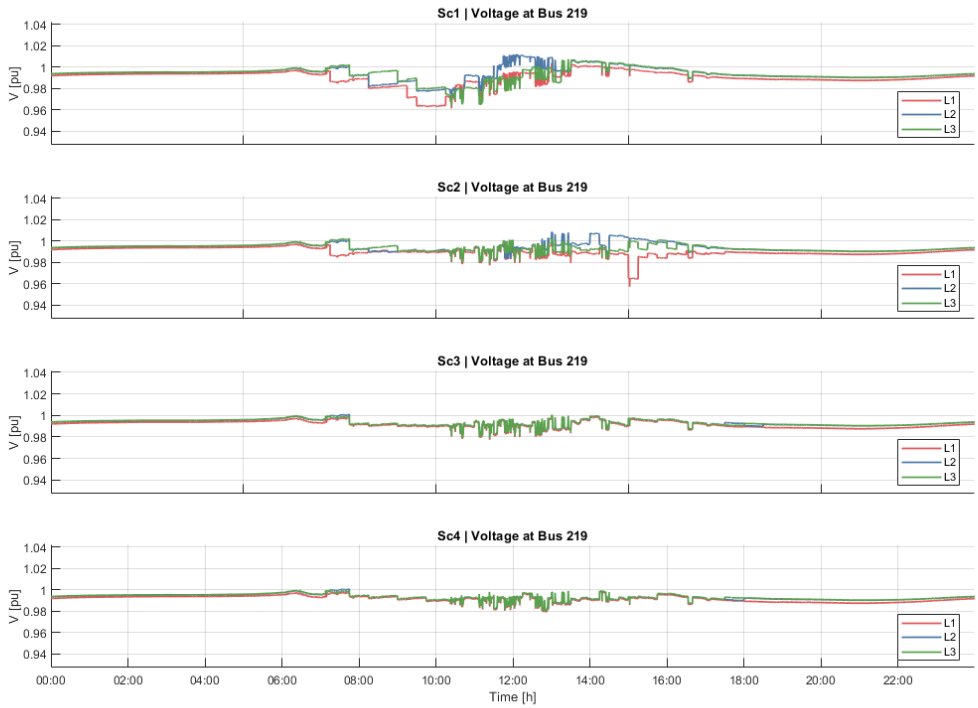

Figure B.1: Voltage profile comparison of the 4 scenarios in main simulation

B.2. EV Charging Load Profile Effect Towards Charging Cost

Subsection 5.2.4 identified that the load adjustment towards PV fluctuations by the inner loop reduces unnecessary/unscheduled power import and increase from the grid. Figure B.2 and B.3 show the responsiveness of Sc4 which enables the cost reduction in the main simulation.

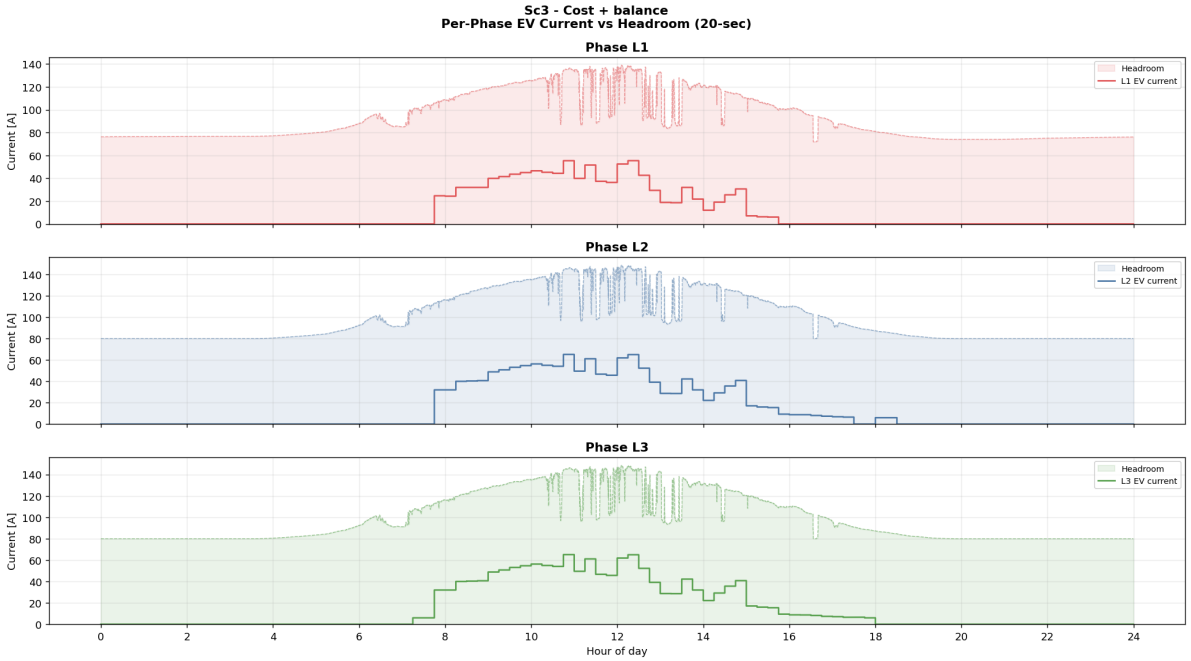

Figure B.2: Total EV charging power Sc3 : 11 EVs PV capacity 50kWp.

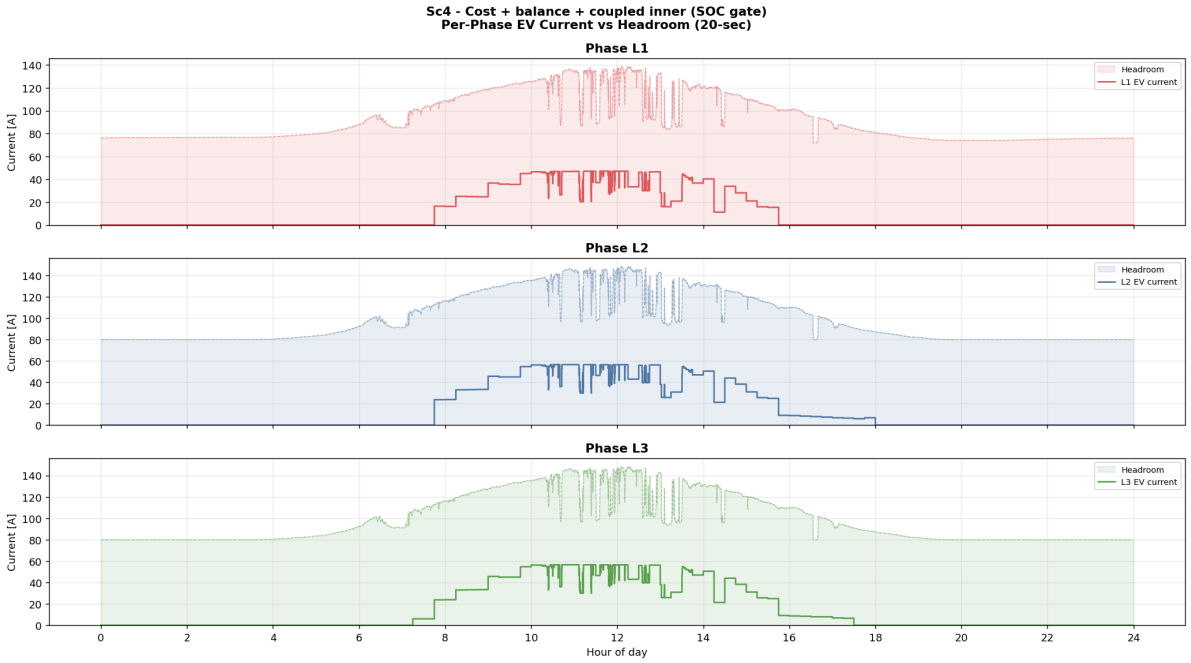

Figure B.3: Total EV charging power Sc4 : 11 EVs PV capacity 50kWp.

Continuing the finding from Subsection 5.2.4, Subsection 5.4.1 reaffirm the robustness of Sc4 to a certain point. The cost reduction flipped at low EV to PV ratio, which is attributed to the inability of the inner to adjust to fluctuation at that region. The compounding effect that flatten the outer loop is so far unidentified. Comparison between figure B.4 and B.5 shows the passiveness of inner loop and flattening of outer loop.

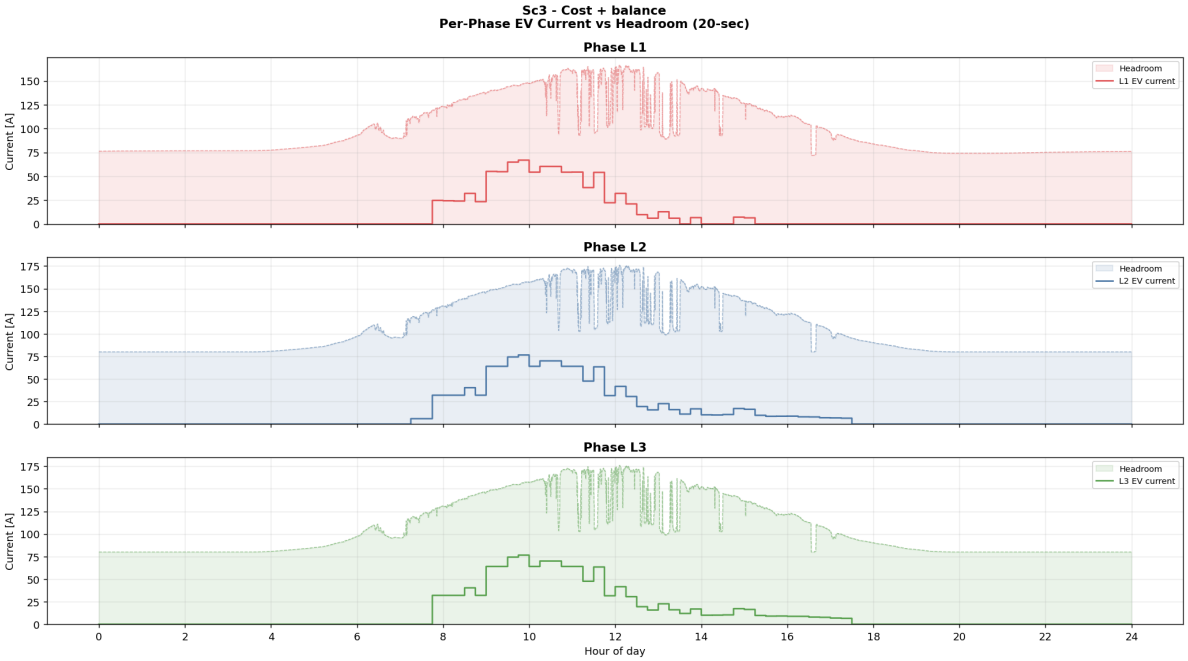

Figure B.4: Total EV charging power Sc3 : 9 EVs PV capacity 70kWp.

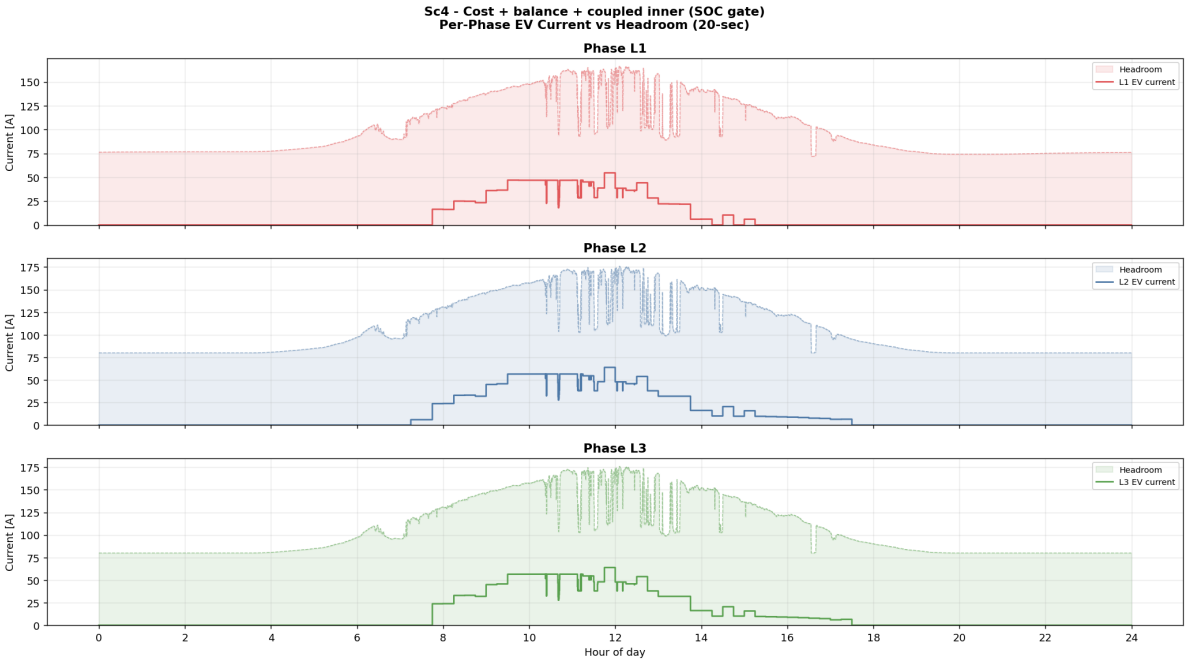

Figure B.5: Total EV charging power Sc4 : 9 EVs PV capacity 70kWp.

C

Appendix : PHIL Testbed Equipment

This appendix documents the physical equipment used in the Power Hardware-in-the-Loop setup of Section 6.1. The components are presented in the order in which they appear in the signal and power path: the real-time simulator that computes the grid state, the amplifier that realizes it as physical voltage, the charger under test, and the emulator that presents a load to that charger.



Figure C.1: Overall view of the PHIL testbed.

Figure C.1 shows the assembled testbed. The controller, the simulated grid, the charger and the emulator are physically separate but operate on a shared simulation clock: the setpoint scheduler reads the timestamp from the real-time model and dispatches commands against it, so that the hardware executes the optimizer schedule in simulated rather than wall-clock time.

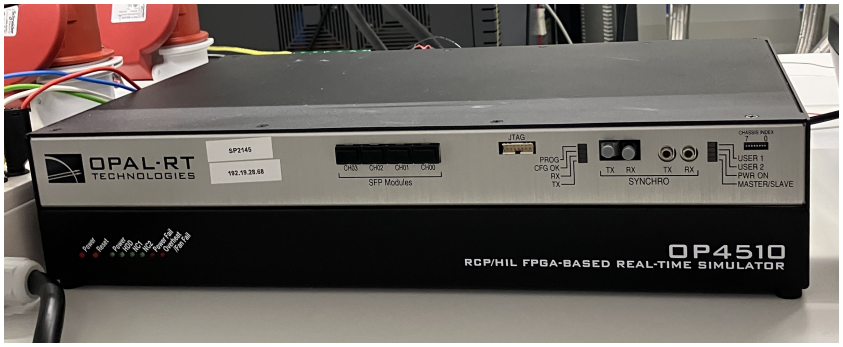

Figure C.2: OPAL-RT digital real-time simulator running the three-phase load flow.

The distribution grid model is built in MATLAB/Simulink on the host computer and compiled to the real-time target through RT-LAB. Once deployed, the simulator solves the three-phase forward-backward sweep load flow at a fixed timestep of $0.1\,\mathrm{s}$ and outputs the resulting per-phase node voltages as analog references to the amplifier. The timestep was met for the entire duration of every run reported in Chapter 6, with no computation overruns recorded by RT-LAB, which is the precondition for attributing the observed behavior to the hardware rather than to degradation of the simulation itself.

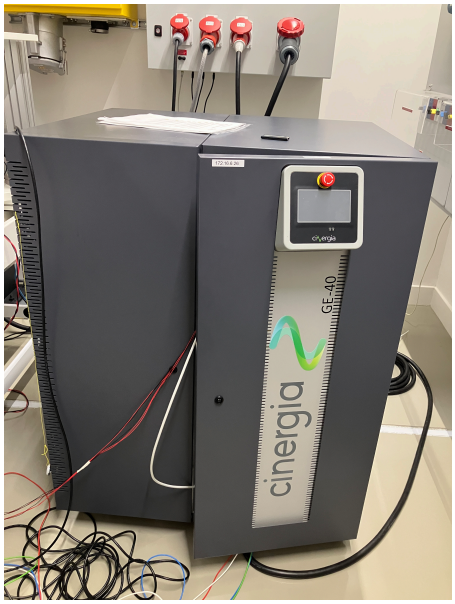

Figure C.3: Cinergia GE-40 power amplifier.

The amplifier converts the per-phase voltage references computed by the simulator into physical three-phase voltages at the charger terminals. A single unit translates the three analog references simultaneously, rather than requiring one amplifier per phase, and supports bidirectional power flow. The latter is a requirement rather than a convenience in this setup: because the emulator draws single-phase power and returns it three-phase, only a third of the circulated power closes the loop and the remainder is exported back to the grid through the amplifier.

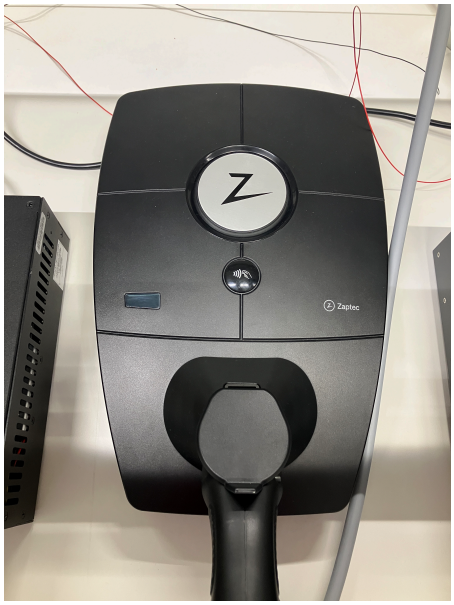

Figure C.4: Zaptec Pro charger, the device under test.

The Zaptec Pro is a single-outlet, wall-mounted AC charger rated at $22\,\mathrm{kW}$ (three-phase, $32\,\mathrm{A}$ per phase) with an IEC 62196 Type 2 socket. It is one of the few commercially available units with native phase-switching capability, which is what makes it the natural device under test for a framework whose central decision variable is the per-EV phase assignment. Charging follows the IEC 61851-1 control pilot protocol: the charger generates a $1\,\mathrm{kHz}$ PWM signal whose duty cycle encodes the maximum permitted current. Its current modulation resolution of $0.1\,\mathrm{A}$ is substantially finer than the $\pm 0.6\,\mathrm{A}$ tolerance permitted by the standard, and the PWM update latency from command receipt to control pilot change is approximately $400\,\mathrm{ms}$. Current and phase commands are issued remotely through the cloud-accessible REST API. An internal certified meter logs the delivered current, and was found to agree with an independent AC measurement to within 1.3%.

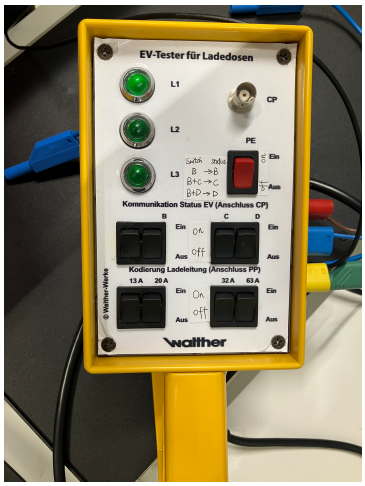

Figure C.5: Walther-Werke EV signal tester.

The charger will not energize its output without a valid control pilot handshake from a connected vehicle. The signal tester supplies that handshake, presenting the ready-to-charge state on the control pilot line so that the charger closes its contactor and the emulator behind it can draw power. In the realized setup the tester establishes the handshake only. it does not decode the duty cycle, so the charging current is set by the scheduler rather than derived from the control pilot signal.



(a) SM70-22



(b) SM1500-30

Figure C.6: Back-to-back Delta Elektronika supply pair forming the on-board charger emulator and power circulation unit.

The emulator stands in for the vehicle on-board charger. The first unit, a Delta Elektronika SM70-22, receives single-phase AC from the charger output and presents the controllable DC load; its rating limits the emulated charging power to $1.5\,\mathrm{kW}$, which is the constraint that required the PHIL scenario of Section 6.2.1 to be re-run with one EV capped at $6\,\mathrm{A}$. The second unit, a Delta Elektronika SM1500-30, returns the power three-phase to a terminal downstream of the amplifier, forming the circulating power routing that avoids dissipating the test energy. The load setpoint is sent to the supply over Ethernet.

Bibliography

[1] I. IEA, “Global EV Outlook 2025,” International Energy Agency, Paris, Technology Report, 2025. [Online]. Available: https://www.iea.org/reports/global-ev-outlook-2025

[2] “Capaciteitskaart elektriciteitsnet | Regionale netbeheerders - Afname.” [Online]. Available: https://data.partnersinenergie.nl/capaciteitskaart/rnb/afname

[3] N. Jabalameli and A. Ghosh, “Online Centralized Coordination of Charging and Phase Switching of PEVs in Unbalanced LV Networks With High PV Penetrations,” *IEEE Systems Journal*, vol. 15, no. 1, pp. 1015–1025, Mar. 2021. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/9122465/citations

[4] S. H. Soltani, M. Rashidinejad, and A. Abdollahi, “Dynamic phase balancing in the smart distribution networks,” *International Journal of Electrical Power & Energy Systems*, vol. 93, pp. 374–383, Dec. 2017. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0142061516308870

[5] S. Martinenas, K. Knezović, and M. Marinelli, “Management of Power Quality Issues in Low Voltage Networks Using Electric Vehicles: Experimental Validation,” *IEEE Transactions on Power Delivery*, vol. 32, no. 2, pp. 971–979, Apr. 2017. [Online]. Available: https://ieeexplore.ieee.org/document/7579545/

[6] L. De Herdt, A. Shekhar, Y. Yu, G. R. C. Mouli, J. Dong, and P. Bauer, “Power Hardware-in-the-Loop Demonstrator for Electric Vehicle Charging in Distribution Grids,” in *2021 IEEE Transportation Electrification Conference & Expo (ITEC)*, Jun. 2021, pp. 679–683, iSSN: 2377-5483. [Online]. Available: https://ieeexplore.ieee.org/document/9490098/

[7] O. Gehrke, C. Ziras, K. Heussen, T. V. Jensen, and H. W. Bindner, “Testbeds For Active Distribution Networks: Case Experience from Syslab,” in *CIRED 2021 - The 26th International Conference and Exhibition on Electricity Distribution*, vol. 2021, Sep. 2021, pp. 2299–2303. [Online]. Available: https://ieeexplore.ieee.org/document/9692506

[8] K. Sevdari, S. Striani, P. B. Andersen, and M. Marinelli, “Power Modulation and Phase Switching Testing of Smart Charger and Electric Vehicle Pairs,” in *2022 57th International Universities Power Engineering Conference (UPEC)*, Aug. 2022, pp. 1–6. [Online]. Available: https://ieeexplore.ieee.org/document/9917660/

[9] S.-H. Hwang, Y. Chen, H. Zhang, K.-Y. Lee, and D.-H. Kim, “Reconfigurable Hybrid Resonant Topology for Constant Current/Voltage Wireless Power Transfer of Electric Vehicles,” *Electronics*, vol. 9, no. 8, p. 1323, Aug. 2020. [Online]. Available: https://www.mdpi.com/2079-9292/9/8/1323

[10] “Transparency Platform.” [Online]. Available: https://transparency.entsoe.eu/market/energyPrices

[11] Y. Wesseling, N. Damianakis, G. Ram Chandra Mouli, and P. Bauer, “Grid Impact of Unbalanced Phase Integration of PV Generation, Electrified Mobility, and Heating in LV Distribution Grids,” *IEEE Access*, vol. 13, pp. 56 887–56 907, 2025. [Online]. Available: https://ieeexplore.ieee.org/document/10942319

[12] C. Fetting, “The European Green Deal,” ESDN Office, Vienna, Tech. Rep., Dec. 2020.

[13] “Regulation (EU) 2023/851 of the European Parliament and of the Council of 19 April 2023 amending Regulation (EU) 2019/631 as regards strengthening the CO2 emission performance standards for new passenger cars and new light commercial vehicles in line with the Union’s increased climate ambition (Text with EEA relevance),” Apr. 2023, legislative Body: CONSIL, EP. [Online]. Available: http://data.europa.eu/eli/reg/2023/851/oj

[14] L. Cozzi, “The Future of Heat Pumps,” International Energy Agency, Stockholm, Tech. Rep., May 2023. [Online]. Available: https://heatpumpingtechnologies.org/content/uploads/2023/06/20230525cozziheatpumps-v4.pdf

[15] E. De Winkel, Z. Lukszo, M. Neerincx, and R. Dobbe, “Adapting to limited grid capacity: Perceptions of injustice emerging from grid congestion in the Netherlands,” *Energy Research & Social Science*, vol. 122, p. 103962, Apr. 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S221462962500043X

[16] “EN 50160: Voltage Characteristics of Electricity Supplied by Public Electricity Networks,” European Committee for Electrotechnical Standardization (CENELEC), Brussels, European Standard, 2010.

[17] M. T. Bina and A. Kashefi, “Three-phase unbalance of distribution systems: Complementary analysis and experimental case study,” *International Journal of Electrical Power & Energy Systems*, vol. 33, no. 4, pp. 817–826, May 2011. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S014206151000205X

[18] N. Kongtrakul, W. Wangdee, and S. Chantaraskul, “Comprehensive review and a novel technique on voltage unbalance compensation,” *IET Smart Grid*, vol. 6, no. 4, pp. 331–358, Aug. 2023. [Online]. Available: https://ietresearch.onlinelibrary.wiley.com/doi/full/10.1049/stg2.12106

[19] S. Weckx and J. Driesen, “Load Balancing With EV Chargers and PV Inverters in Unbalanced Distribution Grids,” *IEEE Transactions on Sustainable Energy*, vol. 6, no. 2, pp. 635–643, Apr. 2015. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/7057546

[20] S. V. Saleh and M. A. Latify, “Coordinated Unbalance Compensation and Harmonic Mitigation in the Secondary Distribution Network Through EVs Participation,” *IEEE Transactions on Smart Grid*, vol. 15, no. 5, pp. 4450–4462, Sep. 2024. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/10445782

[21] Y. Yu, L. De Herdt, A. Shekhar, G. R. C. Mouli, and P. Bauer, “EV Smart Charging in Distribution Grids–Experimental Evaluation Using Hardware in the Loop Setup,” *IEEE Open Journal of the Industrial Electronics Society*, vol. 5, pp. 13–27, 2024. [Online]. Available: https://ieeexplore.ieee.org/document/10387706/

[22] N. Brinkel, M. Gerritsma, T. AlSkaif, I. Lampropoulos, A. Van Voorden, H. Fidder, and W. Van Sark, “Impact of rapid PV fluctuations on power quality in the low-voltage grid and mitigation strategies using electric vehicles,” *International Journal of Electrical Power & Energy Systems*, vol. 118, p. 105741, Jun. 2020. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0142061519319994

[23] “Global Solar Data | Meteonorm.” [Online]. Available: https://meteonorm.com/

[24] “Profielen Elektriciteit,” 2018. [Online]. Available: https://web.archive.org/web/20181202215314/https://www.nedu.nl/documenten/verbruiksprofielen/

[25] “Charging profile simulation.” [Online]. Available: https://charging.elaad.nl/

[26] K. Knezović and M. Marinelli, “Phase-wise enhanced voltage support from electric vehicles in a Danish low-voltage distribution grid,” *Electric Power Systems Research*, vol. 140, pp. 274–283, Nov. 2016. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0378779616302231

[27] E. Akhavan-Rezai, M. F. Shaaban, E. F. El-Saadany, and F. Karray, “Managing Demand for Plug-in Electric Vehicles in Unbalanced LV Systems With Photovoltaics,” *IEEE Transactions on Industrial Informatics*, vol. 13, no. 3, pp. 1057–1067, Jun. 2017. [Online]. Available: https://ieeexplore.ieee.org/document/7865962/

[28] “Zaptec Developer Guides.” [Online]. Available: https://docs.zaptec.com/

[29] J. F. Franco, M. J. Rider, and R. Romero, “A Mixed-Integer Linear Programming Model for the Electric Vehicle Charging Coordination Problem in Unbalanced Electrical Distribution Systems,” *IEEE Transactions on Smart Grid*, vol. 6, no. 5, pp. 2200–2210, Sep. 2015. [Online]. Available: https://ieeexplore.ieee.org/document/7031457/

[30] M. A. Fotouhi Ghazvini, G. Lipari, M. Pau, F. Ponci, A. Monti, J. Soares, R. Castro, and Z. Vale, “Congestion management in active distribution networks through demand response implementation,” *Sustainable Energy, Grids and Networks*, vol. 17, p. 100185, Mar. 2019. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S2352467718300699

[31] F. Shen, S. Huang, Q. Wu, S. Repo, Y. Xu, and J. Østergaard, “Comprehensive Congestion Management for Distribution Networks Based on Dynamic Tariff, Reconfiguration, and Re-Profiling Product,” *IEEE Transactions on Smart Grid*, vol. 10, no. 5, pp. 4795–4805, Sep. 2019. [Online]. Available: https://ieeexplore.ieee.org/document/8454501/

[32] S. Shivashankar, S. Mekhilef, H. Mokhlis, and M. Karimi, “Mitigating methods of power fluctuation of photovoltaic (PV) sources – A review,” *Renewable and Sustainable Energy Reviews*, vol. 59, pp. 1170–1184, Jun. 2016. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S1364032116000897

[33] M. Abdel-Basset, L. Abdel-Fatah, and A. K. Sangaiah, “Metaheuristic Algorithms: A Comprehensive Review,” in *Computational Intelligence for Multimedia Big Data on the Cloud with Engineering Applications*. Elsevier, 2018, pp. 185–231. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/B9780128133149000104

[34] K. Girigoudar and L. A. Roald, “Linearized Three-Phase Optimal Power Flow Models for Distribution Grids with Voltage Unbalance,” in *2021 60th IEEE Conference on Decision and Control (CDC)*, Dec. 2021, pp. 4214–4221, iSSN: 2576-2370. [Online]. Available: https://ieeexplore.ieee.org/document/9683780/

[35] Y. Wu, Z. Wang, Y. Huangfu, A. Ravey, D. Chrenko, and F. Gao, “Hierarchical Operation of Electric Vehicle Charging Station in Smart Grid Integration Applications — An Overview,” *International Journal of Electrical Power & Energy Systems*, vol. 139, p. 108005, Jul. 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0142061522000503

[36] H. Liang, B. J. Choi, W. Zhuang, and X. Shen, “Stability Enhancement of Decentralized Inverter Control Through Wireless Communications in Microgrids,” *IEEE Transactions on Smart Grid*, vol. 4, no. 1, pp. 321–331, Mar. 2013. [Online]. Available: https://ieeexplore.ieee.org/document/6451180/

[37] M. Restrepo, J. Morris, M. Kazerani, and C. A. Cañizares, “Modeling and Testing of a Bidirectional Smart Charger for Distribution System EV Integration,” *IEEE Transactions on Smart Grid*, vol. 9, no. 1, pp. 152–162, Jan. 2018. [Online]. Available: https://ieeexplore.ieee.org/document/7442180/

[38] G. F. Lauss, M. O. Faruque, K. Schoder, C. Dufour, A. Viehweider, and J. Langston, “Characteristics and Design of Power Hardware-in-the-Loop Simulations for Electrical Power Systems,” *IEEE Transactions on Industrial Electronics*, vol. 63, no. 1, pp. 406–417, Jan. 2016. [Online]. Available: https://ieeexplore.ieee.org/document/7177085/

[39] “IEC 61851-1:2017.” [Online]. Available: https://webstore.iec.ch/en/publication/33644

[40] N. H. Refa, Nazir, “Impact of Smart Charging on EVs Charging Behaviour Assessed from Real Charging Events,” Lyon, France, May 2019. [Online]. Available: https://www.researchgate.net/profile/Nazir-Refa/publication/333262496_Impact_of_Smart_Charging_on_EVs_Charging_Behaviour_Assessed_from_Real_Charging_Events/links/5ce4ff7b458515712eba7809/Impact-of-Smart-Charging-on-EVs-Charging-Behaviour-Assessed-from-Real-Charging-Events.pdf

[41] J. R. Carson, “Wave propagation in overhead wires with ground return,” *The Bell System Technical Journal*, vol. 5, no. 4, pp. 539–554, Oct. 1926. [Online]. Available: https://ieeexplore.ieee.org/document/6768096/

[42] W. Kersting and W. Phillips, “Distribution feeder line models,” in *Proceedings of 1994 IEEE Rural Electric Power Conference*, Apr. 1994, pp. A4/1–A4/8. [Online]. Available: https://ieeexplore.ieee.org/document/326257/

[43] K. Zhao, P. Ciufo, and S. Perera, “Rectifier Capacitor Filter Stress Analysis When Subject to Regular Voltage Fluctuations,” *IEEE Transactions on Power Electronics*, vol. 28, no. 7, pp. 3627–3635, Jul. 2013. [Online]. Available: https://ieeexplore.ieee.org/document/6357322/

[44] A. Wilkins, J. Veitch, and B. Lehman, “LED lighting flicker and potential health concerns: IEEE standard PAR1789 update,” in *2010 IEEE Energy Conversion Congress and Exposition*, Sep. 2010, pp. 171–178, iSSN: 2329-3748. [Online]. Available: https://ieeexplore.ieee.org/document/5618050/

[45] N. Damianakis, G. R. C. Mouli, P. Bauer, and Y. Yu, “Assessing the grid impact of Electric Vehicles, Heat Pumps & PV generation in Dutch LV distribution grids,” *Applied Energy*, vol. 352, p. 121878, Dec. 2023. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0306261923012424

[46] A. Smets, K. Jäger, O. Isabella, R. v. Swaaij, and M. Zeman, *Solar Energy: The physics and engineering of photovoltaic conversion, technologies and systems*. Bloomsbury Publishing, Jan. 2016, google-Books-ID: gj7JEAAAQBAJ.

[47] W. B. Mol, W. H. Knap, and C. C. v. Heerwaarden, “High resolution solar irradiance variability climatology dataset part 2: classifications, supplementary data, and statistics,” Sep. 2022. [Online]. Available: https://zenodo.org/records/7092058