DTU Wind and Energy Systems

Coordinated control of EV smart chargers through microcontroller integration

Carl Aggernæs Thornberg &
Jonathan Vincents Hellemann Eriksen

DTU Wind-M-0912 June 2025

Author:
Carl Aggernæs Thornberg &
Jonathan Vincents Hellemann Eriksen

Title:
Coordinated control of EV smart chargers through microcontroller integration

DTU Wind & Energy Systems is a department of the Technical University of Denmark with a unique integration of research, education, innovation and public/private sector consulting in the field of wind energy. Our activities develop new opportunities and technology for the global and Danish exploitation of wind energy. Research focuses on key technical-scientific fields, which are central for the development, innovation and use of wind energy and provides the basis for advanced education at the education.

DTU Wind-M-0912
June 2025

ECTS: 30

Education: Master of Science

Supervisor(s):

Pietro Zunino
Simone Striani
Mattia Marinelli
Jan Engelhardt

DTU Wind & Energy Systems

Remarks:
This report is submitted as partial fulfillment of the requirements for graduation in the above education at the Technical University of Denmark.

Technical University of Denmark
Department of Wind and Energy Systems
Frederiksborgvej 399
4000 Roskilde
Denmark
www.wind.dtu.dk

Coordinated control of EV smart chargers through microcontroller integration

Master Thesis
June 10, 2025

By
Carl Aggernæs Thornberg
Jonathan Vincents Helleman Eriksen

Cover photo: WIND, DTU
Published by: DTU, Department of Wind and Energy Systems, Frederiksborgvej 399, 4000 Roskilde Denmark
https://wind.dtu.dk/
ISSN: [N/A] (electronic version)
ISBN: [N/A] (electronic version)

ISSN: [N/A] (printed version)
ISBN: [N/A] (printed version)

Approval

This thesis has been prepared over five months at the Section for E-mobility and Prosumer Integration, Department of Wind and Energy Systems, at the Technical University of Denmark, DTU, in partial fulfillment for the degree Master of Science in Engineering.

Carl Aggernæs Thornberg - s185825
Jonathan Vincents Helleman Eriksen - s193875

Signature *Signature*

10-06-2025

Date

Abstract

To meet the objectives set forth by the Paris Agreement, substantial decarbonization of the energy sector through integration of renewable energy sources (RES) is crucial. This shift presents significant challenges, notably in managing intermittent energy generation, diminished grid inertia, and reduced flexibility. Simultaneously, the rapid uptake of electric vehicles (EVs) presents both challenges, in terms of heightened grid stress during peak hours, and opportunities through the potential for EVs to provide grid-supporting flexibility via smart charging and vehicle-to-grid (V2G) technologies.

This thesis builds upon findings from the DTU ACDC project, which successfully demonstrated autonomous smart charging, though with certain limitations, notably delays due to externally implemented controllers. To address these limitations, this thesis introduces an integrated charging solution featuring a BeagleBone Black (BBB) microcontroller embedded within EV chargers, enabling direct local measurement acquisition, setpoint dispatching, and bidirectional communication among chargers through the Energydata.dk platform.

The main contributions include the design and implementation of autonomous decision-making algorithms specifically tailored to fulfill four primary objectives: ensuring grid compliance at the point of common coupling, efficiently charging vehicles according to dynamic user priorities, aligning consumption with day-ahead spot market bids, and facilitating participation in frequency containment reserve markets.

Physical testing was conducted at DTU's Lyngby and Risø campuses, involving four chargers equipped with identical hardware setups comprising the BBB microcontroller, DEIF measurement unit, Phoenix Contact charge controller and modem all connected in a local network via a switch. Results confirmed significant improvements in response times, with upregulation frequency control activations averaging 840 ms, meeting the activation requirement of 1300 ms for Denmark's fastest reserve, Fast Frequency Reserve (FFR). Preliminary tests for downregulation also indicated compliance with the required reaction times for Frequency Containment Reserve Down (FCR-D). Additionally, system performance analysis revealed that control delays are primarily caused by physical component-induced delays, and that loop time scales linearly with the number of chargers.

Acknowledgements

We would like to thank our team of supervisors for involving us in their engaging research and for framing a project that was both full of potential and rich in learning opportunities. Throughout the process, their ongoing feedback and insightful discussions were a great help in shaping our work. A special thanks to the two PhD students: Co-supervisor Pietro Zunino, as well as to our unofficial sparring partner Kristoffer Laust Pedersen, for their valuable input and support along the way.

We would also like to extend a big thanks to the PLK technicians Magnus Fich Rabjerg, Adam Juel Madsen, and Dennis Tryk. Thank you for letting us occupy your workspace for five months, for patiently answering all our—at times—less than brilliant questions, and for your vital contribution to the hardware setup.

A special thanks as well to Tom Luc Bailly-Salins for his patience and invaluable help in configuring the BeagleBone Blacks and their network connections.

Table of Contents

Figures

Tables

List of Abbreviations

Abbreviation	**Full Term**	**Explanation**
AC	Alternating Current	Electrons switches direction back and forth at regular intervals. The type of electrical current used for charging in this project.
ACDC	Autonomously Controlled Distributed Chargers	A DTU research project on distributed EV charging.
aFRR	Automatic Frequency Restoration Reserve	Ancillary service with moderate speed and energy demand, used to restore frequency to 50 Hz.
BBB	BeagleBone Black	Microcontroller used in each charger to perform control calculations locally.
BMS	Battery Management System	Internal EV system responsible for battery safety and control.
BRP	Balancing Responsible Party	Market actor responsible for balancing energy consumption and production.
BSP	Balancing Service Provider	Entity offering flexibility to the TSO by delivering reserve capacity or balancing energy.
CC	Charge Controller	Phoenix Contact device controlling charging via PWM communication with the EV.
CP	Control Pilot	Communication line carrying PWM signals between charger and EV.
CPU	Central Processing Unit	Primary component of a computer that runs the machine's operating system and apps.
DC	Direct Current	One-directional flow of electric charge.
DEIF	Electrical Multimeter	Installed at each charger to measure voltage, current, and frequency.

Abbreviation	**Full Term**	**Explanation**
DIP	Dual In-line Package	A set of small manual electronic switches that are designed to be packaged with other circuits.
EDDK	EnergyData.dk	Real-time cloud-based data platform for communication.
EV	Electric Vehicle	The charging subject in this study.
EVSE	Electric Vehicle Supply Equipment	All charger hardware components between the power source and EV.
FCR-D	Frequency Containment Reserve – Disturbance	Ancillary service requiring fast power reactions during frequency events outside a dead-band.
FCR-N	Frequency Containment Reserve – Normal	Ancillary service requiring continuous frequency-sensitive operation.
FFR	Fast Frequency Reserve	Ancillary service requiring very fast, temporary frequency-based power reduction.
IEC	International Electrotechnical Commission	International standards for all electrical, electronic and related technologies.
IoT	Internet of Things	System of connected devices sharing data, including chargers in this setup.
IP-address	Internet Protocol address	Unique number linked to network activity.
Mac	MacBook Pro M3	Used to compare computation times with local microcontrollers.
mFRR	Manual Frequency Restoration Reserve	Ancillary service with slow activation for frequency recovery.
MQTT	Message Queuing Telemetry Transport	A lightweight, publish-subscribe messaging protocol designed for IoT.
OBC	Onboard Charger	Converts AC power to DC power inside the EV for charging.

Abbreviation	**Full Term**	**Explanation**
P_IoT	Power IoT Device	Includes a BBB and DEIF used for publishing electrical measurements to EDDK.
PCC	Point of Common Coupling	Connection point to the external grid, typically with limited current capacity.
PWM	Pulse Width Modulation	Method used to encode current setpoints in the charging control signal.
PX	Proximity Plug	Cable used by the charge controller to detect cable current rating.
RES	Renewable Energy Sources	Intermittent power sources such as wind and solar.
SSH	Secure Shell	Cryptographic network protocol for operating network services securely over an unsecured network.
SOC	State of Charge	Percentage of energy stored in the EV battery.
TCP	Transmission Control Protocol	Transmission medium for Modbus TCP/IP messaging
TSO	Transmission System Operator	Responsible for maintaining grid stability and activating balancing services.
VPN	Virtual Private Network	Encrypted tunnel between a device and the internet.
V2G	Vehicle-to-Grid	Operation mode where EVs return power to the grid.
ϕ	Phases	Number of electrical phases used by an EV for charging.
ρ	Priority Metric	Key variable used to determine an EV's eligibility to charge based on need and system constraints.

1 Introduction

In the process of complying with the Paris Agreement on holding the increase in the global average temperature to below 2°C above pre-industrial levels[1], the decarbonization of the energy sector is inevitable. In practice, this means replacing conventional, fossil fuel based power plants with renewable energy sources (RES) [2]. However, this transition introduces challenges due to the intermittent nature of RES, reduced system inertia, and limited flexibility in the power grid.[3]

Simultaneously, the rapid adoption of electric vehicles (EVs) presents challenges and opportunities for grid management. Without smart coordination, uncontrolled EV charging may amplify grid stress by concentrating demand during peak periods, potentially leading to capacity violations and even blackouts[4]. This issue is pertinent in Denmark, where the number of registered EVs has surged past 400,000 as of April 2025 [5] and is expected to exceed one million by 2030 [6]. The growing scale of the EV fleet underscores the urgency of implementing intelligent charging strategies that ensure grid stability.

Conversely, EVs hold substantial potential to support grid stability through smart control strategies. Ultimately, vehicle-to-grid (V2G) systems could enable EVs to function as flexible energy storage units [7]. Unlocking this flexibility remains a key challenge and is the focus of ongoing research.

The DTU project *Autonomously Controlled Distributed Chargers (ACDC)* [8], conducted from April 2020 to September 2023, successfully developed and tested autonomous smart charging technology. The project introduced a charge controller for electric vehicles capable of reacting to frequency deviations, along with a virtual aggregator that coordinates these controllers via broadcast signals [9, 10].

Additional research related to the ACDC project includes quantifying the flexibility potential of EVs [11, 12, 13] and assessing the AC to DC conversion efficiency of onboard chargers (OBCs) across different EV models [14].

1.1 Project Scope

This thesis is grounded in the findings of the ACDC project. The most influential contributions to this work are further detailed in the background section. One limitation of the ACDC project was that the controllers could not be implemented directly into the chargers, which lead to some suboptimal workarounds with increased control delay. This project equipped EV chargers with an internal microcontroller, specifically a BeagleBone Black (BBB), which acts as the primary controller of the device. The controller interacts in a local network through Ethernet connection to a switch. This thesis involved developing and implementing several control tasks that allow the chargers to make autonomous decisions and coordinate with other chargers. The primary control tasks are:

- Reading local measurements via a DEIF unit.
- Sending setpoints to a Phoenix Contact charge controller (CC).
- Enabling bidirectional communication with Energydata.dk (EDDK) for coordination between chargers.
- Develop a main control script that integrates all components and supports autonomous decision-making.

This charger setup, featuring an integrated microcontroller, eliminates some of the system delays reported in a previous study [15]. With an average upregulation activation time for a single EV of less than 1 second (see subsection 5.2), it meets the activation requirement for participation in Denmark's fastest frequency reserve (FFR).

To support development and testing, a simulation model has been built in Python. This model runs the same control algorithm but bypasses physical components, allowing for rapid prototyping and analysis. It enables testing of control strategies, investigation of system delays, and evaluation of performance under a wide range of constraints.

The control algorithms developed in this project focus on the following four prioritized objectives:

- Ensuring compliance with the current capacity at the point of common coupling (PCC) for the EV cluster.
- Charging vehicles efficiently based on a priority system that considers dynamic user inputs: energy demand and time of departure.
- Aligning power consumption with the aggregated spot market bid to avoid imbalance fees.
- Compatibility with frequency regulation services based on the power bid submitted to the market.

The contribution of the study lies in the evaluation and improvements of EVs' ability to meet the requirements for frequency containment reserves. Testing compatibility with frequency restoration reserves (aFRR, mFRR) was excluded from the scope because the investigation in that case would change from response times to evaluating charging flexibility.

Physical testing was conducted on four chargers: two located at DTU Lyngby Campus and two at DTU Risø Campus. The control architecture is designed and built for scalability. The reference perspective for the control design is a large EV charger cluster with a capacity of roughly 1 MW. Adopting the perspective of a large, aggregated system introduces the opportunity of active participation in the day-ahead and frequency regulation markets, i.e., submitting spot- and frequency reserve bids to a Balancing Responsible Party (BRP).

Determining the optimal spot- and frequency reserve bid based on the aggregated charger capacity exceeds the scope of this project. However, gaining knowledge about market structures and requirements has been an essential part of the project to develop a control system that is compatible with the present markets.

2 Background

This section provides the foundational knowledge underpinning the work in this thesis. It begins with an overview of EV charging in subsection 2.1, covering hardware, communication protocols, charging dynamics, and user behavior relevant to smart charging systems. Next, subsection 2.2 presents the main categories of control architectures—centralized, decentralized, and distributed—highlighting their respective advantages, drawbacks, and relevance for EV integration. Finally, subsection 2.3 outlines the technical and regulatory requirements for participating in the Danish day-ahead and frequency regulation markets, with a focus on services applicable to aggregated EV flexibility.

2.1 EV Charging

This subsection focuses on the underlying EV charging equipment, communication protocols and other charging dynamics. The majority of EV charging is expected to occur through AC chargers, primarily in residential areas and workplaces where user behavior tends to be consistent [16]. AC charging offers inherent flexibility, as vehicles are typically parked for longer than the time required to fully charge [10]. This flexibility is crucial for implementing coordinated control strategies in smart charging systems. Therefore, this thesis concentrates exclusively on AC chargers.

2.1.1 AC Charging Equipment

Other than the EV itself and its OBC, the main equipment to facilitate charging is an AC connector and a CC. This project works with the European standard IEC 62196 Type 2 charging inlet and Pheonix Contact charge controller shown in Figure 1 and Figure 2. The project also incorporates a DEIF device, a microprocessor-based multimeter used to monitor the charging process by measuring key electrical parameters. Configuration details for both the DEIF and the CC are provided in section 3. The cable shown in Figure 1 contains five power conductors: three for the phases, one neutral, and one ground. Additionally, the two smaller conductors at the top are the Control Pilot (CP) and Proximity plug (PX) lines, which enable essential communication for smart charging functionality.

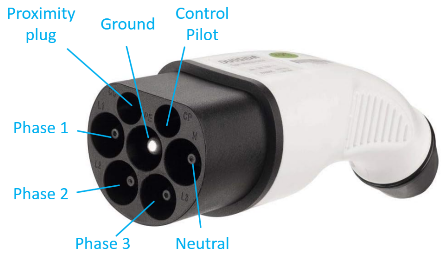

Figure 1: IEC 62196 Type 2 charging cable plug [17]

Figure 2: Phoenix Contact Charge Controller [18]

2.1.2 Communication Protocols

The CC manages communication with the EV via the CP line by generating Pulse Width Modulation (PWM) signals. The duty cycle of the PWM signal defines the maximum current the EV is permitted to draw. However, the EV may draw less than this setpoint in specific situations—for example, when the battery is near full state of charge (SOC) or operating outside its optimal temperature range. PWM is a square-wave voltage signal whose duty cycle represents the proportion of time the signal stays in the high-voltage state during each cycle. A detailed overview of the CP wiring is shown in Figure 3.

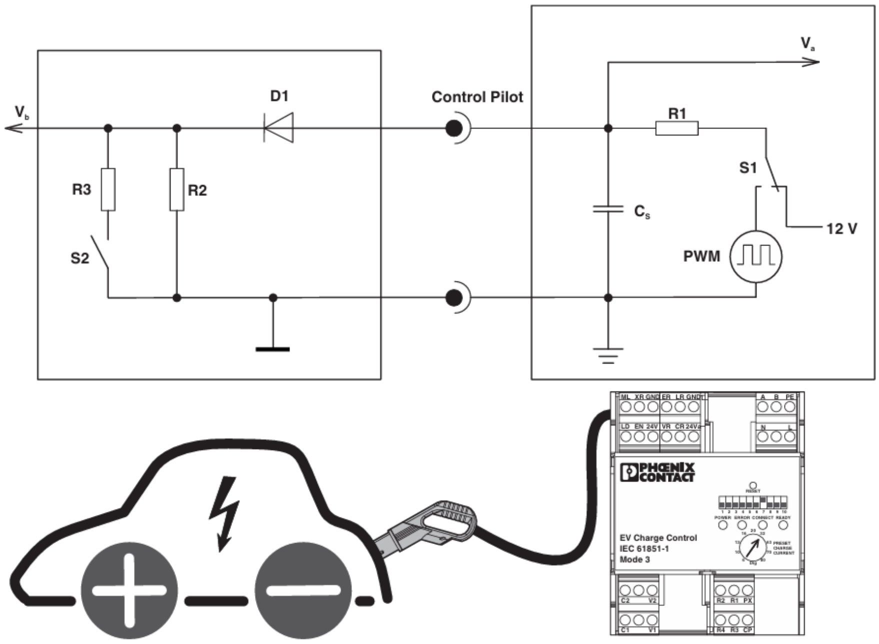

Figure 3: Control Pilot Wiring. EV side to the left and CC side to the right [18].

Resistors R2 and R3 in the EV are used to communicate EV status by altering the voltage seen by the CC. The nominal voltage values and their interpretation are presented in Table 2.

Table 2: Vehicle states (A-F) based on Control Pilot signaling [18].

Vehicle status	**Connected vehicle**	**Charging possible**	**Va**[1]	**Description**
A	No	No	12 V	Vb[2] = 0 V
B	Yes	No	9 V	R2 detected; B1 (9 V DC): EVSE[3] not ready; B2 (9 V PWM): EVSE ready
C	Yes	Yes	6 V	R3 = 1.3 kΩ ±3%; ventilation not required
D	Yes	Yes	3 V	R3 = 270 Ω ±3%; ventilation required
E	Yes	No	0 V	Vb = 0; EVSE short-circuit or no power supply
F	Yes	No	N/A	EVSE not available

[1] *Va = measured voltage in the EV charge control*

[2] *Vb = measured voltage in the vehicle*

[3] *EVSE = Electric Vehicle Supply Equipment (charging station)*

Table A.7 from the IEC 61851-1:2017 [19] standard states how the duty cycle of the PWM signal directly correlates with the maximum charging current I_{max} that the EV is permitted to draw. The relationship is defined by the following equations:

$$I_{max} = (\text{duty cycle percentage}) \cdot 0.6\text{ A, for signals between 10-85\%} \tag{1}$$

$$I_{max} = (\text{duty cycle percentage - 64}) \cdot 2.5\text{ A, for signals between 85-96\%} \tag{2}$$

E.g. 50% and 96% duty cycles correspond to 30 A and 80 A charging current limits, respectively. The setpoint range set by this standard is 6-80 A with only integers allowed. Hence, the control granularity is limited to 1 A steps [19].

In addition to the CP line, the CC can utilize the PX signal to ensure that only charging cables with adequate current-carrying capacity are used. The resistance value on the PX line indicates the rating of the connected cable and plug. When properly configured, the CC can use this information to reject under-rated cables by activating a locking mechanism, thereby preventing potential overloads [18].

The CP and PX signals represent the only communication channels between EVs and CCs under present standards. As a result, important parameters such as vehicle model, SOC, maximum possible charging current, and the number of available phases cannot be exchanged through this interface alone [20]. These parameters are critical for implementing coordinated control of EV charging. Alternative methods for deducing this information will be discussed in section 4.

2.1.3 EV Charging Dynamics

EV Charging Flexibility

EVs offer flexibility when parking time exceeds charging time, making it important to quantify the flexibility potential of EV charging clusters for smart grid integration. However, this potential remains uncertain. Striani et al. [12] address this by introducing an evaluation tool and five flexibility indexes. The paper's sensitivity analysis highlights how various factors influence this flexibility. Grid connection—which is often lower the than summed outlet capacity of the charging cluster to minimize grid connection costs [21]—shows to have significant influence on power-based flexibility services, but only a marginal effect on energy flexibility. Storage capacity of EVs (or energy requested) is identified as the most influential factor for energy and time flexibility, with larger requests allowing greater flexibility gains. The study further concludes that EV clusters at work places tend to offer more flexibility at the beginning of shifts and less throughout the day due to fewer new connections. The number of EVs alone does not significantly impact flexibility potential; instead, connection patterns and energy requests have a much greater influence [12].

The paper [22] used data from 179,000 charging sessions in Sweden to generate synthetic weekly charging profiles for residential and workplace charging. They find that many drivers do not fully deplete their battery each day, leaving an energy buffer that could be shifted or shared. The authors quantify up to 51.5 kWh per EV per day of potential energy balancing. Such unused storage capacity and flexible charging behavior translate into substantial opportunities for smart charging strategies.

Charging Behavior in Relation to SOC

SOC is crucial in determining whether an EV is ready to charge. The internal charging logic evaluates SOC, connection status, and system load to initiate charging [23]. Commercial EVs with lithium-ion batteries typically follow a constant-current/constant-voltage (CC–CV) charging profile: high current is applied until a voltage limit is reached, after which the voltage is held constant and current gradually decreases [24]. This causes charging power to ramp down as SOC approaches 80–90%, entering a saturation phase to avoid overcharging [25].

Kostopoulos et al. (2020) observed this behavior in a BMW i3 using real-world data: maximum power (11 kW) was sustained until 80% SOC, after which power declined sharply. Charging efficiency significantly dropped at an SOC above 80%, with energy losses nearly doubling compared to the 20–80% range. The study concluded that charging beyond 80% often yields diminishing returns and increased battery stress [25] when keeping other impacting factors constant such as battery temperature [26].

Charging Efficiency

The article by Sevdari et al. [14] presents a comprehensive experimental evaluation of AC to DC conversion efficiency in EV OBCs. Efficiency is assessed by comparing the grid-side AC power draw with battery side DC power, the latter measured via the vehicle's OBD-II port. The study reveals significant variation in OBC efficiency across vehicle models and production years. Notably, lowest efficiencies are observed at minimum charging currents (6 A) with a general trend that lower currents yield lower efficiencies. The losses are attributed to the relatively constant power losses in the power electronics [27]. A heatmap summarizing OBC efficiencies for various EV models is shown in Figure 4.

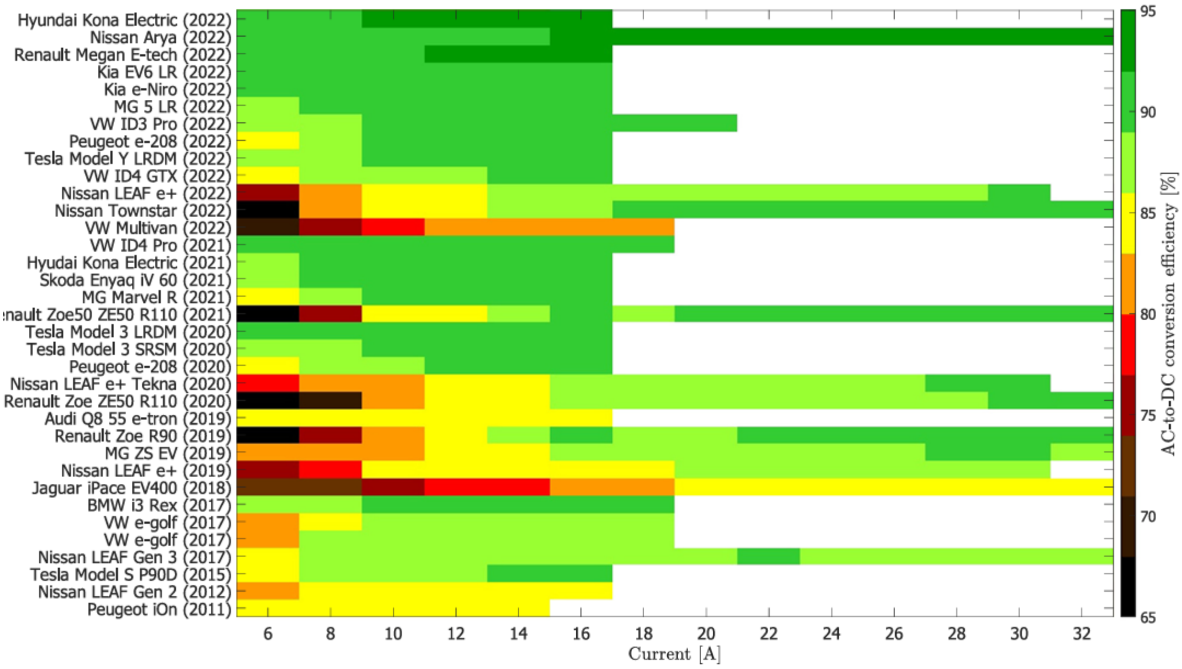

Figure 4: EV OBC characteristics on AC to DC conversion efficiency. Graphic is taken from [14]

The authors highlight that smart charging strategies can unintentionally increase total energy consumption in EVs by 1–10%, as they may induce operation in low efficiency regions. Another interesting finding by the study include that the battery's SOC was found to have negligible impact on efficiency, which indicates that the efficiency drops highlighted by [25] are due to low charging currents. Based on observed trends in newer models, the study anticipates continued improvements in OBC efficiency, with the market average potentially reaching 95% by 2030 and stabilizing between 90% and 95% in 2035 [14].

The variability in charging efficiency presents a challenge for smart charging optimization. In the final report of the ACDC project, they categorize 6 clusters of OBC behavior. Further, it is summarized that decision-making for efficient smart charging should be made based on the individual vehicle models [8]. Restricted by the time available, this thesis incorporates as much as possible of this important knowledge into the developed control algorithm, with the main focus of maximizing the current setpoint and thereby charging efficiency.

User Behavior

User behavior is a fundamental determinant of EV charging characteristics and the flexibility potential of EV clusters. Charging sessions are typically defined by user input data such as connection times, disconnection times, and the energy demanded [28]. The timing of this flexibility is significantly influenced by user patterns, e.g. working shifts [12].

Recent empirical analysis by Ziras et. al [29] of over 5,500 residential AC chargers in Denmark shows that charging behavior is evolving rapidly with increasing EV adoption. Users tend to charge more frequently and consume more energy per session, especially during winter months. Delayed charging is common, often overnight to align with low spot prices, but this has led to new early morning peaks. The increase in duration and volume of charging sessions provide increased room for load shifting, although rising peak loads highlight the need for coordinated smart charging strategies.

[30] presents survey-based experiments indicating that EV users in the Netherlands are generally willing to adjust their charging behavior, when offered modest monetary incentives or improved convenience. Charging preferences are sensitive to factors such as time of use, pricing, charging speed, and charger availability. Even small incentives can encourage off-peak or alternative location charging, suggesting strong potential for demand-side flexibility, when user motivations are appropriately addressed.

2.2 Control Architecture

The integration of EVs into the power grid requires intelligent coordination strategies to ensure grid stability and optimal operation [31]. Control architectures define how decisions regarding EV charging are made and communicated across the system. This section presents three main categories of control architecture: centralized, decentralized, and distributed. Figure 5 presents these control architectures alongside expansions done in literature. Each has unique characteristics, advantages, and limitations. These architectures differ in terms of data processing location, communication needs, scalability, and resilience. The appropriate choice depends on the use case, infrastructure maturity, and desired operational objectives [32, 33].

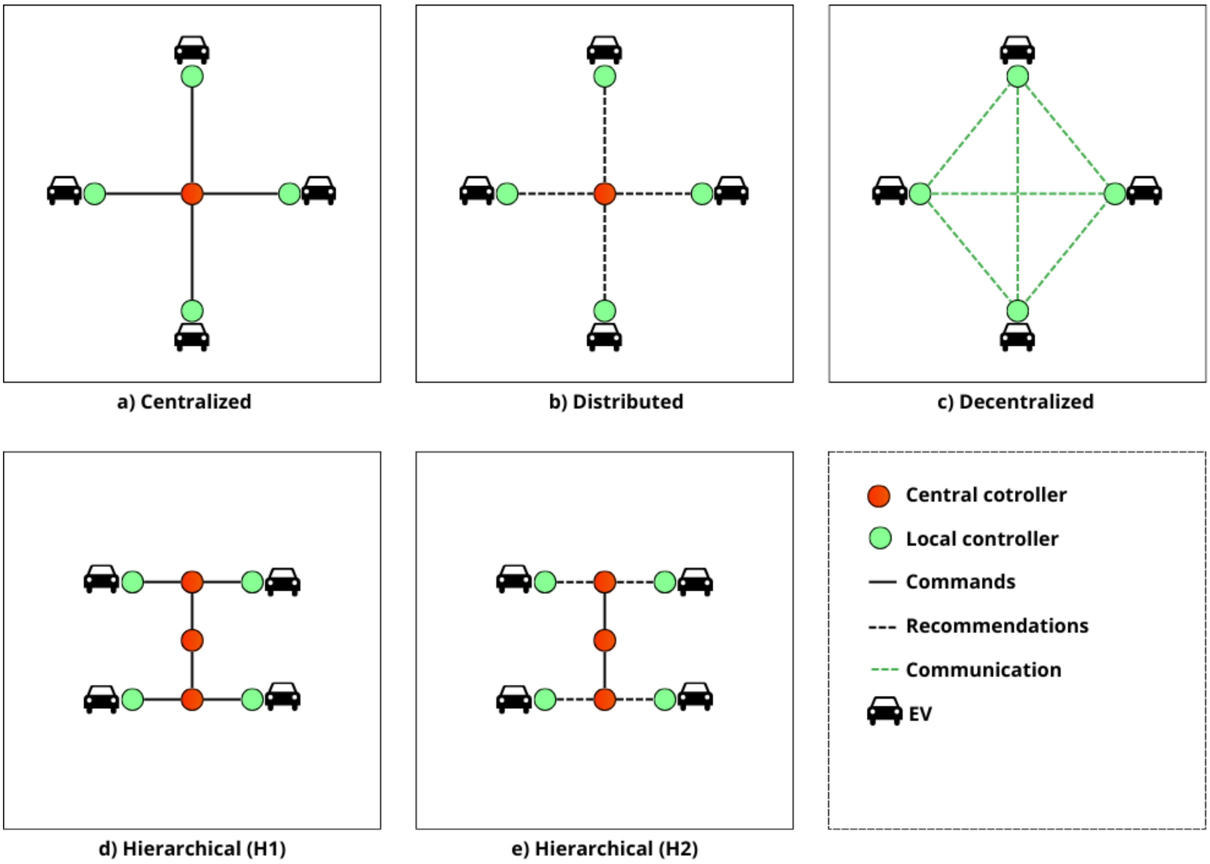

Figure 5: Various control architectures. This thesis focuses on the architectures visualized in a), b) and c) [34].

Centralized Control

In a centralized control architecture, a single central controller gathers data from all participating EV chargers and computes optimal charging setpoints, which it then distributes to the chargers. This central aggregator has a global view of the system and can optimize objectives such as cost, grid balance, or battery SOC synchronization [35].

Advantages: Centralized control offers operational transparency and can deliver globally optimal solutions using advanced optimization techniques. It facilitates coordinated responses and simplifies high-level planning and monitoring [9].

Disadvantages: This approach is highly dependent on reliable communication infrastructure and is vulnerable to single points of failure. It faces significant scalability issues due to increasing computational and data handling demands. There are also concerns related to data privacy, latency, and the costs of maintaining high performance cloud-based systems [11].

Decentralized Control

In decentralized control architectures, each charger acts as an independent decision-maker, solving local optimization problems without relying on central commands. These systems rely on either local only information or limited communication with other chargers. As such, decentralized schemes do not always lead to globally optimal charging outcomes due to incomplete system knowledge at each node. Nonetheless, they are gaining interest for their practical field applicability and computational scalability [34].

Advantages: Decentralized control schemes are highly scalable in terms of both computation and deployment. Since decision-making is distributed, the system is robust to individual node failures and resilient to network disruptions. These setups require less central infrastructure and offer enhanced data privacy [23]. [34] presents two types of decentralized control. Type 1 architectures enable direct coordination between EVs, while Type 2 models reduce communication overhead by using an indirect aggregator to broadcast coordination signals.

Disadvantages: A key limitation is the potential suboptimal system-wide performance due to limited information at the local level. In T1 architectures, extensive peer to peer communication can create high overhead as the number of chargers increases. In contrast, T2 architectures, while reducing communication needs, introduce partial central elements, somewhat blurring the line between decentralized and distributed control [33].

Distributed Control

Distributed control combines centralized coordination with decentralized autonomy. It typically involves a cloud-based central control and local control embedded in each charger. The central control manages global objectives and transmits commands and/or recommendations to the local control. The local control can apply local implementation based on central information or respond autonomously to local conditions e.g. alternating grid frequency [10, 11].

Advantages: This hybrid approach leverages the optimization capability of centralized systems while preserving the robustness and scalability of decentralized systems. Communication and computation loads are distributed, reducing system vulnerability and operational cost. If cloud communication fails, local controllers can continue operating based on local data, preserving a basic level of functionality [23, 10] .

Disadvantages: The control design is more complex, potentially increasing individual charger costs. Coordination between local- and central controllers must be carefully tuned to avoid instability. Experimental validations have so far been limited to small-scale implementations with a few chargers [11].

Maturity

Each control architecture has its place in the evolving smart charging ecosystem. Centralized schemes are suitable for tightly managed fleet operations or research environments, but their real-world scalability is limited. Decentralized systems offer resilience and privacy but struggle to meet complex grid objectives. Distributed control presents a promising compromise and is currently gaining traction in advanced pilot projects [11]. While centralized and decentralized control are well understood, distributed systems are at the forefront of research and development, with several field demonstrations. One example is the ACDC project at DTU [8] showcasing practical feasibility and effectiveness for future EV integration.

2.3 Participation in Electricity Markets

As previously mentioned, the project adopts the perspective of a large, aggregated EV charging cluster acting as an active participant in the electricity market.

This section outlines the regulatory obligations and opportunities for such an EV cluster to engage in both the Danish day-ahead market and selected frequency regulation markets. The focus is on compliance, technical requirements, and the minimum market participation threshold relevant to a large EV cluster. It should be noted that participating in multiple markets allow for multiple revenue streams [36].

Denmark is split into two electrical zones: DK1 (West Denmark, connected to Continental Europe) and DK2 (East Denmark, part of the Nordic synchronous area). DK2 is characterized by lower system inertia, resulting in a less stable frequency profile [37]. This motivates the need for the fast Frequency Containment Reserve Disturbances (FCR-D) and Fast Frequency Reserve (FFR), which are unique services for the Nordic synchronous area.

2.3.1 Day-Ahead Market

All electricity consumption and production must be accounted for by a BRP. A BRP is financially responsible for any differences between scheduled and actual consumption. The BRP submits bids to the electricity market and is settled based on the deviation from their cluster's planned schedule [38]. A lot of consumers do not worry about how much and when they consume electricity. But with complete flexibility comes higher electricity costs. The more intermediaries (suppliers, aggregators, etc.) between the consumed electricity and the BRP, the greater the profit margin is factored into the electricity bill. As the capacity of a consumer increases, it becomes feasible to take on the extra complexity of forecasting consumption to avoid supplier markups. On the one hand, a large electricity consumer is not obliged to bid on the day-ahead market. However, once the capacity accumulates to roughly 1 MW, it contributes significantly to the BRP's imbalance risk, and it becomes feasible to actively participate in the market to reduce these risks and associated costs [39].

If actual consumption deviates from the submitted bid, the BRP incurs a dynamic imbalance cost based on the maximum between the mFRR and aFRR price in the given 15-minute period. These are typically unfavorable compared to spot prices and incentivize accurate forecasting and alignment with the planned schedule [40].

2.3.2 Frequency Regulation Services

If a company wants to deliver frequency regulation in Denmark, it needs to either have an agreement with an approved BRP or become one. Alternatively, if the company only wants to deliver services with negligible energy delivery FFR, FCR and/or FCR-D it can be approved as a Balance Service Provider (BSP) [41].

Besides being approved as a BRP or BSP, participation in the markets requires prequalification of the involved assets. The prequalification process is different for each service. For all services, it includes performance testing of response times, a verification of power meter accuracy of +/- 5 %, and data logging at a resolution of 0.0167-10 Hz (1 measurement every 60 s-100 ms) depending on the service [42].

Table 3 summarizes the specifications and requirements for each primary reserve. According to the danish transmission system operator—Energinet—response time is defined as the time to deliver 5% of the required response [41], and it is specified for FCR, FCR-D and FCR-N. Ramping constraints indicate the minimum time to reach a certain power level (as a percentage of the required response). The table also lists under- and overshoot limits representing the maximum allowed deviation from steady state response. Under- and overshoot limits are listed in the table for upregulation services and are inverted for downward regulation. Notably, the asymmetric services FCR-D upregulation and FFR only require the EV cluster to reduce charging power. This is advantageous for EV charging, as it avoids the need to reserve headroom for increasing charging power (as would be required for downregulation). As a result, the available capacity at the PCC is effectively reduced for downregulation services. Brief descriptions of each service and their operational frequency ranges are presented after the Table 3.

Table 3: Technical overview of selected frequency regulation services [42, 41].

Service	**FFR**	**FCR-N**	**FCR-D**	**FCR**
Price zone	DK2	DK2	DK2	DK1
Response time	None	2.5 s	2.5 s	2 s
Ramping constraints	0.7-1.3 s (100 %)	60 s (63 %), 180 s (95 %)	7.5 s (86 %), *Linear	15 s (50 %) 30 s(100 %)
Undershoot limit	5 %	5 %	5 %	5 %
Overshoot limit	35 %	**20 %	**20 %	15 %
Minimum bid	0.3 MW	0.1 MW	0.1 MW	1 MW
Bid Type	Only upregulation	Symmetric	Asymetric	Symmetric
Load Factor (2021)	≈ 0 %	0.5 %	0.05 %	0.05 %
Requires BRP	No	Yes	No	No
Measurement rate	10 Hz	1 Hz	1 Hz	1 Hz
Frequency measurement resolution	5 mHz	5 mHz	5 mHz	5 mHz
Power measurement resolution	1 kW	1 kW	1 kW	1 kW

* *The power ramp is required to be at least linear for FCR-D. Which is further described in 2.3.2*

** *FCR-N and FCR-D only allow 10 % overshoot for linearity test [42].*

Fast Frequency Reserve

FFR is used to stabilize the grid frequency during major outages, particularly under low-inertia conditions. Its purpose is to limit large frequency deviations and prevent them from exceeding 1 Hz. If the deviation surpasses this 1 Hz threshold, Energinet initiates LFDD (Load Frequency Dependent Disconnection), which disconnects large loads to prevent a system collapse. FFR is only designed for upregulation, since the grid is more prone to large frequency dips rather than jumps. The activation of FFR providing units is binary and activates at frequency deviations greater than 300 mHz compared to the nominal 50 Hz. Table 4 lists three alternative activation levels for FFR. Depending on the activation level, one has to deliver a full reaction within 0.7-1.3 s.

Alternative	Activation Level [Hz]	Full activation time [s]
A	49.7	1.3
B	49.6	1.0
C	49.5	0.7

Table 4: FFR Alternatives.

As the only service, FFR has a restoration time of 15 minutes before the next delivery is required. FFR can have a provision period of either 5 or 30 s. If the 5-second provision is chosen, the deactivation requirement is a maximum of 20 % pr second and is fully deactivated within 30 seconds of activation start. If the 30-second provision alternative is chosen, there are no deactivation requirements [42].

Frequency Containment Reserve - Normal

FCR-N is used to stabilize the frequency close to the reference frequency. FCR-N units must be fully activated at +/- 100 mHz frequency deviations [43]. Between 0-100 mHz deviation, the units are activated proportionally to the frequency change (0-100%) with no dead-band allowed. Because of these characteristics, the FCR-N is the primary reserve with the highest load factor and the only one requiring a BRP [42, 38].

Frequency Containment Reserve - Disturbances

FCR-D is used to reduce frequency deviations larger than 100 mHz, as these units are activated below 49.9 Hz and above 50.1 Hz for up- and downregulation, respectively. As opposed to FCR-N, the service is asymmetric, meaning that there are two separate markets for up- and downregulation. A unit providing FCR-D must be fully activated at +/- 400 mHz from 49.9 Hz and 50.1 Hz. In the ranges 49.9-49.5 Hz and 50.1-50.5 Hz, the reaction should be proportional to the frequency deviation. FCR-D units should reach 86 % of required power within 7.5 s. Additionally, FCR-D units have an energy delivery requirement:

"The supplied energy must from the start of the ramp to 7.5 seconds after the start of the ramp, be equivalent to minimum 3.2 seconds multiplied with the theoretical steady state response."[42]

Essentially, the power has to ramp linearly or faster.

Frequency Containment Reserve

FCR in DK1, much like FCR-N in DK2, is a symmetric service designed to stabilize the frequency close to 50 Hz. However, FCR units are allowed a dead-band of +/- 20 mHz [41] and are fully activated at +/- 200 mHz [42].

For all services, aggregation is permitted, provided that units operate under the same BRP/BSP and respond in a coordinated manner. Non-compliance—such as failure to deliver or respond within specified times—results in forfeiture of availability payments and potential removal from the market until re-qualification is complete [42, 41].

3 System Development

This section introduces the charger and system-wide setup of the two EV parks in operation, as well as the two websites developed for the project: the *Control Board* and the *Visuals* webpage.

Initially, subsection 3.1 provides a detailed overview of the charger architecture, outlining the internal hardware components—BBB, DEIF multimeter, modem, and CC—and their communication structure. This subsection also introduces EDDK, which serves as the central data exchange platform.

Building on this, subsection 3.2 describes the experimental setup used for frequency response testing, including the use of an amplifier and oscilloscope to measure the system's dynamic behavior following frequency deviations. In addition subsubsection 3.2.1 details the oscilloscope data processing used to extract grid frequency, RMS current and current setpoint metrics from oscilloscope recordings. The processed data forms the basis of the reaction time presented in subsection 5.2.

The physical implementation is then presented in subsection 3.3, which describes the deployment of 4 chargers evenly distributed between Lyngby and Risø, including EV specifications and the grid connection at each site.

The subsection 3.4 introduces the two web-based interfaces developed to support control and monitoring. The *Control Board* facilitates user interaction and parameter configuration, while the *Visuals* page provides real-time insight into charger operation. This section also includes a description of the *Central Control* script developed to evaluate the trade-offs between local and centralized dispatch.

Finally, a simulation environment is introduced in subsection 3.5. It allows flexible testing of the control logic independently of the physical hardware. This model enables rapid development and debugging while preserving the key control dynamics.

Together, these subsections establish the technical foundation for the experimental work and control strategies examined in the subsequent chapters.

3.1 Charger Setup

Figure 6 and Figure 7 below illustrate our setup, where a reference charger is connected to the PCC of the EV cluster. Broadly, the charger consists of a BBB, responsible for local control and forwarding relevant data to the cloud via the modem; a DEIF, which measures e.g. current, power, and grid frequency; and a CC, which receives setpoints from the BBB and transmits them to the connected EV. Note that each charger only has a single plug. The figure also includes a Powerlab Internet of Things (P_IoT)—a DEIF multimeter measuring the PCC which is published to EDDK by a local BBB with internet connection. Other devices with access to these streams can subscribe to specific ones and receive real-time data.

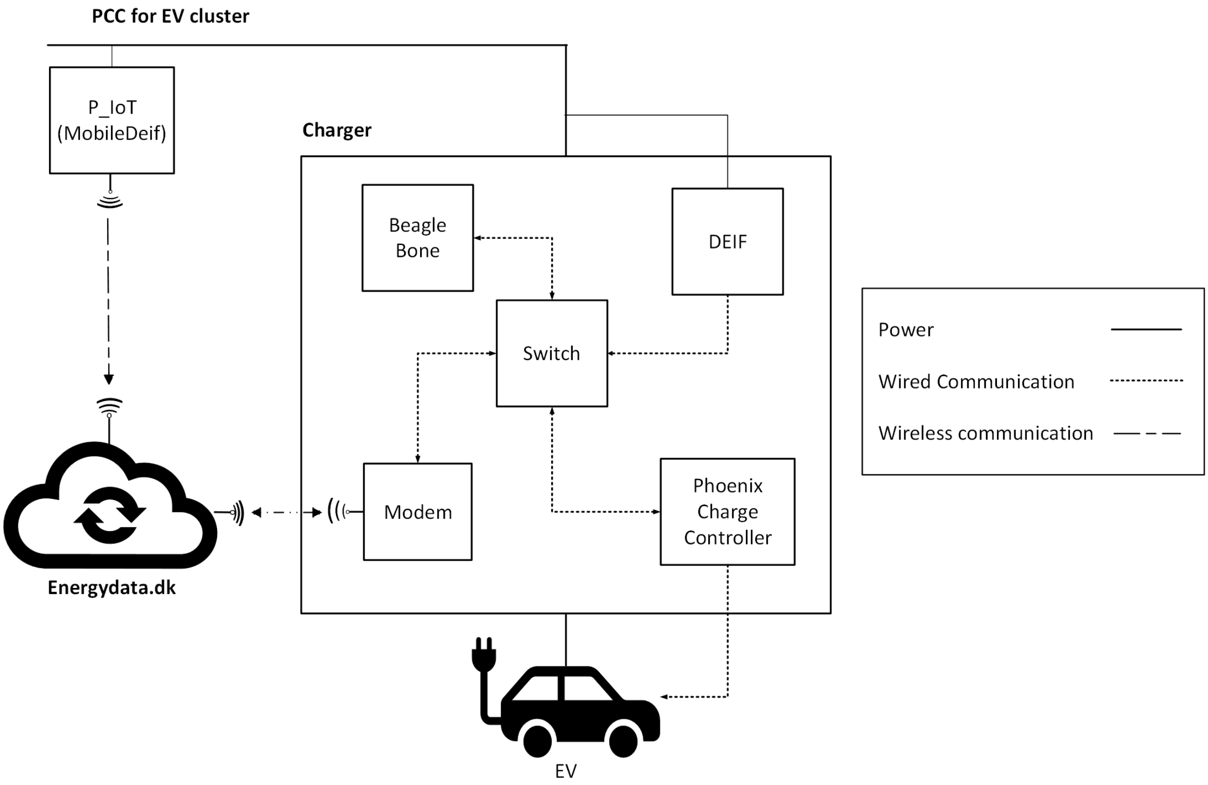

Figure 6: Physical setup of a single charger connected to the PCC.

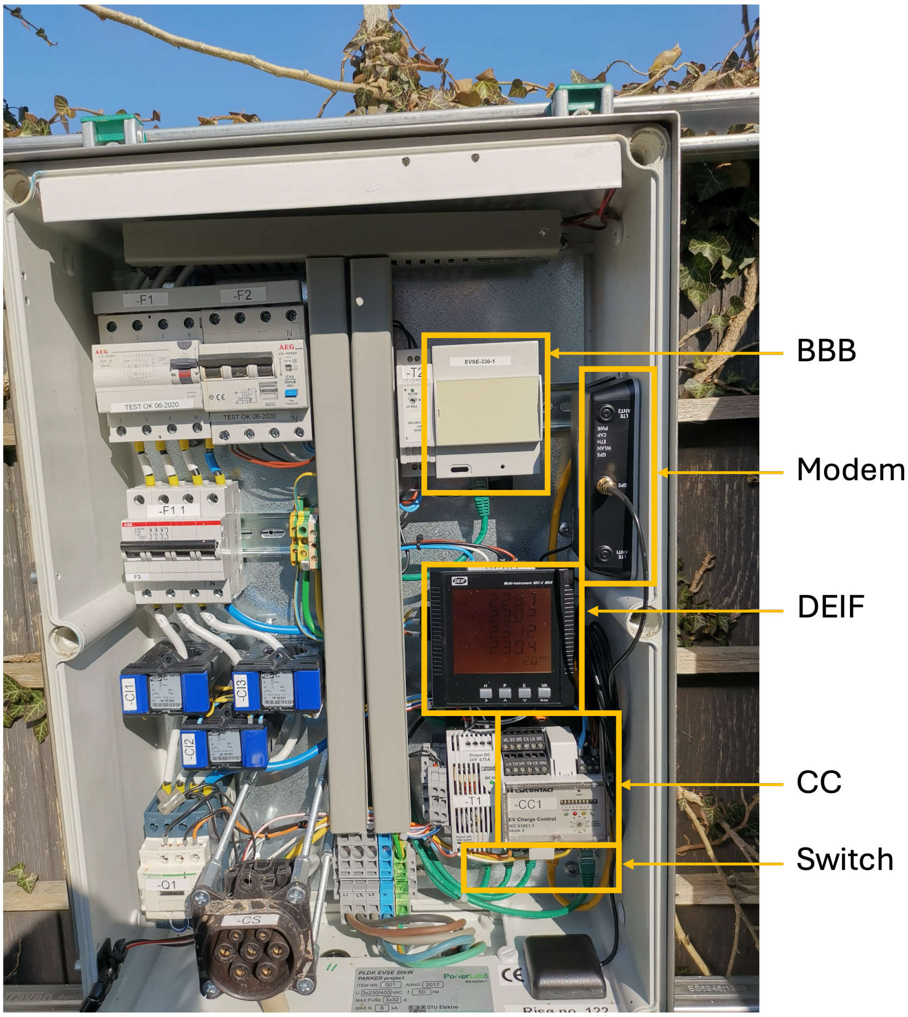

Figure 7: Photo of one of the chargers at Lyngby. Note that each charger only has a single plug.

The switch acts as the central communication hub, connecting the BBB, DEIF, modem and CC via Ethernet on the local network. By linking the charger's internal devices, the switch enables fast data exchange for control and monitoring. To allow the devices to communicate effectively, they are all configured to be on the same local IP network. The specific IP configurations appear in Table 5.

Table 5: Network configuration. Each modem has a public IP of the form YY.YY.YY.X, where X ranges from 80 to 83 for chargers 1 through 4, respectively.

Port in switch	**Component**	**IP address**
1	BBB	192.168.88.10
2	DEIF	192.168.88.100
3	CC	192.168.88.200
4	Modem (Gateway)	192.168.88.1
5	Modem (Public)	YY.YY.YY.X

3.1.1 BeagleBone Black (BBB)

The BBB is a low-cost single-board computer (SBC) running the open-source operating system Linux. It can run Python and connect to the internet via a modem, making it well suited for coordinated charge control [44]. Each BBB in the chargers is connected to a switch via Ethernet and powered by a dedicated 5 V power supply.

The BBBs were updated to run Python version 3.11.8, which is both recent, well documented, and compatible with the required packages. Since the BBBs are solely used for the work described in this report, no virtual Python environments were created.

The clocks on all BBBs are synchronized to the timezone Europe/Copenhagen using `ntp` at every reboot to ensure accurate documentation and visualization of the chargers' state.

3.1.2 Modem

The primary role of the modem is to assign a public IP address that enables SSH tunneling and/or VPN access to the local network behind the switch, independently of the BBB. Initially, internet access was provided by a USB modem with a SIM card coupled to the BBB, but when upgrading the BBBs to a newer OS version, the tunnel connection broke, the public IP changed, and the BBBs had to be accessed and reconfigured manually via a serial interface.

The updated setup using an autonomous modem enables stable VPN access and separates BBB specific configuration issues from those related to internet connectivity, making troubleshooting and maintenance more manageable.

3.1.3 DEIF

The DEIF MIC-2 MKII is a microprocessor-based multimeter used to measure key electrical parameters including current, voltage, active, reactive and apparent power across all three phases, as well as frequency.

The BBB communicates with the DEIF via Modbus TCP using the `ModbusTcpClient` from the `pymodbus` Python library. A `ModbusTcpClient` instance is initialized with the DEIF's IP address and the default Modbus port (502). Measurements are retrieved using the `read_holding_registers` method, which reads two 16-bit registers and converts them into a 32-bit float. The DEIF supports a refresh rate of 100 ms and offers Class 0.1 accuracy, corresponding to a maximum error of $\pm 0.1\%$. The measurement resolution is 10 mHz for frequency and 1 W for power. These specifications meet Energinet's prequalification requirements outlined in subsection 2.3.2, with the exception that the frequency resolution is 5 mHz below the required threshold [45].

To protect the DEIF from excessive current, a current divider with a 60:5 ratio is installed before the multimeter. This means that for every 60 A flowing into the charger, only 5 A reach the DEIF. As a result, current and power measurements must be multiplied by 12, while voltage and frequency remain unaffected. Notably, the DEIF registers power flowing into the EV as negative, which is reversed in the Python script.

3.1.4 Phoenix Contact Charge Controller

The key role of the CC in the developed system is using the CP to transmit current setpoints and specify the vehicle status (vehicle connected, vehicle ready for charging) translated from voltages as described in subsubsection 2.1.2. The CC has a communication interface from which control signals can be read and written. In this project, the control signals are transmitted by the BBB using the Modbus TCP protocol, exchanging data using the same protocol as described for the DEIF.

The CC has 10 dual in-line package (DIP) switches that can be either on or off to activate certain control functions. For this project the following DIP switch configuration is used:

- DIP 7 - ON: Charging process enabled
- DIP 10 - ON: Control via Ethernet (Modbus TCP) enabled

The project does not utilize the PX signal to evaluate the current carrying capacity of the cables, nor is the locking mechanism of the cable a part of the setup. Hence, the rest of the switches are kept off to avoid redundant checks and ensure compatibility between the specific hardware and control methods in use. This helps reduce unnecessary errors while keeping the charging process efficient and responsive.[18]

3.1.5 EnergyDataDK

EDDK serves as the database for online communication between chargers and the *Central Control* system. Within EDDK, datasets act as superclasses for individual data streams. We created the Automatic EV Charging (AUTEC) dataset, which contains multiple data streams, all using the same MQTT token — effectively serving as a password for read and write access. In real-time operation, EDDK functions primarily as a data hub, where information is published and retrieved via subscriptions at a resolution of 1 Hz [46]. The datastreams in use comprise:

- `CentralSetpoints`
- `Charger_1`
- `Charger_2`
- `Charger_3`
- `Charger_4`
- `Charger_5-30` - only active when performing the experiments listed in subsection 5.4.
- `UserInput`

Publishing and subscribing to data streams in Python requires an MQTT client, and we use the `paho-mqtt` package for this purpose. Each client instance is randomly assigned a unique client ID, as duplicate IDs cause disconnections. A secure connection is established using the `tls_set` method, and `on_connect` and `on_message` callback functions define how incoming messages are handled. Communication begins with `loop_start`, a non-blocking method that allows the BBB to continue computations while listening for new messages. Unlike `loop_forever`, this method does not automatically attempt to reconnect if the connection is lost. To ensure data is published reliably, our implementation includes checks on connection status before publishing and timeout mechanisms for reconnection.

All messages published to EDDK must be formatted as serialized JSON dumps. Messages can either contain a single value (e.g., an integer) or multiple values, the latter being widely used in our setup to transmit data as dictionaries with additional context.

Publishing Content of Chargers

Published data each second from each BBB includes:

- Current on each of the 3 phases [A] (even though some or all may be 0 A).
- Phase voltages on each of the 3 phases [V].
- EV power [W].
- EV apparent power (S_{sum}) [VA].
- System reactive power (Q_{sum}) [VAr].
- Frequency [Hz].
- Status of EV. See Table 2.
- Number of phases. See subsubsection 4.2.3.
- Setpoint signal sent to CC in integer amperes.
- Priority as calculated in subsubsection 4.2.4.
- Energy charged during current charging session.
- Max current of EV. See subsubsection 4.2.5.
- Test variable. May be Step, Sine or No test as received by the *Control Board*.
- Simulation boolean as received from the *Control Board*.

3.2 Experimental Setup for Frequency Response Testing

To test the response time of the cars to frequency changes, the experimental setup was expanded with an amplifier and an oscilloscope, as illustrated in Figure 8.

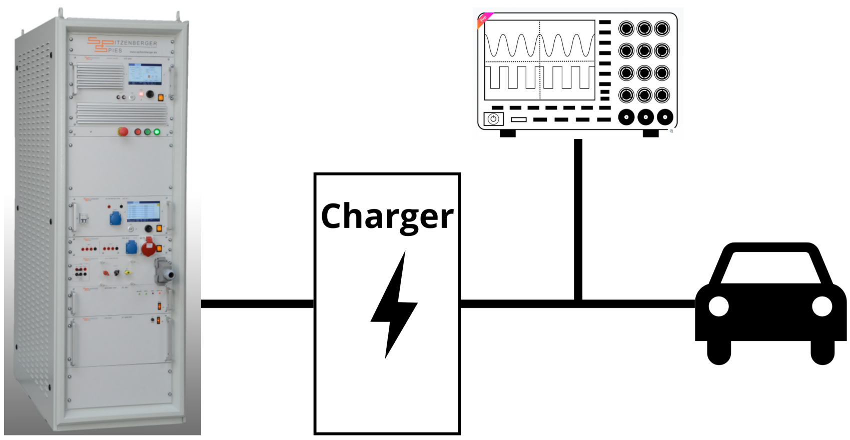

Figure 8: Illustration of experimental setup with amplifier and oscilloscope.

The amplifier was configured to 230 V per phase and used to shift the frequency from 50 Hz to 49.45 Hz to make sure an FFR event was triggered. This allowed testing of both the activation time of the EV's onboard charger and the speed of the control software running on the BBB reacting to the changes in DEIF measurements. The oscilloscope served as an external measurement device to verify the activation time.

Measurements on the oscilloscope were taken with:

- Scope 1 measuring AC voltage on phase 1 - to retrieve frequency.
- Scope 2 measuring AC current on phase 1 - to retrieve changes in current drawn by the EV.
- Scope 3 measuring voltage on the control pilot to capture the PWM signal - to retrieve current setpoint.

An example measurement in a 5 second scope is shown in Figure 9 with voltage (yellow), current (green) and PWM (blue).

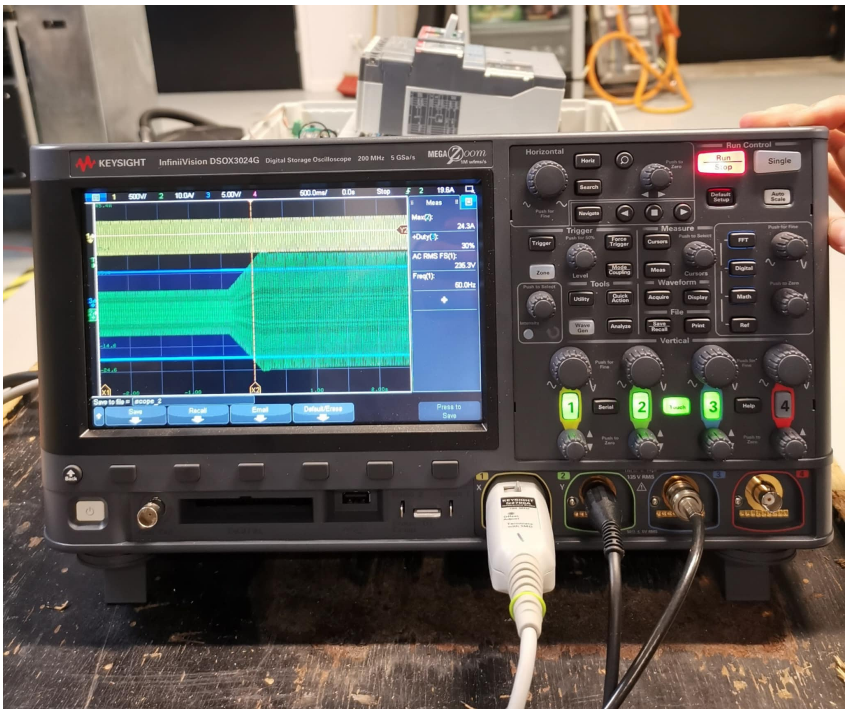

Figure 9: Example oscilloscope measurement. Scope 1 (Yellow): AC voltage, Scope 2 (Green): AC current, Scope 3 (Blue): Voltage (PWM signal).

These three signals enabled quantification of the system delay—time difference between frequency drop and power reaction—consisting of:

1. **Control delay** – the time between the frequency change and the PWM duty cycle adjustment by the BBB.

2. **Onboard Charger delay** – the time from PWM duty cycle change to the OBC reacting by adjusting the charging current.

The first delay is independent of the EV model but depends heavily on the measurement device and control software, including factors like loop redundancy, sleep durations, and CPU performance. The second delay varies significantly across EV models and between the car ramping charging power up or down.

3.2.1 Oscilloscope Data Processing

To generate the results presented in subsection 5.2, data was collected from an oscilloscope over a 5-second interval, resulting in 64,488 measurements sampled at 12.9 kHz. A sample of the first five entries is shown in Table 6.

Table 6: Sample head of data from oscilloscope. The AC voltage is used for calculating the frequency, AC current is converted to RMS current, and the PWM is translated to a current setpoint.

Second	**AC Voltage, phase 1 (Frequency) [V]**	**AC Current, phase 1 [A]**	**PWM in CP [V]**
-2.4900	275.25	-15.02	1.92
-2.4899	279.27	-15.37	-12.40
-2.4898	283.29	-15.53	-12.35
-2.4898	287.31	-15.57	-12.20
-2.4897	291.33	-15.68	-12.20

A maximum sample rate of 12.9 kHz on the oscilloscope was on the verge of being too low to perform the desired test. Hence, it was important to find a balance between accuracy of the frequency measurement and the temporal resolution. The accuracy of the frequency determined from the the phase voltage benefits from many samples while temporal resolution is defined by the size of our time window. The goal of the test is to prove an activation within 1.3 s, which led us to consider 100 ms temporal resolution as the minimum requirement. With the data segmented into windows of 100 ms, each contained N = 1,289 measurements. This resulted in 50 grouped samples per test. Although a higher sample rate could have improved accuracy, the oscilloscope did not support this without shortening the total measurement duration. Therefore, the above settings were selected as a compromise.

Determining the Frequency

Initially, frequency estimation was performed using a zero-crossing method. This involved applying a rolling mean (sub-window size 5), counting sign changes to estimate the number of zero crossings, dividing by two to determine the number of cycles, and dividing by the window duration to obtain the frequency. However, this approach is sensitive to noise and smoothing artifacts.

To improve robustness and accuracy, frequency was instead estimated using the Fast Fourier Transform (FFT) via `numpy`. Each 100 ms window was multiplied by a Hann window to reduce spectral leakage, after which the FFT was computed. From here, the average frequency in each window was calculated.

A sample of the alternating voltage of the frequency can be seen in Figure 10. Note that the x-axes are not aligned, i.e. voltage is only shown for 40 ms, whereas the translated frequency is shown for the full 5 seconds:

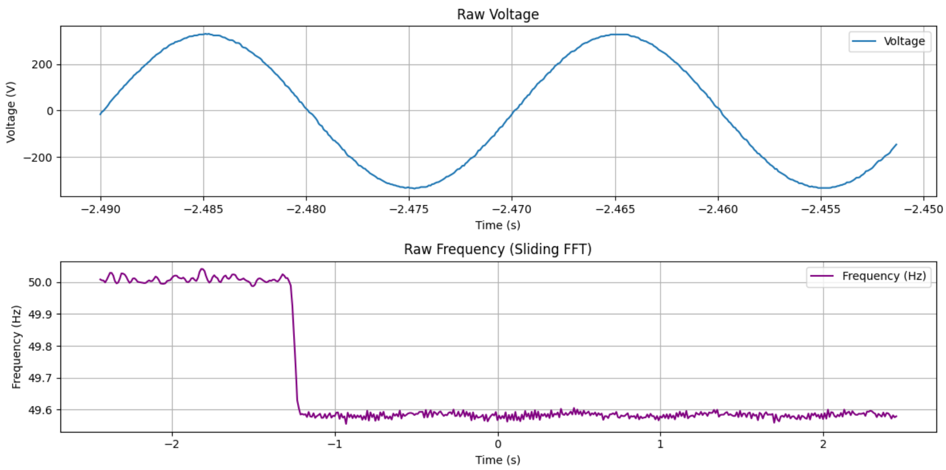

Figure 10: Raw voltage and the processed frequency signal.

Calculating the Setpoint

As detailed in subsubsection 2.1.2, the charging setpoint is transmitted via a PWM signal on the CP line. With the duty cycle of the CP having a fixed frequency of 1 kHz and a 12.9 kHz sampling rate on the oscilloscope, each cycle contained approximately 13 samples. Since the PWM signal alternated between high voltage (e.g., 6 V when charging is allowed) and low voltage (−12 V), the duty cycle was computed by summing all positive values within each 100 ms window and dividing by the total number of samples:

$$\text{Duty cycle percentage} = \frac{\sum_{n=1}^{N} V_{n,\text{PWM}}}{N} \cdot 100\% \quad \forall V_{n,\text{PWM}} > 0 \tag{3}$$

Where N is the number of measurements in each window. The resulting duty cycle was then converted to a current setpoint using Equation 1 and Equation 2.

Calculating the RMS Current

To compare the current setpoint with the measured current, the AC current signal was converted to RMS values for each 100 ms window using:

$$I_{\mathrm{RMS}} = \sqrt{\frac{1}{N}\sum_{n=1}^{N} I_{n,\mathrm{AC}}^2} \tag{4}$$

Alignment of Time Axis and Smoothing of Frequency

To make it easier to visually compare how the EV charger responds to a frequency drop over 10 test runs, the data shown in Figure 38 was preprocessed in several ways. First, all test runs were synchronized by aligning them to a common reference point: the first 100 ms time window in which the frequency dropped below 49.7 Hz was defined as $t = 0$. This allowed consistent alignment of the response curves across all experiments. To further simplify the visualization, frequency values above 49.7 Hz were uniformly set to 50 Hz, while those below were mapped to 49.65 Hz. Although this does not reflect the true shape of the frequency signal—as shown in Figure 35—this simplification is justified because the FFR logic is binary: It activates fully as soon as the frequency falls below 49.7 Hz, independently of the following frequency activity. The goal was not to analyze frequency shapes, but to clearly illustrate how quickly and consistently the chargers react.

3.3 EV Parks

The control algorithm has been developed on 2 car parks being located at the DTU departments in Lyngby and Risø. All chargers are set up as visualized in Figure 6. Both parks are compatible with different kinds of cars, but when performing the tests that are in this report, the cars were as listed in Table 7 and depicted in Figure 11:

Figure 11: 2 charging EVs at Risø. The setup in Lyngby is identical except for the EVs used (see Table 7).

Table 7: Specifications for the EVs used for the results section 5.

Car number	Car model	Location	Max current	Number of phases
1	Nissan Leaf (test model)	Lyngby	28A	1
2	Volkswagen ID4	Lyngby	16A	3
3	Nissan Leaf	Risø	16A	1
4	Volkswagen ID3	Risø	16A	3

Both EV parks were coupled to the grid with a PCC capacity of 32 A. Both parks had their own P_IoT measuring and publishing phase current, voltage and frequency to EDDK once every second. A complete visualization of the system wide setup is shown in Figure 12.

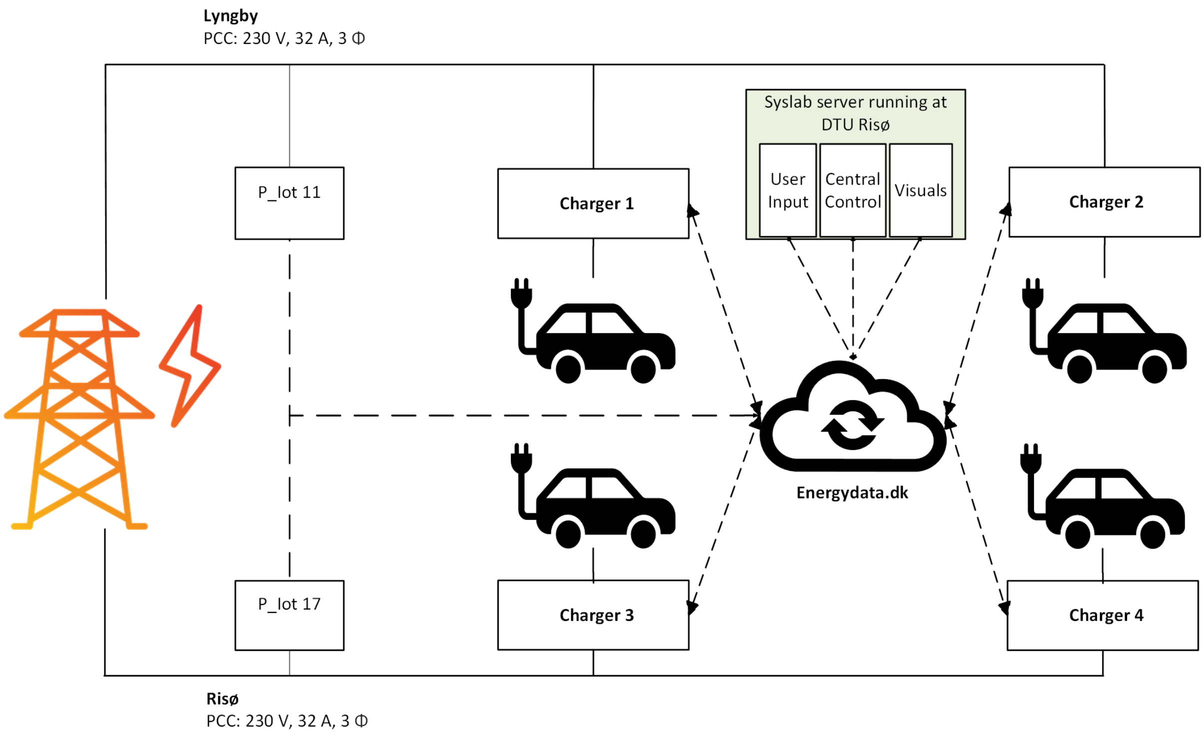

Figure 12: System wide setup of 2 EV parks coupled to the grid. As shown by the dashed lines, all communication goes through EDDK. Solid lines represent power.

3.4 Websites

3.4.1 Control Board

To simulate a user interface for initiating EV charging, the *Control Board* website was developed and is hosted on a server in Risø with the IP-address `http://130.226.55.35:5000`. To accommodate users with limited time and technical expertise, a simple interface was designed to collect information about the user's charging urgency. The interface is shown in Figure 13.

User-submitted charging requests are essential inputs to the implemented control algorithms. The effectiveness of the control relies on the honesty of the users and their understanding of energy units, particularly the size of a kWh. This aspect is further discussed in section 6.

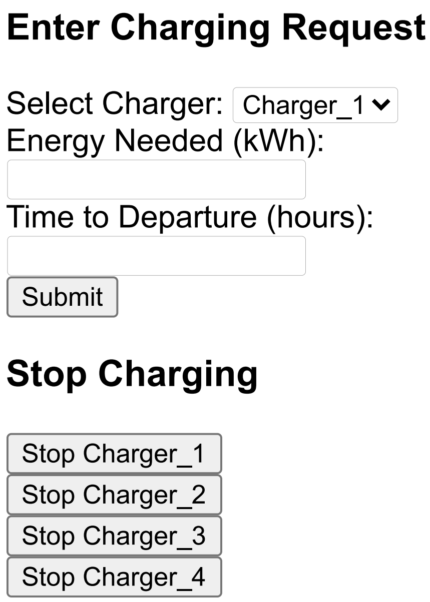

Figure 13: User interface as envisioned on a physical touch display at each charger, as well as on a mobile app. It should be noted that the user should only be able to stop the charging of the chargers they initiated themselves.

In addition to user inputs, the *Control Board* offers advanced functionalities for the car park owner. The website provides an overview of the status of the four chargers and displays the five most recent charging requests. It also enables configuration of several key control parameters through EDDK publications:

- Whether initial current dispatches should be calculated on a central server or locally on each BBB.
- Whether the chargers should respond to deviations in grid frequency.
- Whether the BBBs should simulate a charging session or charge a physical car.
- The ability to initiate or interrupt a sine or step test.
- The type of frequency regulation service to be provided.
- The spot market bid (cumulative power reference) to simulate a new 15-minute period.
- The frequency regulation power bid.

The full interface of the website is shown in Figure 14 and Figure 15.

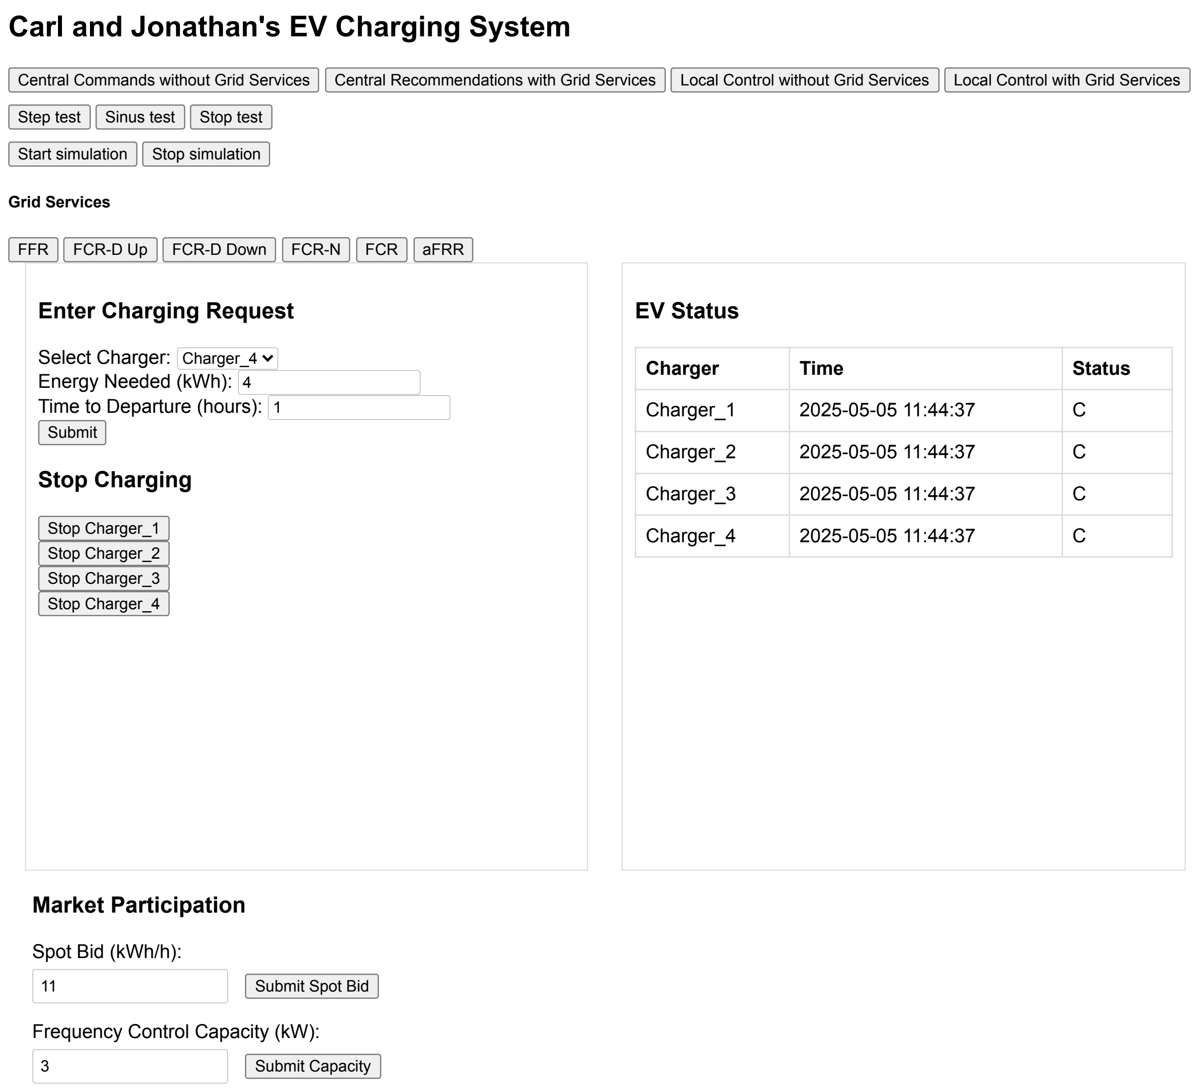

Figure 14: Control board as seen by the car park owner without the overview of the latest charging requests. It should be noted that even though the aFRR button is implemented, the control algorithm is left for future work.

Recent Charging Requests

Charger	Time	Energy (kWh)	Time to departure (h)	Priority (kW)
Charger_4	2025-05-05 11:44:18	4	1	4
Charger_3	2025-05-05 11:44:15	3	1	3
Charger_2	2025-05-05 11:44:11	2	1	2
Charger_1	2025-05-05 11:44:08	1	1	1
Charger_1	2025-05-05 10:00:28	1	1	1

Figure 15: The remaining part of the control board website containing the overview of the latest charging requests.

3.4.2 Visuals

To support both result analysis and debugging, a dedicated web interface named *Visuals* was developed. It is accessible at `http://130.226.55.35:5001` and provides real-time insights into the state of each charger. The interface displays four key plots: measured power and power setpoint alongside the frequency; measured current and current setpoint; charging priority; and cumulative energy charged. The current setpoint can be interpreted as a power setpoint by assuming a nominal voltage of 230 V and considering the number of active phases.

Under normal conditions, frequency measurements are provided by a P_IoT device located in Lyngby. During frequency control experiments for the whole charger cluster, however, the website uses simulated frequency data loaded from a local CSV file. The same file is also used by the BBB units in each charger, replacing live DEIF readings. As a result, these experiments cannot be used to assess reaction rates to frequency deviations. Instead, they serve to demonstrate the functionality and stability of the control algorithm with four active EVs.

Because the EDDK system updates at a frequency of 1 Hz, the Flask-powered website is configured to refresh at the same rate. An example of the interface is shown in Figure 16.

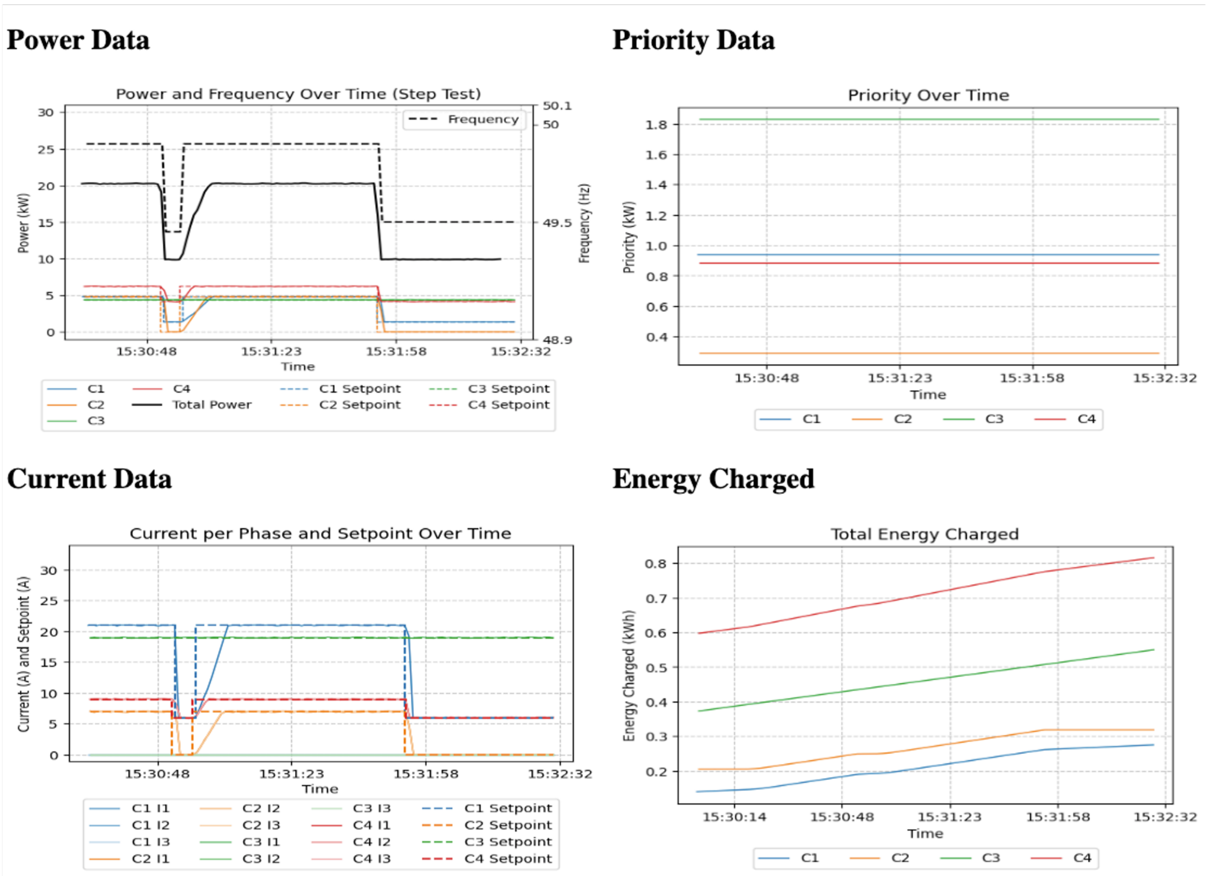

Figure 16: *Visuals* website during a simulated frequency drop.

To support result collection for the analyses presented in subsection 5.1 and subsection 5.3, an additional script was developed. This script replicates the website's functionality by logging the same parameters to a CSV file and generating the four corresponding plots. As with the live interface, the data collection is limited by the 1 Hz update rate of the EDDK system.

3.4.3 Central Control

A common approach to coordinating EV charging within a car park is through centralized control, where a server broadcasts setpoints to all chargers. To investigate this method, a script was deployed on the central server at Risø, performing the same current dispatch calculations as those executed locally on the BBBs.

The aim of introducing *Central Control* is to evaluate the impact of offloading computational tasks from the BBBs to a central unit. This is particularly relevant in the context of frequency control, where reducing the computational load per control loop enables faster iteration cycles and quicker response times to frequency deviations.

The performance differences between central and local control will be analyzed in section 5.

3.5 Simulation Model

To enable flexible testing of the control code, a simulation model was developed. This model mirrors the control algorithms for physical control implemented on the BBBs but operates with the `SIMULATION_BOOL` flag set to `True`, instructing the Python scripts to bypass physical components such as the DEIF and the CC.

To initiate a simulation, a separate script must be executed for each charger to be simulated. The main differences between the simulation and the real-world implementation are as follows:

- Instead of generating a PWM signal through the CC, the setpoint value is stored internally as a variable.
- The frequency is retrieved from the P_IoT device located with the simulated charger unless testing in which case the frequency is read from the local CSV file.
- A JSON file containing realistic technical specifications for each simulated EV is loaded at runtime. These specifications include ramp rates (in both directions), available battery capacity (in kWh), number of phases, and maximum current.

The simulation assumes that each EV is continuously plugged in and that all system components function as intended, since the primary objective is to test the control functionality rather than hardware behavior.

While not deemed relevant for the scope of this project, the model could be extended to accept setpoints below the 6 A minimum controllable current, which could be used to demonstrate the elimination of overshoots as a consequence of the uncontrollable power range. Likewise, the simulation model accepts decimal current setpoints, which removes the imprecision caused by the integer setpoints as restricted by the CC. However, the effects hereof are not investigated in detail, as the imprecision of 1 A is negligible in a large setup.

4 Control Methodology

This section outlines the control methodology developed to coordinate multiple EV chargers across two parks, viewed from the perspective of a large aggregated EV charger cluster. The implementation is driven by the four main control objectives defined in subsection 1.1, which serve as the foundation for the algorithmic design described in this chapter.

We begin by presenting the implemented control architecture in subsection 4.1, describing four operating modes that vary in their level of decentralization and ability to respond to grid frequency deviations. These modes are selected dynamically via the *Control Board* interface and are used to evaluate system performance under different configurations.

The underlying control protocols, described next in subsection 4.2, defines how user inputs are translated into charging priorities and how those priorities guide the dispatch of current setpoints. The methodology for initial current allocation is presented alongside the code structure. Important supporting mechanisms such as phase detection, priority updates, maximum current estimation, and dummy charging are also detailed.

Finally, subsection 4.3 presents the frequency control logic covering unique algorithms for both upregulation and downregulation. This section also details how the system deals with hardware-induced limitations such as uncontrollable power ranges and EV ramp rates.

4.1 Control Architecture

To evaluate the system's performance under different configurations, four control strategies were implemented based on the architecture described in subsection 2.2. These strategies can be toggled via the *Control Board* website, which publishes a control flag to EDDK every 5 seconds. Each BBB listens to this data stream and updates its control behavior according to the most recently received flag.

Flags 1 and 2 represent strategies where the initial setpoints are calculated centrally at the Risø server and then sent to EDDK for the BBBs to receive the information. In contrast, flags 3 and 4 require each BBB to calculate the current dispatch locally, as outlined in subsubsection 4.2.1.

The distinction between the flags lies in whether or not the BBBs respond to grid frequency as well as where the priority-based current dispatches are calculated:

- **Flag 1**: Fully centralized control. The BBBs receive fixed setpoints from the server and do not respond to frequency deviations.

- **Flag 2**: Distributed control. Centralized calculation of current dispatch based on priority. The BBBs may modify the received setpoints based on real-time frequency measurements, as explained in subsection 4.3. This control strategy resembles the Type 2 architecture described in [34].
- **Flag 3**: Decentralized control without frequency response. Each BBB calculates its own setpoint based on shared priority data, but does not consider grid frequency. The control strategy is similar to Type 1 presented in [34].
- **Flag 4**: Decentralized control with frequency-based adjustments. The BBBs both calculate their own dispatch based on priorities and adjust it based on frequency. This control strategy is similar to Type 1 presented in [34].

In this setup, flag 1 represents a purely centralized strategy with no local intelligence or communication between chargers, while flags 2 to 4 reflect increasing levels of distributed control and autonomy. This is summarized in Table 8.

Control Flag	Responsibility for calculation of priority-based setpoints	Frequency responsive
1	Central server	No
2	Central server	Yes
3	Local BBB	No
4	Local BBB	Yes

Table 8: Specifications of the 4 implemented control flags and their corresponding strategies.

As a fallback, if a BBB does not receive an updated flag or a centrally calculated setpoint for 15 seconds, it will default to flag 3.

4.2 Control Protocols

The core of the proposed control methodology is based on assigning priorities to EVs derived from user inputs upon connection to the charger. These priorities form the foundation of the decision-making algorithms and determine which EVs are charged first and which should adjust their power uptake during frequency control.

The *Control Board* interface allows users to enter two key parameters: the desired energy (E_{desired}) and the expected departure time (t_{dep}). These inputs are used to calculate the initial charging priority ρ in units of kilowatts, corresponding to the minimum average power required over the charging period:

$$\rho_{\text{init}} = \frac{E_{\text{desired}}}{t_{\text{dep}}} \tag{5}$$

If the user does not complete the input within two minutes of plugging in the EV, default values are applied: 1 kWh of desired energy, a departure time 1 hour after connection, and a resulting priority of 1 kW. The user may still submit updated inputs later, which will overwrite the default and update the priority accordingly.

The values ρ_{init}, E_{desired}, and t_{dep}, along with the charger ID, are published to the `UserInput` EDDK datastream, to which all local BBBs subscribe. The relevant BBB decodes the received JSON message and stores the input locally for use in control decisions.

4.2.1 Priority-Based Current Dispatch

The objective of current dispatch without frequency control is to fully utilize the power of the spot bid while respecting the PCC capacity constraint. This is achieved by creating a dictionary of current setpoints and updating the remaining spot bid power and PCC current capacity accordingly.

In each iteration of the control loop, the BBBs publish relevant data, including priorities, to EDDK on their individual channels (`Charger_i`), which all other BBBs and *Central Control* subscribe to. If priority based current dispatch is computed locally at each charger, all BBBs fetch all priority dictionaries at the beginning of the next iteration, sort them by priority magnitude, and use it for dispatching. If the dispatch is calculated centrally, each BBB will fetch the latest dispatch data from the EDDK datastream `CentralSetpoints`. These setpoints have been calculated on a central server using the same algorithm as the BBBs perform locally.

Next, the current readings from the P_IoT located at the PCC of each car park are compared to the aggregated EV currents published by all chargers at each respective site. The maximum error across all three phases determines the reduction applied to the nominal PCC capacity of 32 A at each separate park. These deviations are primarily attributed to the idle current of the chargers, which is approximately 150 mA per charger. Another source of error arises when a charger draws current without publishing data—a situation we define as dummy charging, which is further explained in subsubsection 4.2.6. This method of comparing the P_IoT measurements with the sum of local charger currents, in combination with the control strategy outlined in subsubsection 4.2.5, eliminates the need for a feedback loop as implemented in [10].

Based on the information of adjusted PCC and priorities, a current dispatch is performed by iterating through the sorted priority dictionary and allocating the highest allowable current to each charger with sufficiently high priority, as this results in the most efficient charging [14]. The dispatch respects and is restricted by the remaining PCC current capacity ($I_{PCC,\mathrm{remain}}$), the remaining spot bid power, and the minimum controllable current of 6 A, as visualized in Figure 17. The default maximum current is initially assumed to be 32 A, and updates according to the individual car capacities following the protocol described in subsubsection 4.2.5.

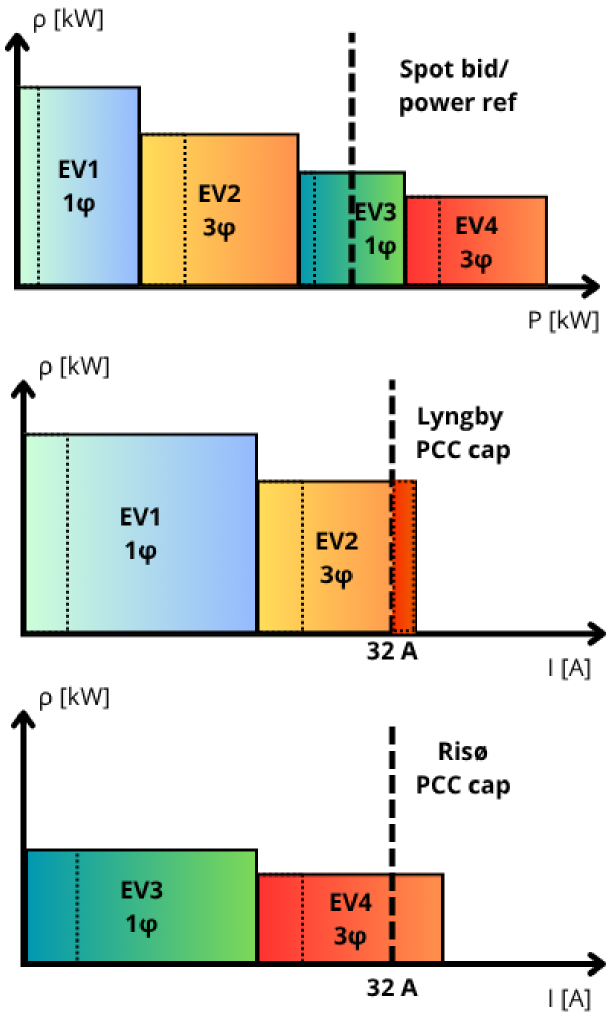

Figure 17: Visualization of the prioritized scheme and the constraints during initial current dispatch. In this specific case, EV2 is limited by the PCC capacity at Lyngby, resulting in EV3 consuming more power to satisfy the spot bid. Note that the y-axis shows priority in kW and the x-axis shows power in the first figure and current in the two latter. The dotted lines on each EV block represent the uncontrollable power range (6 A on each phase), whereas the dashed lines display the spot bid and PCC capacity in the power and current plots respectively.

Before updating the dispatch dictionary and remaining capacities, the algorithm checks whether the preliminary setpoint of the present charger would cause the remaining spot bid power to be below the controllable power of the next EV in line. If so, it evaluates whether the present charger can reduce its setpoint enough to allow the next EV to begin charging at minimum controllable power without itself entering the uncontrollable power range. This behavior is illustrated in Figure 18.

If reduction of a charger's power to allow the next charger with marginally lower priority to charge improves the clusters ability to adhere the spot bid, the current setpoint is reduced by the minimum required amount; otherwise, the present charger keeps its original setpoint, and the remaining power will not be fully utilized.

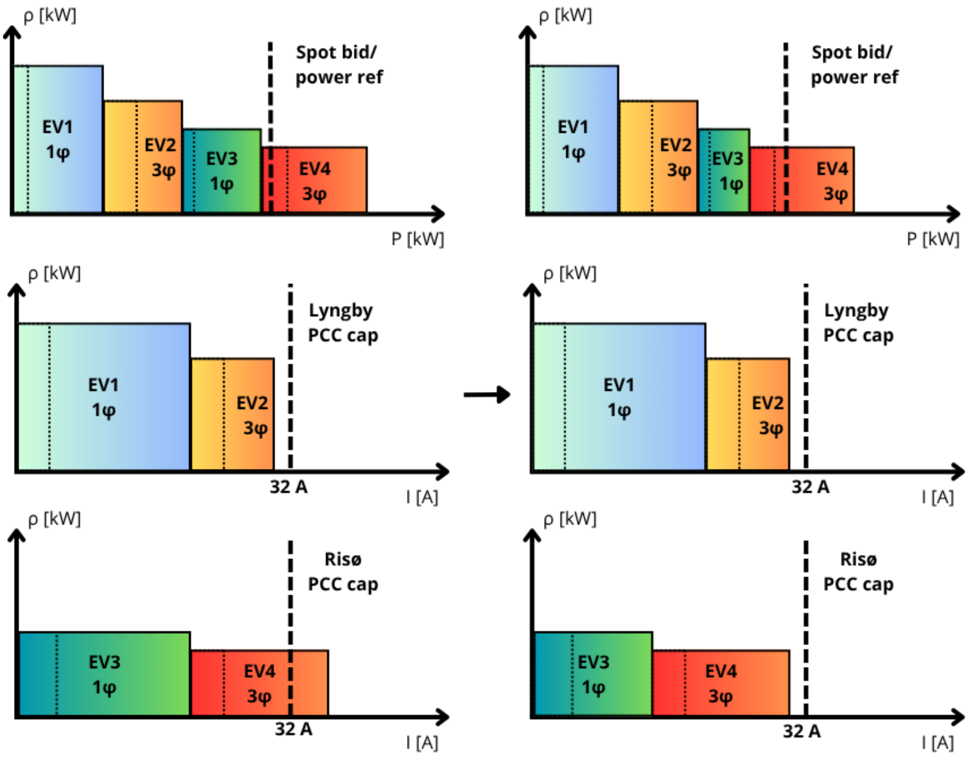

Figure 18: EV3 reducing setpoint, allowing EV4 to charge. The dotted lines on each EV block represent the uncontrollable power range (6 A on each phase), whereas the dashed lines display the spot bid and PCC capacity in the power and current plots respectively.

After these checks, the entry for the dispatch dictionary is finalized, and the updated values for PCC capacity and spot bid power are recorded. The overall decision-making process is illustrated in the flow chart in Figure 19.

Following this, the control logic sets the allowed current via the CP pin. Initially, we included a 5-second delay when ramping up from 0 A to a non-zero setpoint, but this was removed because the cars ramp down much faster than they ramp up.

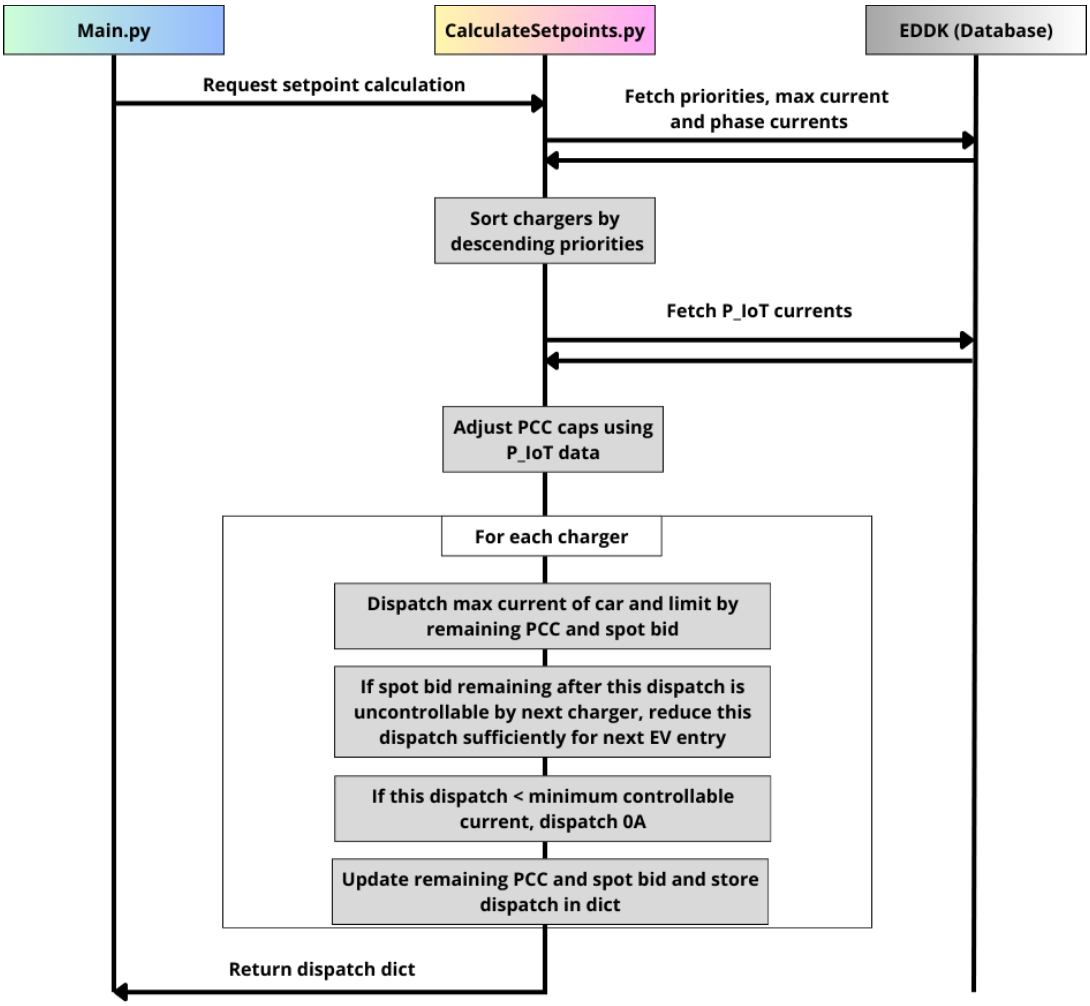

Figure 19: Flowchart illustrating the current dispatch process performed by the *Central Control* or locally by the BBBs.

4.2.2 Code Structure

All the scripts developed in this thesis and the overall code structure are illustrated in Figure 20, where arrows indicate the direction of data flow. Each charger in the cluster is equipped with a BBB running a set of local Python scripts, all coordinated by a single main script operating in a continuous loop. In addition to the local execution, the SYSLAB server hosts the two websites described in subsection 3.4 and runs the *Central Control* scripts responsible for current dispatch when control flag 1 and 2 are active as described in subsection 2.2.

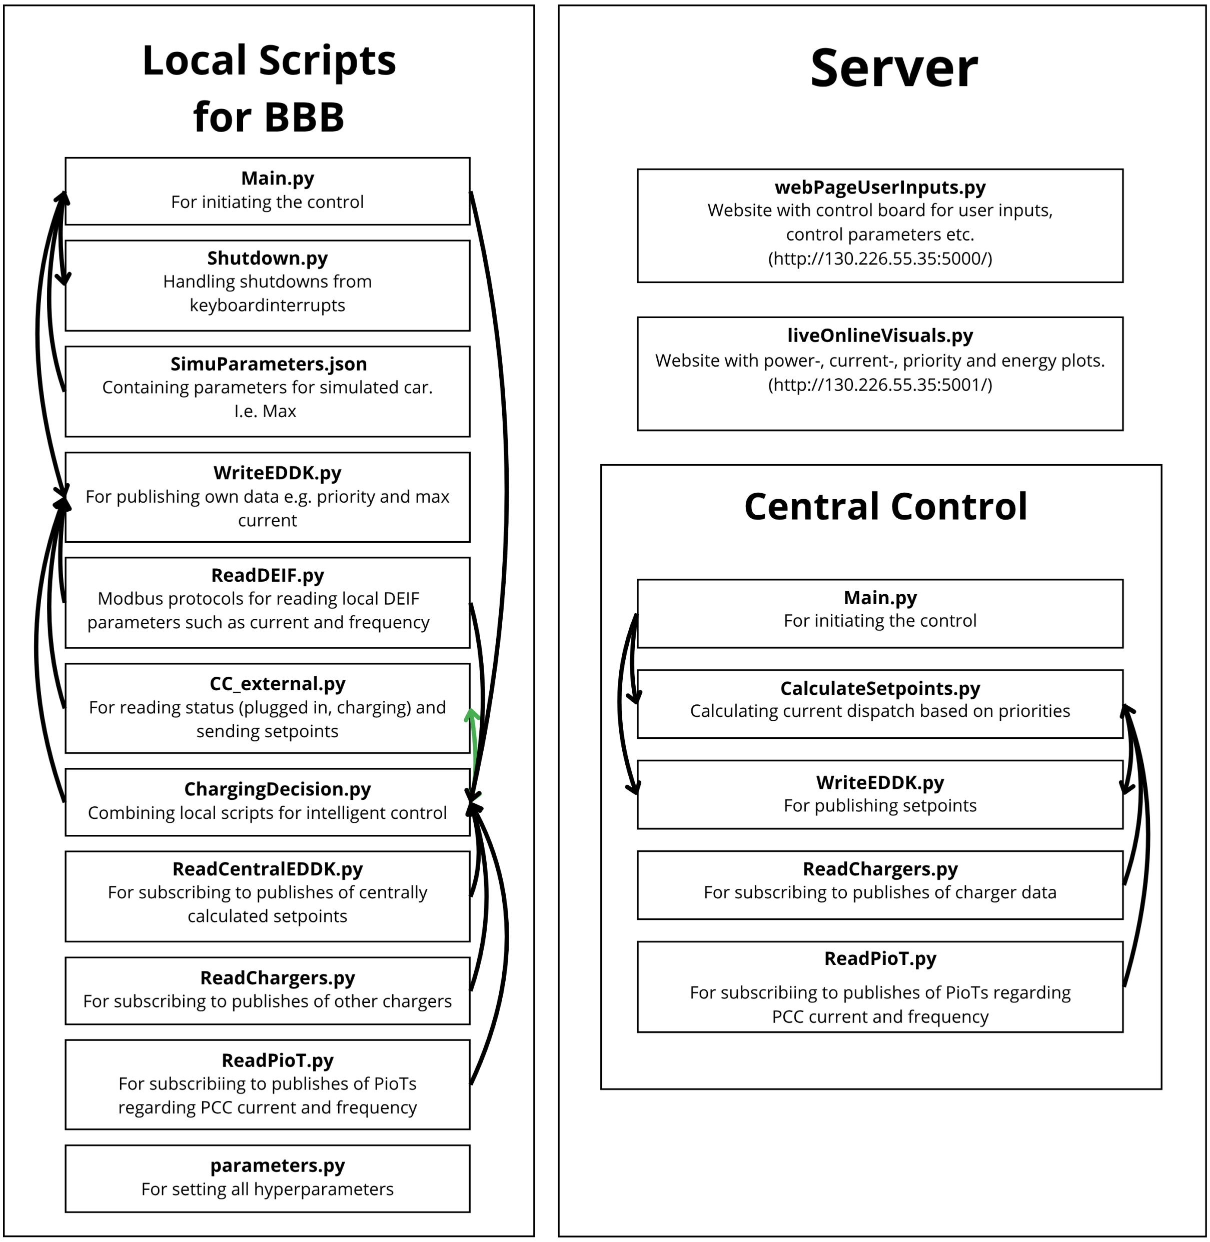

Figure 20: Overview of the code structure, including all scripts. Arrows indicate data flow: black arrows are unidirectional, whereas the green arrow is bidirectional, as the BBB both reads the EV status and writes the current setpoint through the CC.

4.2.3 Phases

This project defines priority as the average charging power per phase throughout the charging period to achieve the requested energy. Hence, the number of phases utilized by the individual EV is an important factor. This is because 3 phase EVs can charge at three times the power of 1 phase EVs while drawing the same current.

As a default, EVs are assumed to be 1 phased until otherwise measured. However, if the DEIF detects three active phases just once, the priority is divided by three and the current dispatch is calculated accordingly with regards to the higher share of the spot bid. If three phases are measured, the algorithm will not check number of active phases again, but simply remember that the relevant car is 3 phased. Upon unplugging or digitally stopping the charging session, the number of phases is reset to 1. After confirming the number of phases, the code ensures that the priority is below the maximum power per phase.

4.2.4 Updating the Priority

In addition to being divided by 3 upon detection of three active phases, the priority is recalculated at regular intervals. In this project, we chose an update interval of 5 minutes, which strikes a balance between responsiveness and system stability. Initially, a shorter interval of every loop was implemented. However, this lead to oscillations between EVs with nearly the same priority, alternating between on and off states, as their priorities continually overtook each other. When 2 EVs by turn oscillate between on and off, the utilization of available power decreases due to the EVs' ramp rates. This behavior was also discussed in [10], where the authors demonstrated the functionality with a 2-minute update interval, but recommended longer intervals for practical applications to avoid such inefficiencies.

Even though the priority is only updated every 5 minutes, the energy consumed is stored for each iteration of the main loop. Upon updating the priority, the consumed energy the last 5 minutes is then subtracted from the energy needed, and the time remaining before departure is used to calculate the new priority:

$$E_{i+1,\text{desired}} = E_{i,\text{desired}} - \sum_{j}^{J} P_j \cdot t_{j,\text{loop}} \tag{6}$$

Where J is number of loop iterations completed since last priority update. Next, the remaining time is updated:

$$t_{i+1,dep} = t_{i,\text{dep}} - t_{\text{sinceLastUpdate}} \tag{7}$$

Which allows the new priority to be calculated:

$$\rho_{i+1} = \frac{E_{i+1,\text{desired}}}{t_{i+1,\text{dep}}} \tag{8}$$

It should be noted that as the hardware does not allow for extraction of information regarding charging efficiency, an efficiency of 100% is simply assumed. This assumption is discussed in subsubsection 6.2.3.

4.2.5 Max Current

When dispatching currents for each EV, it is essential to estimate the maximum allowable current for each individual vehicle to avoid unutilized space on the PCC. This is achieved by monitoring the deviation between the implemented current setpoint and the measured current at the local DEIF. Provided the deviation exceeds 1 A, a value is added to a deviation list. The code evaluates whether to update the maximum current of the car based on the criterion that the current deviation is steady and that the length of the deviation list is at least 20. A steady current deviation is defined as a difference of less than 0.5 A between the current reading and the one 20 samples ago. This typically occurs when a car consistently draws less than the initially assumed maximum of 32 A, allowing the system to adjust to a more accurate, lower maximum current which is simply set to the latest measured current.

When a new, different setpoint is issued, the deviation list is reset to avoid mistakenly setting the maximum current too low.

4.2.6 Dummy Charging

Dummy charging ensures that an EV can continue charging even in the event of a lost internet connection, assuming the local BBB is still operational. As a design choice, the dummy setpoint is 6 A, which is also the minimum controllable current. This value is also used as a fallback when determining the available PCC capacity at a location where the P_IoT has stopped publishing data. If the latest P_IoT message is older than 15 seconds, the system assumes the corresponding PCC capacity is reduced by 6 A for each EV that is not actively reporting.

4.3 Frequency Control

The frequency control strategy is divided into two branches depending on whether an increase or decrease in consumption is required. The algorithm begins by determining the sign of the required power change, which depends on the grid frequency, the magnitude of the frequency control power bid, and the type of frequency regulation service provided, as outlined in subsubsection 2.3.2.

Under normal operation and in the activation time tests shown in subsection 5.1 and 5.2, the frequency is read by the DEIF (see subsubsection 3.4.2). However, when performing tests on the EV fleet's ability to coordinate during frequency control events as in subsection 5.3, the BBBs read frequency values from a locally stored CSV file. The *Control Board* broadcasts the test start time, enabling each BBB to identify the corresponding frequency in the simulation for each loop iteration. The control logic then proceeds based on whether the frequency deviation indicates the need for up- or downregulation.

As not only the time between reading a new frequency and setting a new setpoint is of importance, but also the cars' reaction to the new setpoint play a role in the overall reaction time, a minor test was conducted to evaluate the delay of different car models. The models of choice were the Tesla Model Y and the Nissan Leaf, as the Tesla is a relatively slowly up-ramping car, whereas the Nissan Leaf is much faster with regards to current. It should be noted that the Tesla is 3 phased and therefore ramps power with thrice the speed of the current as well as the fact that the Nissan Leaf of interest is a test model which may impact its charging performance. The result of the tests can be seen in Table 9.

Table 9: Ramp rates of Tesla Model Y and Nissan Leaf.

Test type	**Tesla Model Y**	**Nissan Leaf**
0-16A	31 seconds	4.7 seconds
10-16A	6.8 seconds	1.3 seconds

As seen in the table, the long ramp time of the Tesla is particularly slow when starting from 0 A. This is due to the tardy activation of all 3 phases, as phase 1 ramps to 5 A within a few seconds, whereas the other two phases stay at currents below 1 A, which may be a safety factory default encoded into the BMS.

4.3.1 Upregulation

For upregulation, the algorithm of each BBB sums the power dispatch of all lower-priority EVs to determine whether the EV in question must reduce its setpoint. Seven states and their corresponding control responses have been identified, as presented in Table 10 and visualized in Figure 21:

Table 10: States and corresponding control actions during upregulation.

No.	**State Description**	**Control Action**
1	Lower-priority chargers can collectively reduce their power sufficiently.	Keep current setpoint.
2	The charger with marginally lower priority is in state 4, leaving a remaining power change need: $P_{\text{needed,remaining}} = P_{\text{marginal,reduced,needed}} - P_{\text{marginal,reduced,implemented}}$	Reduce setpoint by $P_{\text{needed,remaining}}$. If this results in a setpoint below the controllable range, turn off charging.
3	Remaining reduction is within this charger's controllable range.	Reduce setpoint by $P_{\text{needed,remaining}}$.
4	Remaining reduction falls in the *un*controllable range **and** a higher-priority car is still charging.	Reduce to minimum controllable power.
5	Remaining reduction falls in the *un*controllable range and **no** higher-priority car is charging.	Turn off charging.
6	Required reduction exceeds current setpoint.	Turn off charging.
7	EV not charging	Remain idle.

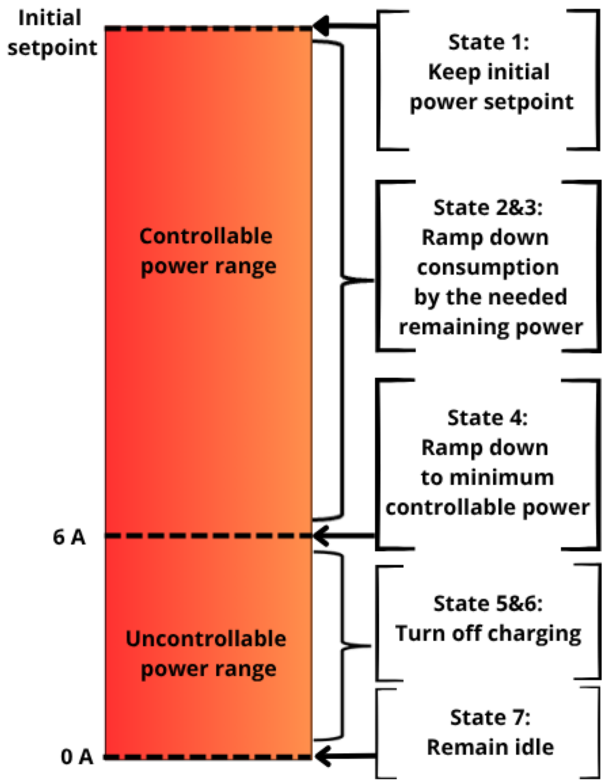

Figure 21: Mapping states to corresponding setpoint reductions. The max of this bar is set to initial setpoint, as the setpoint can only be decreased in this case.

This approach ensures that EVs with the lowest priority are the first to reduce their charging rates, in line with the prioritization strategy defined in subsection 1.1 and as visualized in Figure 22.

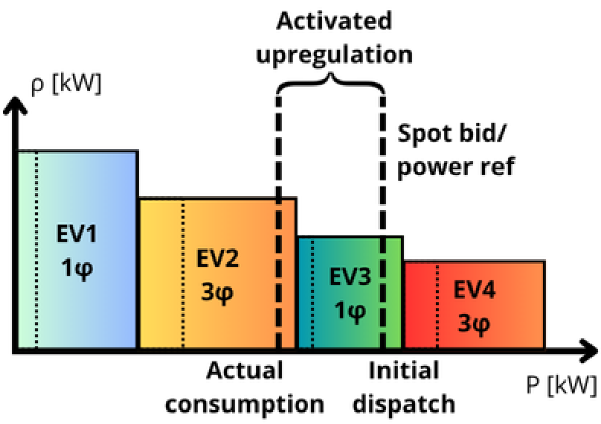

Figure 22: Example of the control of upregulation based on priority. Here, EV1 is in state 1, keeping its setpoint; EV2 is in state 2, ramping down consumption in the controllable power range; EV3 is in state 6, turning off completely; and EV4 is in state 7, remaining idle. The dotted lines on each EV block represents the minimum controllable current, whereas the dashed lines represent initial dispatch and actual consumption.

A potential disadvantage is if the charging car with the highest priority needs to ramp down its power consumption to the uncontrollable range, it will shut off completely, causing an overshoot in the frequency response. The theoretical maximum overshoot of the charger cluster is:

$$3\,\text{phases} \cdot 6\,\text{A/phase} \cdot 230\,\text{V} = 4.1\,\text{kW}.$$

However, with the specifications of over- and undershoot listed in Table 3, it is preferable to overshoot rather than the opposite, as the margin for overshoot is always larger than the one for undershooting. This issue is discussed further in section 6.

The advantage of this control strategy is its responsiveness and accuracy under normal circumstances that are not influenced by the above mentioned disadvantage: the BBBs and EVs used in this study were fast at reading the lower frequency, sending new setpoints and turning down the charging rate; and usually have a precision depending on the number of phases of the marginal car within:

$$0.5\text{A} \cdot \phi \cdot 230\text{V} = 345 \text{ W for } \phi\text{=3 and } 115 \text{ W for } \phi\text{=1}$$

The multiplication with 0.5 stems from the code rounding the setpoint of the car to nearest integer, which is a due to the setpoint granularity limitation mentioned in subsection 2.1.

4.3.2 Downregulation

When designing the control for downregulation, other methods had to be applied, as the current response to a higher setpoint was much slower than when decreasing the charging current. Bearing in mind the requirements in Table 3 for particularly FCR-D and comparing this to the ramp rates in Table 9, we constructed two different reaction algorithms depending on the magnitude of the error between the demanded and actual consumption.

When reading a frequency, the required power change is calculated as with upregulation. If the deviation between required and actual consumption is larger than 20% of the frequency control bid, all EVs ramp up as to meet the speed requirement of FCR-D. This is done by looping through all cars at each location and raising their not yet implemented setpoint by 1 A if their setpoint is above 6 A and by 6 A if their setpoint is at 0 A. This increment is repeated until either the PCC is saturated or no cars at the current location can increase their setpoint any longer. This strategy has the advantage of reacting quicker with a higher power change, but inherits an unfortunate risk of overshooting—a problem that increases with the loop time and update rate of EDDK. This will be visualized in subsection 5.3 and discussed in section 6.

If the deviation is closer than the above-mentioned threshold, the cars will respond in a manner much alike the one described for the upregulation. Firstly, each EV will sum the extra power that could be consumed by EVs with higher priority and from this information, we have identified 4 different states which can be seen in Table 11 and is further visualized in Figure 23.

Table 11: States and corresponding control actions during downregulation

No.	**State Description**	**Control Action**
1	Car is already at max current	Keep setpoint.
2	Needed power change is more than the potential power change of this car.	Set current setpoint to max current.
3	Needed power change is less than the potential power increase of this car **and** within the controllable power range.	Increase power setpoint by remaining power needed.
4	Needed power change is less than the potential power increase of this car **and** is within the *un*controllable power range.	Set current setpoint to 6 A.
5	Cars with higher priority can increase their charging by the required power change.	Keep setpoint.

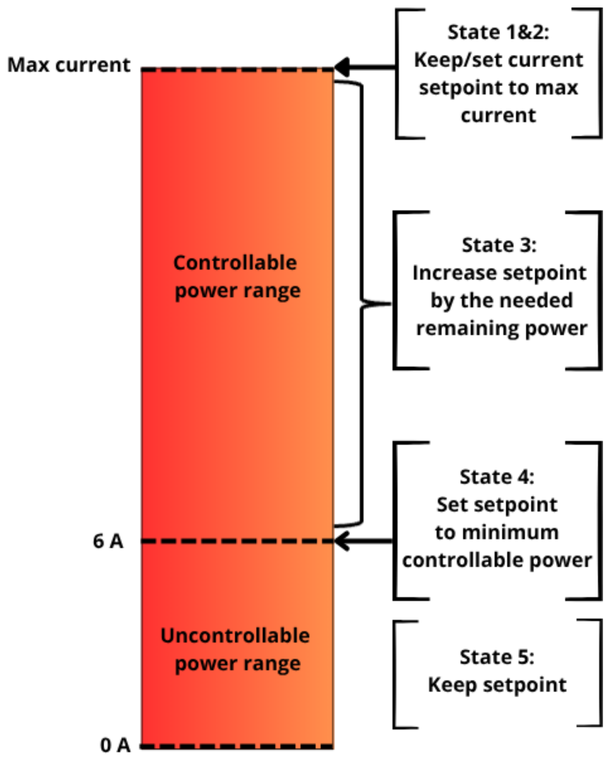

Figure 23: Mapping states to corresponding setpoint increases. State 5 shows no intervals, as the initial setpoint could be anywhere from 0 A to max current.

As seen in the bar plot, knowing the max current of each car is crucial, when performing downregulation. The absence of this knowledge (i.e. assuming a too high max current) could result in an undershoot of the service. However, if the EVs are well diagnosed before initiating the downregulation service, the precision of this control will often match that of the upregulation. With downregulation, the potential overshoot management is not split in two separate states as with the upregulation. Instead, if an EV is suppose to ramp into the uncontrollable power range, it will set a setpoint at the minimum controllable power. This means that the magnitude of the maximum overshoot of 4.1 kW is the same as with upregulation, but the frequency with which an overshot is observed will be higher. This could potentially be averted by more advanced control, but was not prioritized during the project, as the maximum theoretical overshoot would be negligible compared to Energinet's overshoot limits listed in subsubsection 2.3.2 in a large EV cluster of $\approx$ 1 MW.

5 System Performance and Experimental Results

This chapter presents the key experimental results of the developed control strategy for distributed EV charging. The experiments are designed to evaluate system performance under both standard charging conditions and while providing frequency regulation services. The aim is to assess how effectively the control framework balances its dual objectives: supporting the power system and meeting user-defined charging demands. The chapter is structured into four subsections, each highlighting a distinct aspect of system performance. Together, these results provide a comprehensive evaluation of the implemented approach.

Coordinated operation of car parks discussed in subsection 5.1 evaluates system behavior under standard charging conditions without frequency control. This subsection establishes a baseline for the priority-based control logic by assessing the quality of the developed charging functionality as well as the ability to respect global constraints.

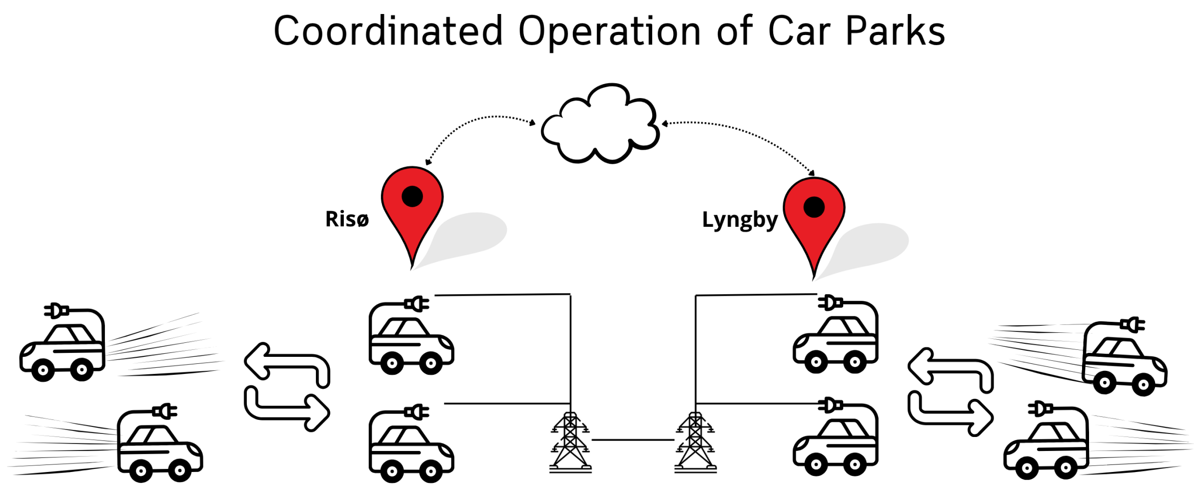

Figure 24: Visualization of tests regarding coordinated operation of car parks presented in subsection 5.1.

Activation time of a single EV, presented in subsection 5.2, quantifies the dynamic behavior of individual EVs in reaction to frequency deviations. High-resolution measurements are used to isolate the delays introduced by the control loop and the onboard charger, providing insight into the responsiveness of the system under rapid regulation scenarios focusing on compliance with FFR.

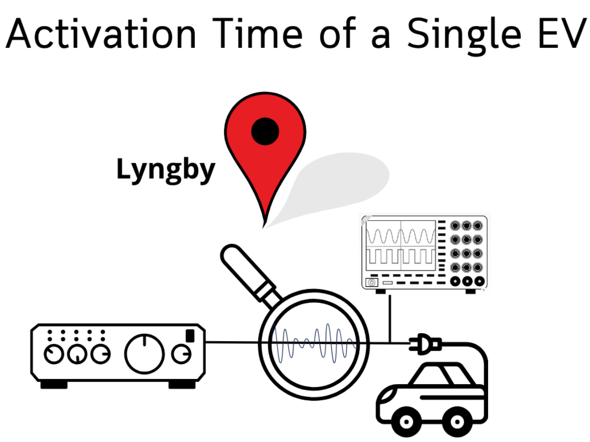

Figure 25: Visualization of tests regarding activation time of a single EV presented in subsection 5.2.

Coordinated frequency control of EV parks, examined in subsection 5.3, evaluates the coordinated behavior of the full system when subjected to simulated grid signals. This includes an assessment of algorithmic accuracy and control stability across multiple chargers operating in parallel under distributed logic.

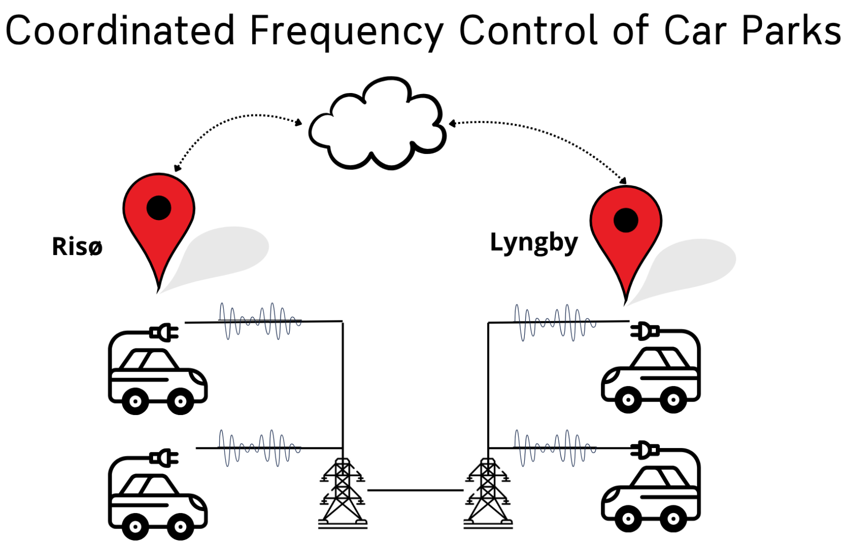

Figure 26: Visualization of tests regarding coordinated frequency control of EV parks presented in subsection 5.3.

Finally, a loop time sensitivity analysis, detailed in subsection 5.4, presents results from an investigation of the differences in average loop time. It was initiated because a significant difference was observed between charger 1 and the remaining three chargers during the testing day. Among other results, the tests include differences between the four physical chargers, simulation vs. physical charging and increase of loop time when the system is scaled with 30 simulated cars.

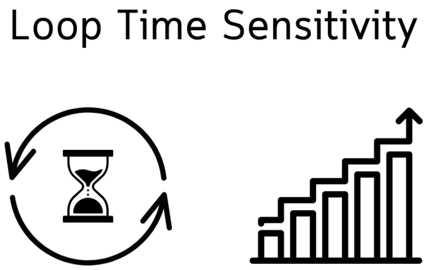

Figure 27: Visualization of tests regarding loop time sensitivity presented in subsection 5.4.

No dedicated test results are presented for FCR, as its requirements and control behavior are very similar to those of FCR-N. The main distinction lies in a wider frequency deviation range and a small dead-band allowance. Therefore, FCR-specific tests were omitted to avoid redundancy.

All results presented in this section are based on experiments conducted with the EVs listed in Table 7 and depicted in Figure 11. Simulations were—apart from subsection 5.4—not used to produce any of the results but served solely as a tool to accelerate control development.

5.1 Coordinated Operation of Car Parks

This section illustrates the performance of the system using multiple test scenarios with real-time current measurements. It shows that the system performs reliably under standard charging conditions without external disturbances.

The objective is to verify whether EVs are charged according to user-defined demands while respecting the global power and current constraints of the system. This includes correctly allocating available power among connected EVs based on their priorities, the aggregated current limit set by the PCC charging infrastructure and the overall power reference of the spot bid. The plots illustrate the allocation logic and demonstrate that the charging power is distributed efficiently and according to priority, laying the groundwork for evaluating more dynamic scenarios in the following sections.

5.1.1 Adjusting Max Current

Figure 28 illustrates the behavior of a 1 phased EV entering a charger cluster with no EVs charging. The EV receives a priority from the *Control Board* website and begins charging. After a series of measurements reveal consistent deviations between the setpoint and the measured current, the system updates the vehicle's maximum charging current accordingly.

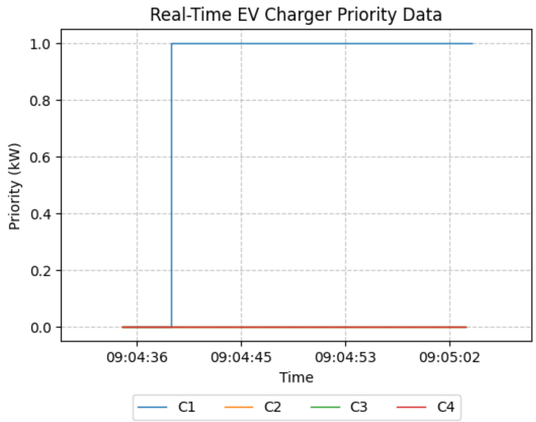

(a) EV receives initial charging priority.

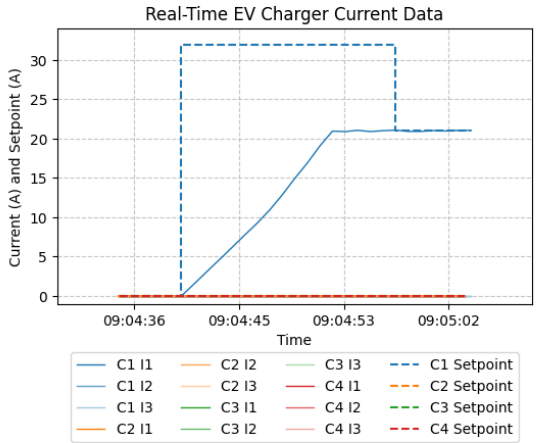

(b) Current setpoint is adjusted based on steady deviations.

Figure 28: A 1 phased EV enters an empty charger cluster and begins charging. Following several measurements showing consistent deviation between setpoint and actual current, the system updates the car's maximum current.

5.1.2 New EV Allowed Charging

In Figure 29, a new 1 phased EV arrives while two other vehicles are already charging. The new EV is granted a sufficiently high priority, allowing it to charge. Consequently, lower-priority vehicles are turned off/down to maintain compliance with the 15 kW spot bid limit.

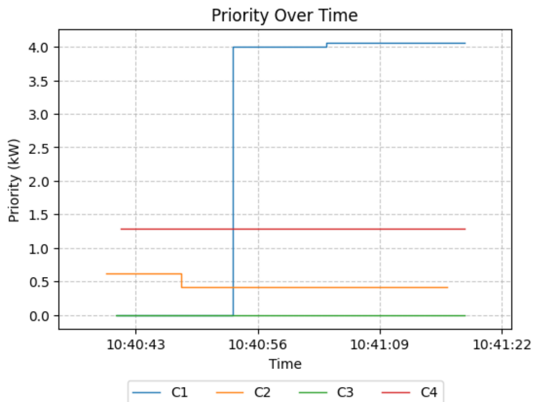

(a) New EV receives highest priority.

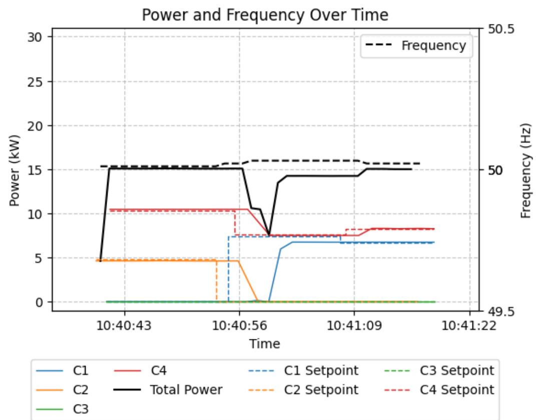

(b) Other EVs reduce power to maintain power cap.

Figure 29: A 1 phased EV (C1) with high priority enters while two other EVs are charging. Due to the 15 kW spot bid constraint, lower-priority EV (C2) stops charging to accommodate the new entrant. The small steps in the priority plot can by accounted for by the regular, interval-based updates of priority described in subsubsection 4.2.4.

5.1.3 Adjusting Priority Based on Phases

Figure 30 shows a 3 phased EV entering an empty charger cluster. Upon recognizing three phased charging, the system divides the EV's priority by three—as priority is defined per phase. The default assumption is that an EV is single-phased, hence the vehicle initially aims to charge at 32 A. However, charging at 32 A across 3 phases would exceed the 8 kW spot bid limit. Thus, the EV's power setpoint is reduced.

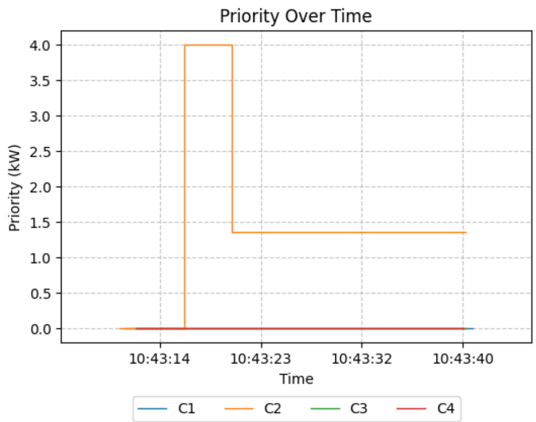

(a) Priority divided by 3 upon detection of 3 phased charging.

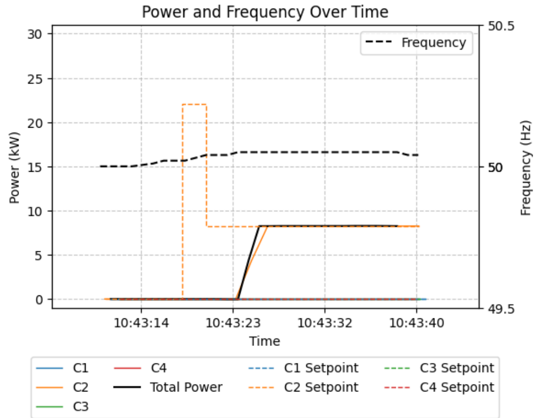

(b) Power setpoint reduced to comply with spot bid.

Figure 30: A 3 phased EV adjusts both priority and power setpoint to remain within the 8 kW spot bid constraint. Priority is normalized per phase, and the setpoint is scaled down to avoid violating the power reference.
It appears that the priority is adjusted before the EV begins drawing significant power, but the actual charging current is very low at that point. The algorithm only updates the priority once current is detected on all three phases.

5.1.4 Periodic Priority Update

In intervals of 5 minutes, priorities are recalculated based on the amount of energy the EVs have charged respectively. Figure 31a demonstrates that C3 overtakes C4 in priority, prompting the system to reallocate charging current. C4 ramps down its charging to enable C3 to charge. After an initial setpoint of 32 A, the max current of C3 is recognized, allowing C4 to charge at a lower power than previously, adhering to the spot bid of 11 kW. It can be noted that at 10:58:39, the summed power consumption is a bit higher than 11 kW and when only C4 was charging. This is due to C3 consuming $16\text{A} \cdot 1\phi \cdot 230\text{V} = 3.68\text{kW}$ leaving $11\text{kW} - 3.68\text{kW} = 7.32\text{kW}$, which in amperes becomes $\frac{7.32\text{kW}}{230\text{V} \cdot 3\phi} = 10.6\text{A}$, which is rounded to 11 A, creating the small overshoot observed.

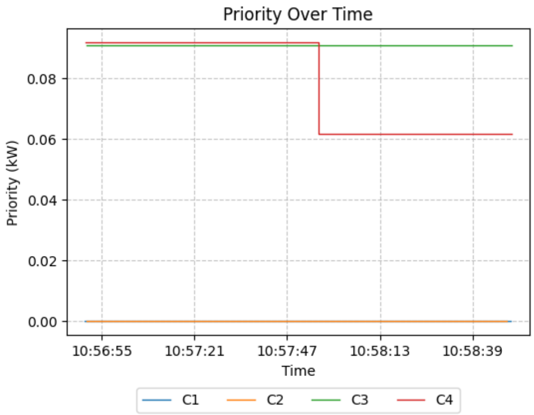

(a) Priority shift in favor of C3.

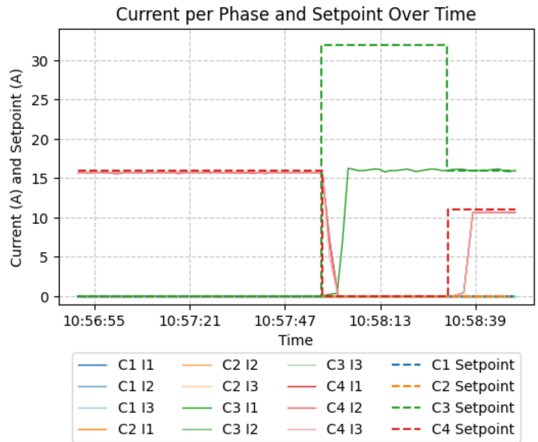

(b) Charging switches from C4 to C3.

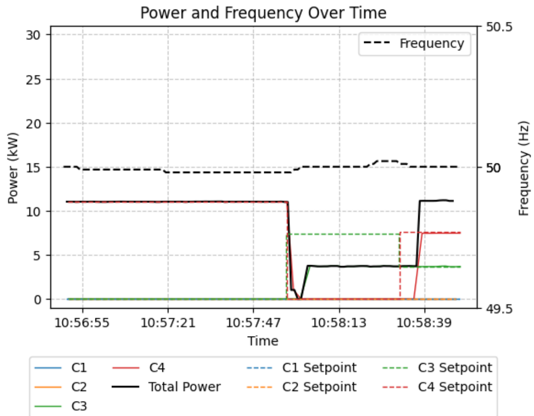

(c) C4 and C3 adhering to spot bid after max current is diagnosed.

Figure 31: Periodic updates in priority result in C3 overtaking C4. Since the 11 kW spot bid cannot accommodate both, C4 ramps down to allow C3 to charge.

5.1.5 Compliance with PCC Limit

Figure 32a shows how the PCC constraint limits total current to 32 A, resulting in a maximum combined power consumption of 15 kW, despite a spot bid of 20 kW. This demonstrates that local constraints can overrule market-based allocations.

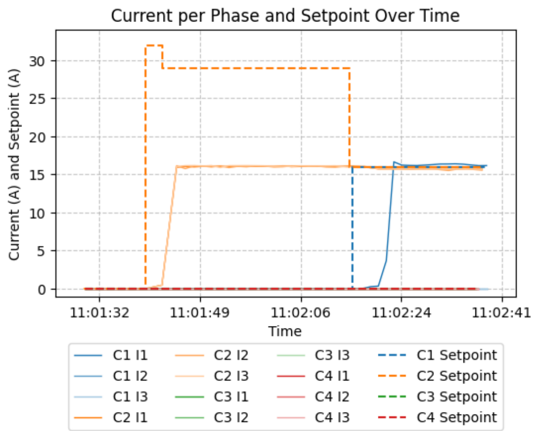

(a) Combined current capped by PCC limit.

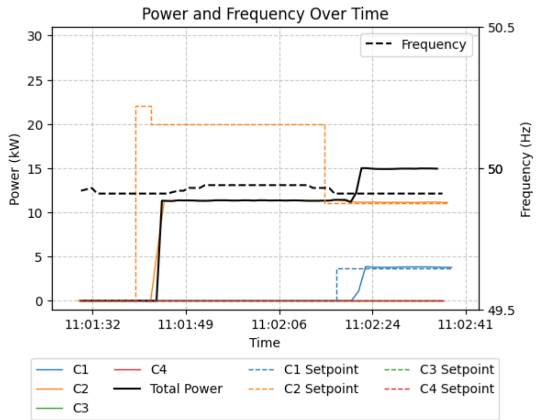

(b) Total power restricted to 15 kW.

Figure 32: Despite a 20 kW spot bid, PCC limits total current to 32 A, leading to a maximum load of 15 kW across C1 and C2. This exemplifies the prioritization listed in subsection 1.1.

5.1.6 Spot Bid Alteration

Figure 33 depicts a scenario where the spot bid is reduced from 20 kW to 15 kW via the *Control Board*. This results in the lowest-priority EV presently charging—C1—ramping down its power consumption accordingly.

Figure 33 additionally illustrates intelligent behavior of the control where a car makes space for the car with marginally lower priority. In this case, C4 ramps down from 16 A to 14 A on three phases to make space for C1 to charge 7 A on 1 phase. This results in an aggregated power closer to the spot bid compared to C4 charging at 16 A and C1 shutting down completely. The algorithm works because information is communicated to all chargers, resulting in the same dispatch list being calculated across chargers.

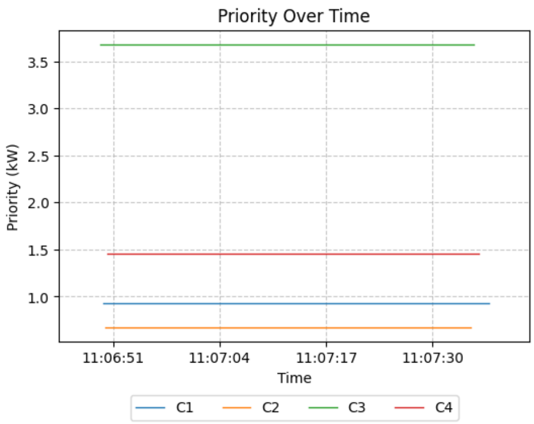

(a) Fleet priorities at time of spot bid reduction.

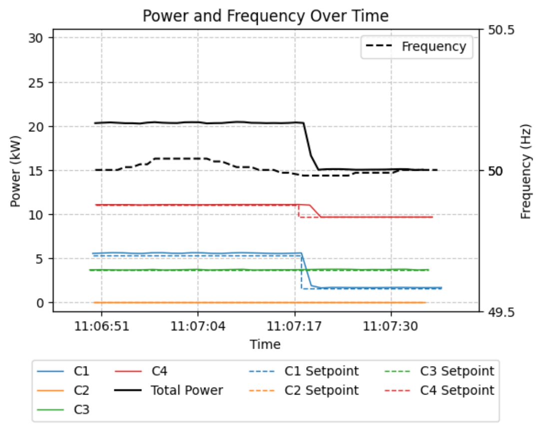

(b) Power reduction to match updated spot bid.

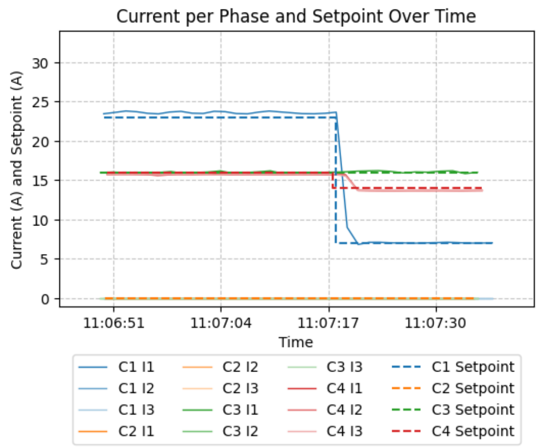

(c) Current reduction to match updated spot bid.

Figure 33: The spot bid is reduced from 20 kW to 15 kW, prompting the two lowest-priority charging EVs (C1 and C4) to reduce their charging power accordingly. This also displays the intelligent control listed in Figure 19.

5.1.7 Reaching Maximum SOC

Maximum SOC can be set manually by the EV owner. It will often be set to 80% to minimize battery degradation and low efficiency charging. Figure 34 illustrates C2 reaching its maximum SOC, causing it to stop charging. The control system responds by detecting a maximum allowable current of 0 A and adjusting C2's priority accordingly. Once this condition is identified, C4 receives a setpoint corresponding to the full spot bid. However, since C4 is limited to charging at a maximum of 16 A, C1 compensates by charging the remainder to fulfill the total spot bid of 16 kW.

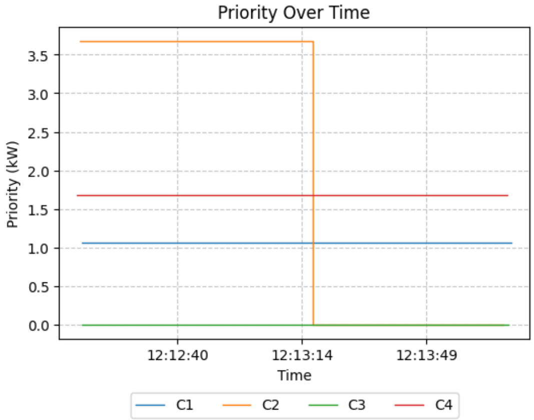

(a) Assigned priorities for the EVs during the charging session.

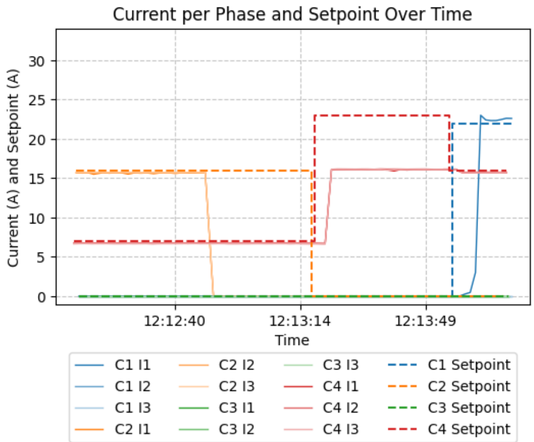

(b) Current of C2 drops after reaching maximum SOC, enabling C4 and finally C1 to increase their setpoints.

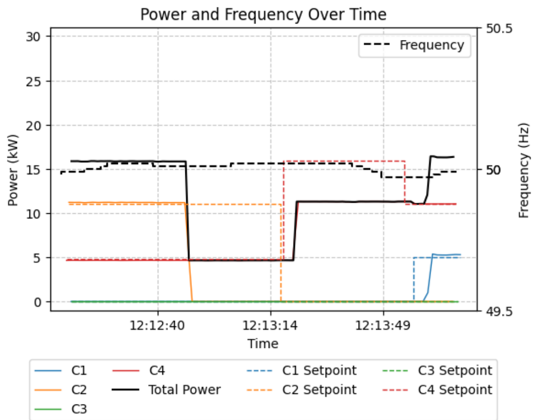

(c) Power of C2 drops after reaching maximum SOC, enabling C4 and finally C1 to increase power consumption.

Figure 34: Behavior of the charging system when C2 reaches its maximum SOC followed by a sharp decrease in power. The system then updates the maximum current of C2 to 0 A, which allows the available charging power to be reallocated to C4.

5.2 Activation Time of a Single EV

The following experiments are based on the setup described in subsection 3.2, focusing on the dynamic response of a single EV to frequency changes. After the initial experiments, the amplifier began triggering the ground fault circuit interrupter, likely due to an internal grounding issue. A replacement amplifier was made available, but it eventually failed to deliver the required voltage on phase two. As a result, only a subset of the planned tests was completed. Nevertheless, reliable data was collected for upregulation tests, where the EVs reduce their charging power. Only limited data was obtained for downregulation (increasing power), consisting of one test with a Tesla Model Y.

5.2.1 Example of Upregulation

Figure 35 presents an example where the C1 charging the Nissan Leaf reacts to a frequency drop. The frequency first decreases, triggering a corresponding reduction in the current setpoint. Shortly thereafter, the Leaf reduces its charging current accordingly. The treatment of data follows the approach described in subsubsection 3.2.1.

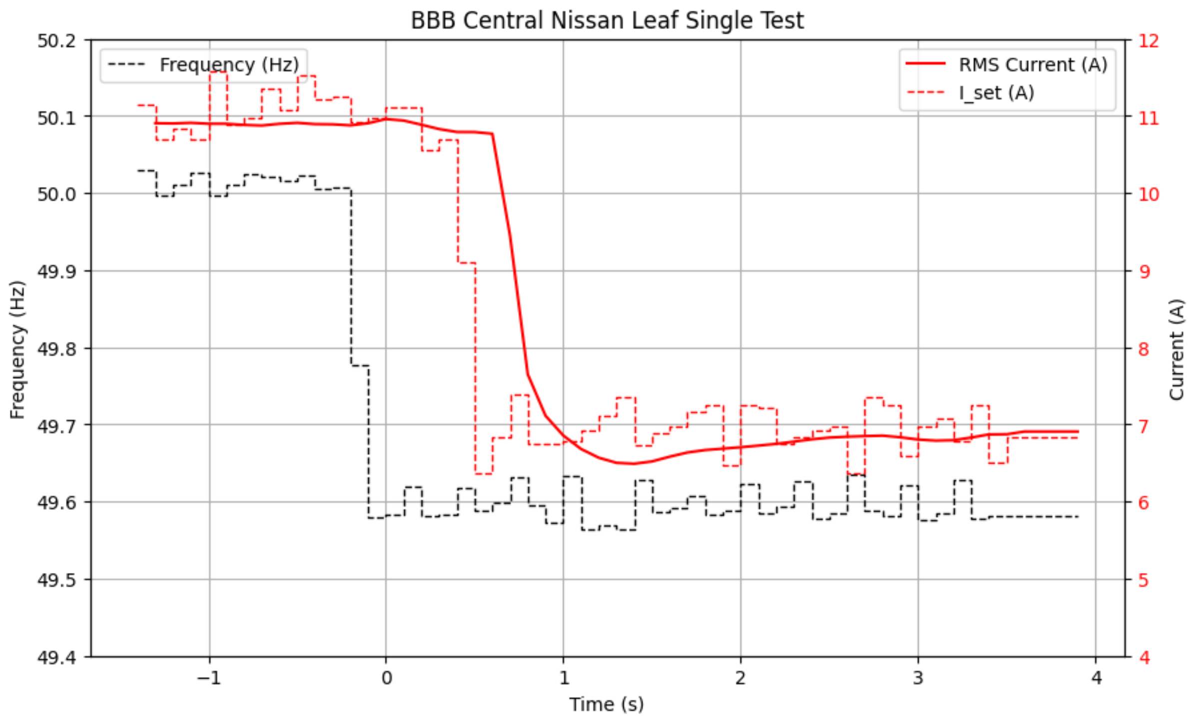

Figure 35: Demonstration of the base-case test setup where the Nissan Leaf responds to a frequency drop induced by the power amplifier. The central server calculates the setpoint using the vehicle's priority, spot bid, and PCC. The frequency-based correction is computed and applied locally on the BBB. No other BBBs were publishing to EDDK during this test.

As to address the apparent instability of frequency and setpoint in Figure 35, we present another plot based on the same raw data output from the oscilloscope, but where Figure 35 applies windows of 100 ms as described in subsubsection 3.2.1, Figure 36 presents the effect of increasing the window size to 330 ms, smoothing out oscillations, but decreasing temporal resolution. Note that the frequency drops to ≈ 49.42 Hz which is closer to the 49.45 Hz set on the amplifier.

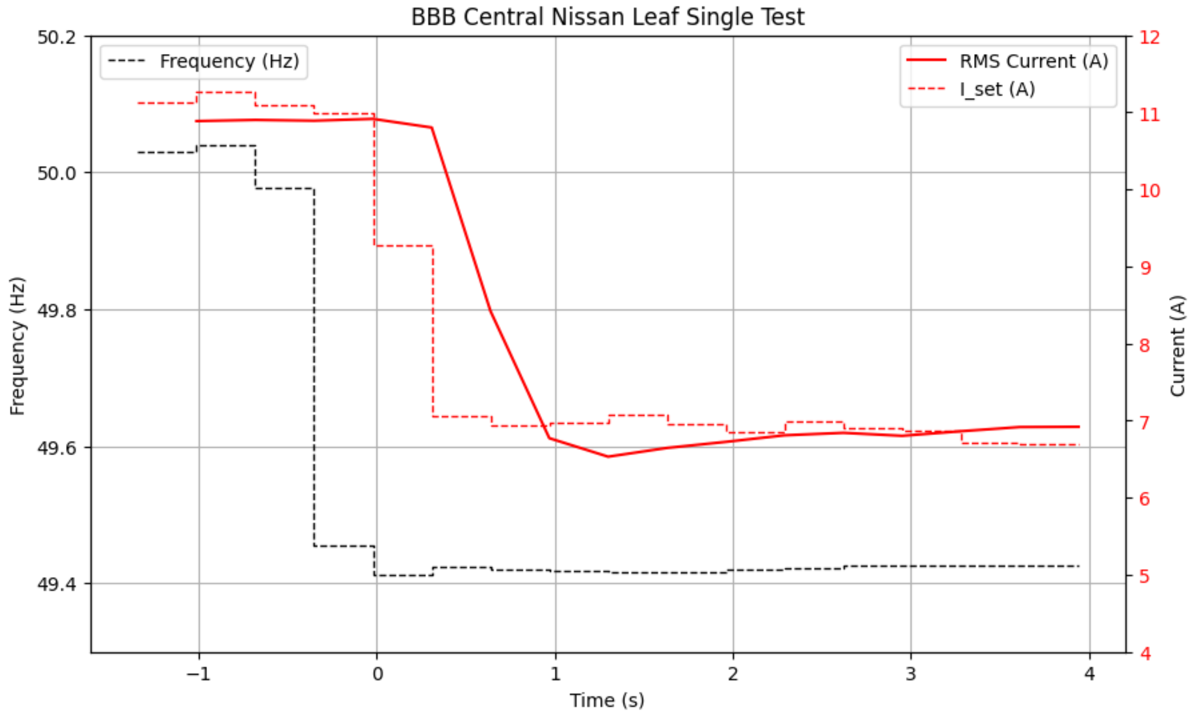

Figure 36: Visualization of the effects of increasing window size to 330 ms, smoothing out uncertainty of measurements, but decreasing temporal resolution.

5.2.2 Average Upregulation across 5 Scenarios

Figure 37 illustrates the average system response, showing both the measured current setpoints and actual currents. In the base case, the Nissan Leaf at Lyngby is running with control flag 2, receiving centrally calculated setpoints and reacting to frequency deviations based on local measurements by the DEIF and calculations by the BBB. Only 1 charger was actively publishing. In this case, the average current activation time was 840 ms, well within the FFR threshold of 1.3 s given in Table 4. The shaded area in Figure 37 visualize ±2 standard deviations of the 10 tests conducted.

All timestamps have been altered based on the methodology described in subsubsection 3.2.1.

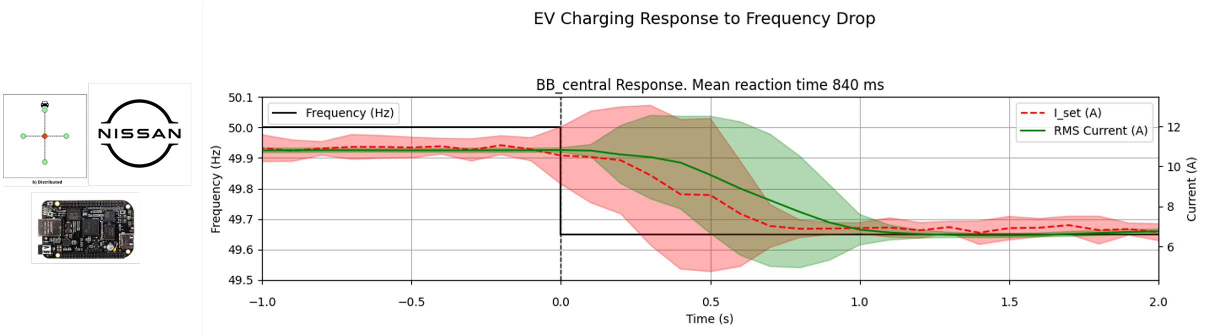

Figure 37: Base case reaction with ±2 standard deviations in the shaded area. The base case consists of the Nissan Leaf, control flag 2, no other active chargers, and applying the BBB as the local computer. These settings are visualized by the figures on the left of the plot.

As to assess the impact of the chosen setup, certain features were varied one at a time:

1. Changing control flag from 2 to 4
2. Replacing the BBB by a MacBook Pro M3 Pro
3. Replacing the Nissan Leaf by a Tesla Model Y
4. Increasing the number of active chargers from 1 to 4

The configuration switching the EV model to a Tesla Model Y showed the largest difference with a reduced activation time of 330 ms (39%). Decentralizing the computation to individual BBBs increased the reaction time by an average of 30 ms. Conversely, offloading computation to the MacBook Pro reduced the average reaction time by 110 ms. Even with four BBBs publishing to EDDK every second, the system reduced the reaction time by 52 ms relative to the base case. The results are visualized in Figure 38.

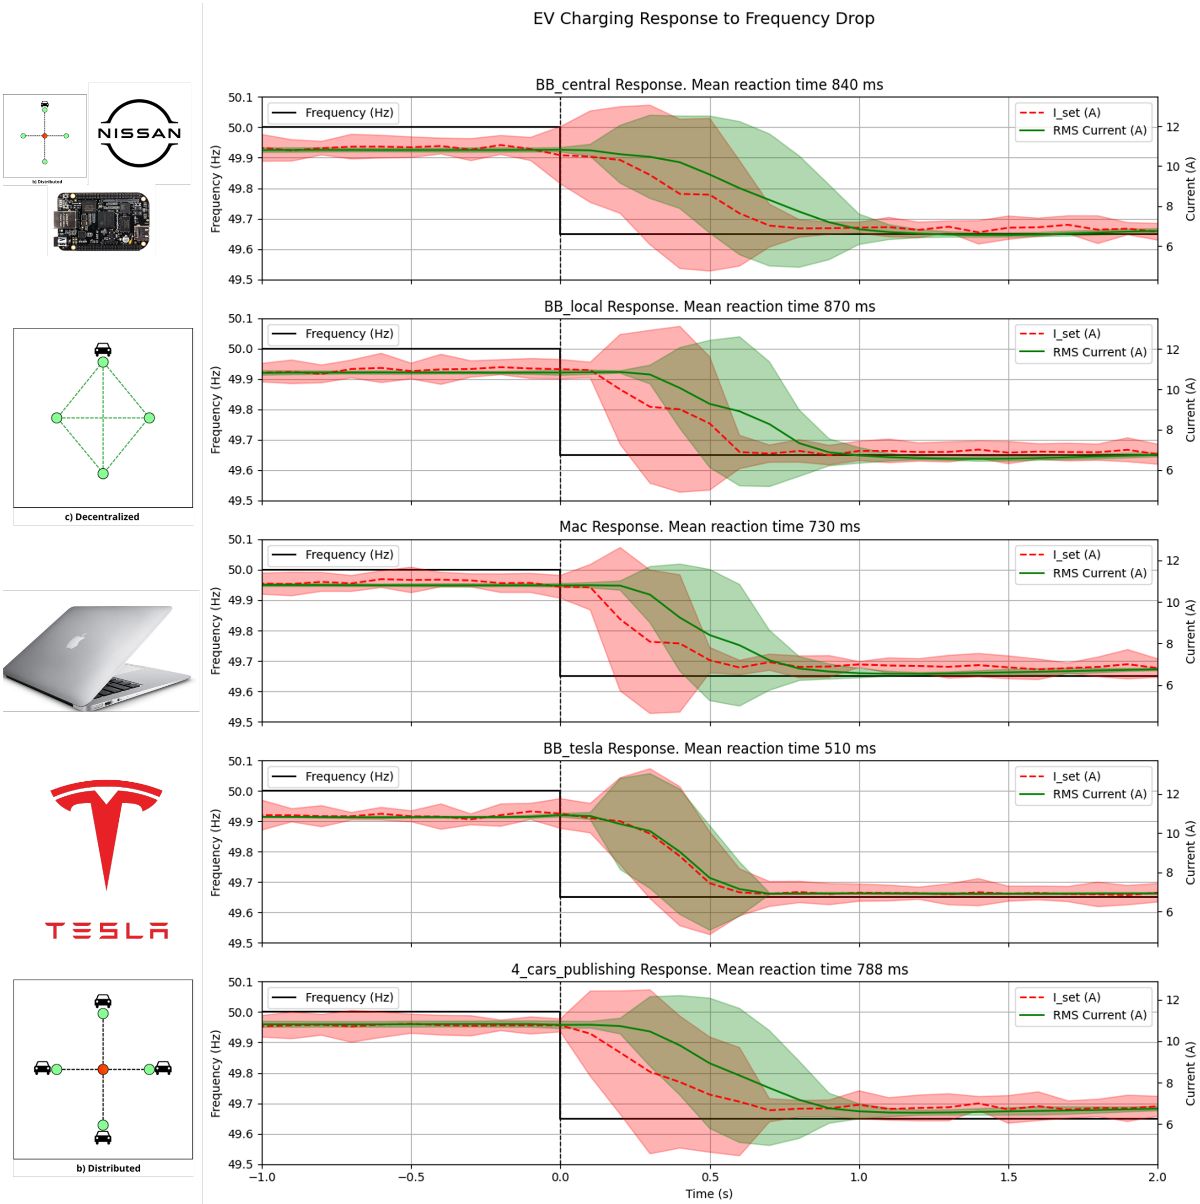

Figure 38: Average upregulation response for various system configurations. Includes different EV models, control architectures, and levels of network activity. Shaded regions represent ±2 standard deviations. The graphics on the left show what feature has been altered for the specific tests.

5.2.3 Downregulation

Due to the amplifier issues described earlier, only one reliable downregulation tests was performed. This involved increasing the current setpoint following a frequency rise. The result is shown in Figure 39, where a noticeably slower current response is observed compared to the upregulation scenario. However, the control setpoint response has similar speed to that seen in the upregulation case presented in Figure 38.

The reaction threshold of 6.8 A is based on Energinet's definition of a valid response, corresponding to 5% of the required steady-state response [41]. The tests suggest that the EV responds fast enough to comply with even the strictest 2-second FCR requirement in DK1 as listed in Table 12. But with only 1 test, nothing can be concluded at this stage regarding downregulation. Additionally, preliminary test showed that the Tesla Model Y ramps slower when ramping from 0 A as displayed in Table 9.

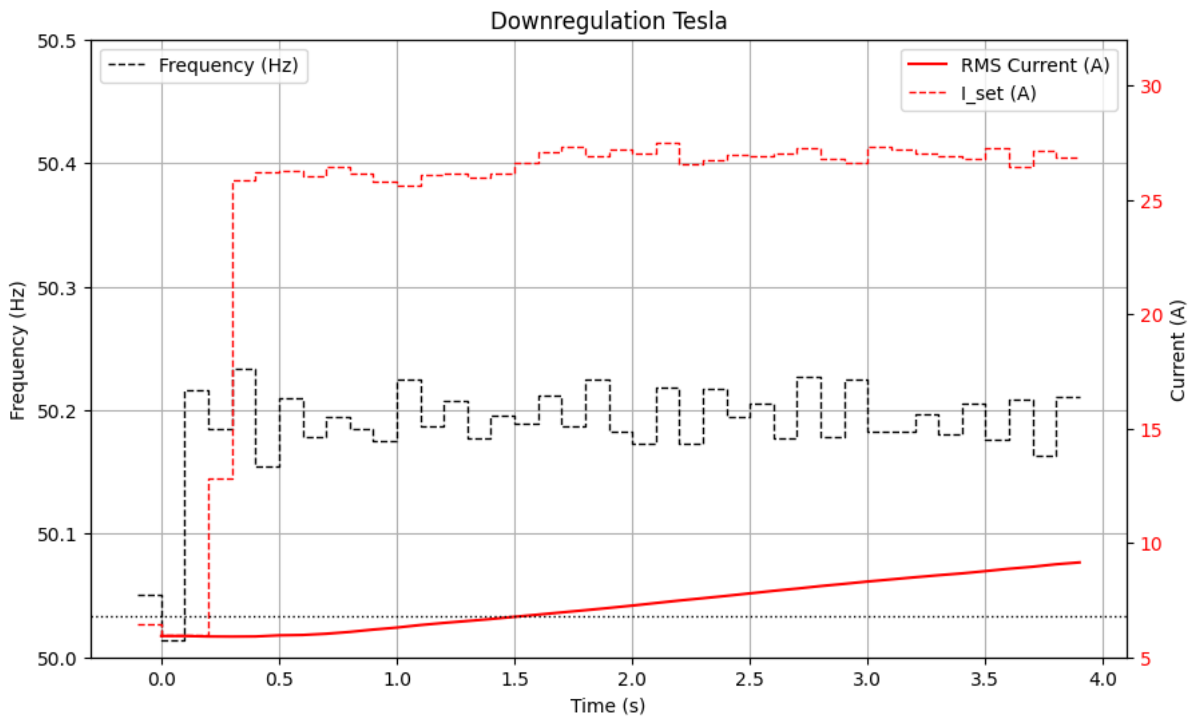

Figure 39: Test where the Tesla Model Y reacts to a frequency increase by ramping up charging power. The current response is considerably slower than during upregulation, while the setpoint response is comparable. The straight dotted line in the bottom depicts the reaction threshold—5 % of required steady state response [41].

Key results and statistics are summarized in Table 12, where it also stands clear that the time between setpoint and current adjustment is around 400 ms for the Nissan Leaf, whereas the delay of the Tesla is around 60 ms, when dealing with upregulation.

Table 12: Summary of test results. All values are in milliseconds except number of tests.

System case	Mean	Std	95% CI	Min	Max	Tests
Current Response						
BBB central	840	185	469 - 1211	600	1100	10
BBB local	870	200	470 - 1270	500	1200	10
Mac	730	142	446 - 1014	600	1000	10
BBB Tesla	510	130	250 - 770	200	700	10
4 cars publishing	788	169	449 - 1126	500	1100	8
Down reg. Tesla	1400	0	-	1400	1400	1
Setpoint Response						
BBB central	440	150	141 - 739	200	600	10
BBB local	480	209	62 - 898	100	800	10
Mac	330	119	93 - 567	200	500	10
BBB Tesla	450	136	178 - 722	200	700	10
4 EVs publishing	375	179	18 - 732	100	700	8
Down reg. Tesla	100	0	-	100	100	1

5.3 Coordinated Frequency Control of EV Parks

While Figure 38 demonstrated individual EV responsiveness focusing on the response time, this subsection evaluates how the control algorithm scales across a larger fleet with less focus on the speed of the reaction. Note that since EDDK (and therefore *Visuals*) updates occur at a frequency of 1 Hz, actual vehicle reaction speeds might be faster than those visualized.

5.3.1 FCR-D Up Step Test

The FCR-D Up step test assesses the fleet's ability to reduce power consumption in response to a frequency drop, as shown in Figure 40. In this scenario, the EV fleet is configured with a total spot bid of 20 kW and a frequency regulation bid of 10 kW.

When the frequency drops below the threshold of 49.9 Hz as specified in subsubsection 2.3.2, the system initiates an upward regulation response by curtailing power consumption of the EVs. Based on the priority allocation, C1 and C3—assigned the lowest priorities—are completely curtailed, resulting in their charging power dropping to 0 kW, corresponding to state 6 as listed in Table 10. C4—being in state 3—contributes to the remaining reduction by slightly decreasing its power draw, ensuring that the full 10 kW reduction requirement is met. C2, already in idle state 7, cannot react to the frequency drop. A second drop in frequency happens at ≈ 11:42:32 causing the same response from the EVs. In this case, the frequency stays low for a longer time to test the steady state response which shows to be completely stable. This test confirms the fleet's ability to deliver fast and prioritized upward regulation by shedding load accordingly.

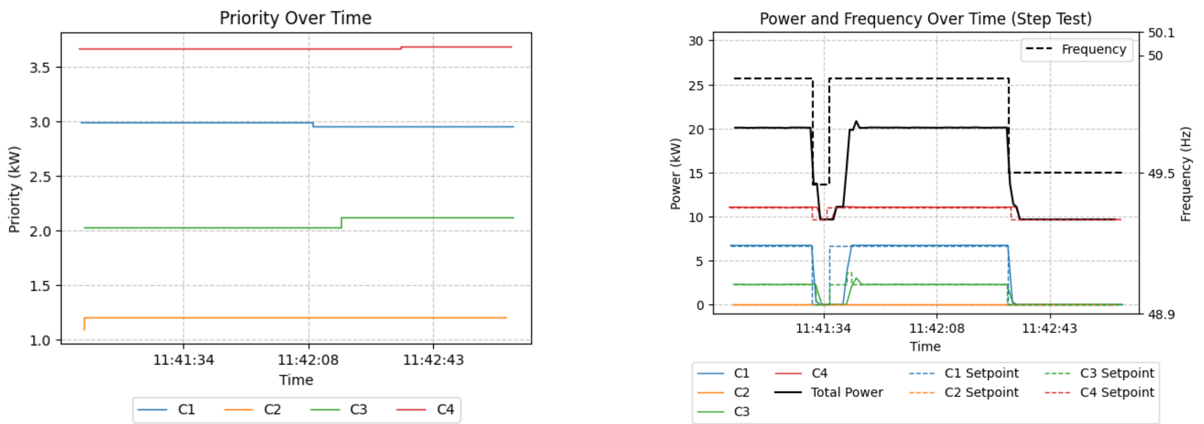

(a) Priorities of the four EVs that are plugged in.

(b) Charging power of C1 and C3 drops to 0 kW, while C4 slightly reduces by 2 kW.

Figure 40: During a simulated frequency drop, the system must reduce 10 kW of charging power. C1 and C3, having the lowest priorities, are curtailed entirely (state 6). Additional reduction is supplied by C4 (state 3) to meet the total frequency bid.

5.3.2 FCR-D Down Step Test

The FCR-D Down step test demonstrates the EV fleet's capability to increase power consumption in response to a frequency rise using the algorithm described in Table 11. In Figure 41, the system is tasked with delivering 10 kW of downward regulation, which involves ramping up the charging load when frequency exceeds the upper limit of 50.1 Hz.

Initially, the fleet responds to the step signal by increasing charging power across all EVs. During the first 5-second step, approximately 75% of the required power is delivered, indicating a fast but not fully saturated ramp. In the following extended step, the fleet stabilizes, and the charging power of C1, C3, and C4 continues to meet the regulation demand. C1 and C4 ends in state 2, whereas C3 is in state 3. C2, assigned the lowest priority, returns to idle (state 5) once the error between measured and required power drops below 20%. The result confirms that the system can provide accurate and stable downward regulation through prioritized power allocation and dynamic rebalancing among the chargers. Furthermore, the test indicates a compatibility with the requirements for FCR-D Down as described in Table 3, where the power should be at least 86% 7.5 s after the frequency deviation.

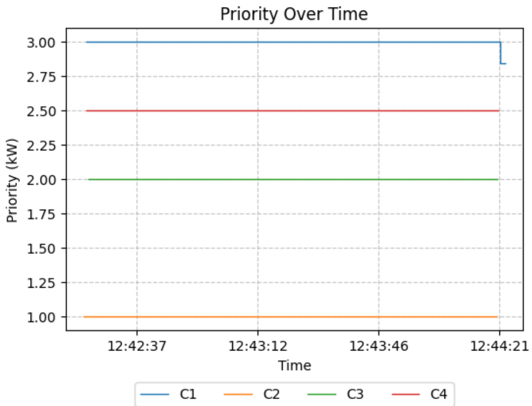

(a) Priorities of the four EVs in the fleet.

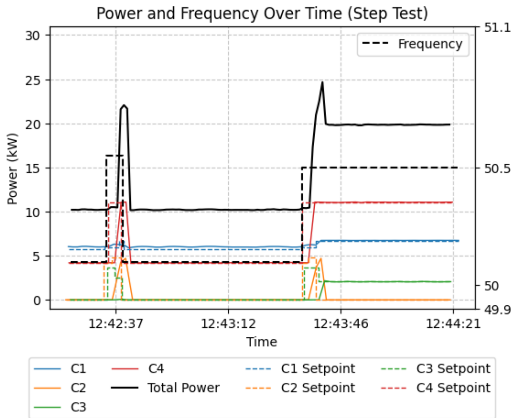

(b) Aggregate power consumption ramps up across all EVs.

Figure 41: Following a frequency increase, all EVs in the fleet ramp up charging to meet the regulation requirement. During the initial 5-second step, about 75% of the required power is achieved. In the longer subsequent step, C1 (state 2), C3 (state 3), and C4 (state 2) maintain increased consumption according to their priority as the error between consumed and required power drops below 20%. C2 returns to idle (state 5).

5.3.3 FFR Step Test

The FFR step test evaluates the system's ability to react instantly to a sudden drop in frequency, as shown in Figure 42. The frequency deviation prompts a rapid power reduction among the charging EVs, according to their predefined priorities. In this particular test, four EVs were plugged in, and the system was configured with a 20 kW spot bid and a 10 kW frequency regulation bid.

As the frequency drops, all charging EVs respond by curtailing their charging power in a prioritized manner. Charger C3, having the lowest priority of charging EVs, turns off charging (state 6). Next EV in the priority line is C4, which ends in state 4, meaning that if C4 was to turn down the required amount of power alone, it would be charging in the uncontrollable range. Therefore, it sets a setpoint of 6 A corresponding to 4.14 kW, which leaves a bit of power reduction for the highest priority EV—C1. C1 being single phased is charging at 6.4 kW and 28 A. Therefore, it is far away from the uncontrollable range and handles the remaining power reduction as described in state 3. Once the 5-second FFR event concludes, the charging behavior is restored to its pre-disturbance levels, demonstrating that the EVs can deliver short-term frequency support with sufficient intelligence to avoid overshoots while maintaining their primary charging task with minimal disruption.

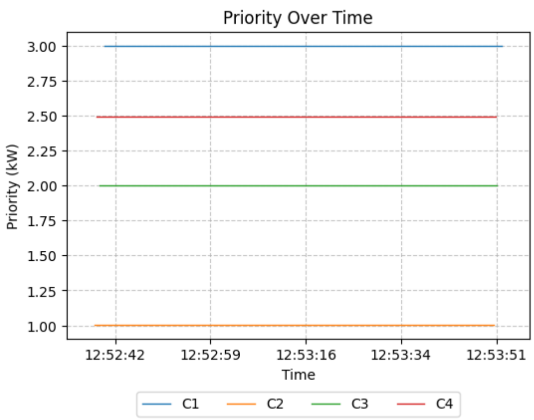

(a) Priorities of the 4 EVs plugged in.

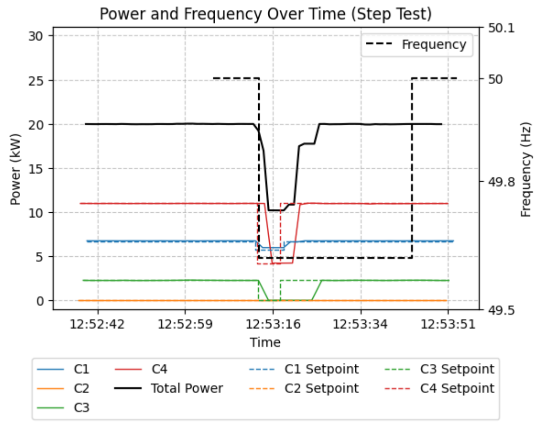

(b) Power drop of all 3 charging EVs.

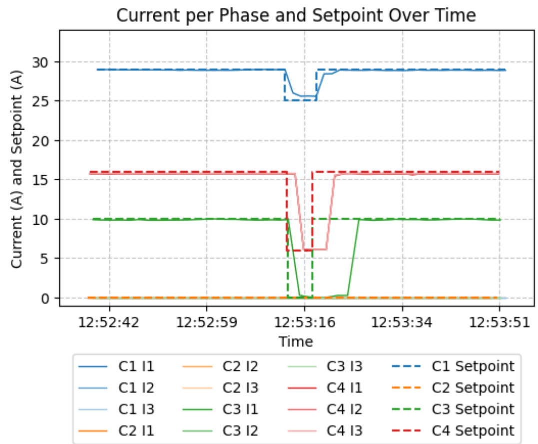

(c) Current of C4 drops to 6 A, forcing C1 to control the remaining power reduction.

Figure 42: Frequency drop forces C3 to stop charging (state 6). C4 ends in state 4 as listed in Table 10 and turns down to minimum controllable current, forcing C1 to control the remaining upregulation (state 3).

5.3.4 FCR-N Sine Test

The FCR-N sine test explores the stability and responsiveness of the distributed EV control strategy under continuously varying frequency conditions. A sine wave profile simulates grid fluctuations, although the range of the oscillations provided by Energinet is rarely observed in the Danish power grid [42].

Figure 43 shows the power draw of four EVs exposed to an oscillating frequency. The system configuration comprised 15 kW spot bid power reference and 5 kW for frequency regulation in both directions. The EVs demonstrate robust and stable behavior, adjusting their charging power dynamically in response to the sine wave.

Notably, the control algorithm shows a more accurate and immediate reduction in current during frequency dips than during frequency recoveries. This asymmetry is consistent with the natural delay in ramping up current due to EV OBC limitations. Nevertheless, the overall tracking performance remains smooth and well-damped, validating the system's effectiveness in participating in continuous frequency regulation (FCR-N) without compromising stability or charging intent. Furthermore, it is very clear that the ramping of currents is much faster, when the current ranges from 10-16 A or 7-11 kW as displayed by C4. Both C1 and C2 are experiencing trouble with setpoints alternating between 0 A and something higher, whereas the stable setpoint of C3 allows a steady charge.

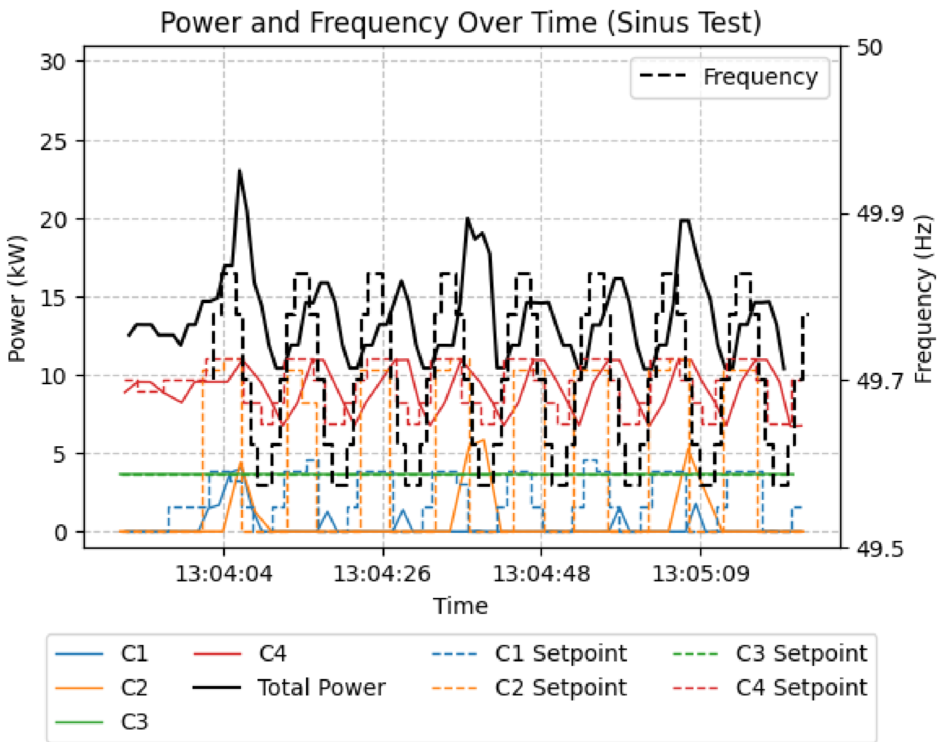

Figure 43: Proof of stability during frequency oscillations. Here in the shape of a sine wave. It can be seen that the EVs are better at turning down their current rather than the opposite.

5.4 Loop Time Sensitivity

As part of the main experiment day described in section 5, it was unexpectedly observed that charger C1, connected to a Nissan Leaf in Lyngby, exhibited significantly faster loop times compared to the other chargers. This observation prompted a dedicated follow-up investigation to identify and isolate the factors contributing to the discrepancy. While these supplementary tests were conducted on a different day—introducing some contextual differences—the resulting data offers valuable insight into the control loop performance and component-induced delays under varying conditions.

To ensure representative results, each test was run for a duration of 10 minutes. Most tests use control flag 3, in which the priority-based current dispatch is calculated locally on each BBB. Although this deviates from the base case configuration used earlier in the results section, one test also explores the effect of performing the dispatch calculation centrally for comparison. For readability, this section presents subsets of the test results, while the complete dataset is available in Table 17 in the appendix.

5.4.1 Simulated Loop Times across BBBs

As seen in Table 13, the first 4 tests show no significant variation in average loop times across chargers, when operating under simulated conditions without frequency control and without additional network traffic. This consistency suggests that the previously observed loop time discrepancy cannot be attributed to fundamental differences between the BBBs and their respective internet connections. Hence, also ruling out the modems and EDDK as potential sources of the discrepancy in loop times.

Table 13: Loop time results across BBBs regarding simulating 1 charger at control flag 3.

Test No.	**Charger No.**	**Number of Chargers Live**	**Simulation vs. Physical**	**Control Flag***	**Average loop time [ms]**	**CPU Load on BBB from `top`** [%]**
1	C1	1	Simulation	3	105	46.6
2	C2	1	Simulation	3	103	45.6
3	C3	1	Simulation	3	103	47.0
4	C4	1	Simulation	3	108	45.9

* Control flags indicating whether or not to respond to frequency deviation and where the priority-based current dispatch is to be calculated. See Table 8.
** Running the command `top` in a Linux system displays the active processes in descending order of CPU usage.

5.4.2 Physical vs. Simulated Loop Time

Introducing the physical layer (i.e., sending and receiving signals to and from the DEIF and CC) causes an increase in loop time. Specifically, Table 14 displays that the difference in average loop time between Test 1 (simulation only) and Test 5 (physical ID4) is approximately 300 ms. This overhead is attributed to hardware interaction latency.

Table 14: Comparing simulation and physical average loop times of C1

Test No.	Charger No.	Number of Chargers Live	Simulation vs. Physical	Control Flag	Average loop time [ms]	CPU Load on BBB from `top` [%]
1	C1	1	Simulation	3	105	46.6
5	C1	1	Physical - ID4	3	407	24.6

5.4.3 Replacing the CC

Charging the ID4 using the other chargers resulted in a significantly increased average loop time of approximately 1100 ms (comparing Tests 5 through 8), thereby confirming the discrepancy initially observed during the main result generation day. Following consultation with the lab technicians, the CC on C2 was replaced with a newer unit of the same model. As shown in Test 9 in Table 15, this replacement reduced C2's average loop time to a value similar to that of C1. This suggests that the delay observed on C2, C3, and C4 was likely caused by differences in firmware or hardware versions of the originally installed CCs.

Table 15: Comparing the average loop times of different chargers charging ID4. Test number 9 displays the effect of replacing the old CC in C2 with a new.

Test No.	**Charger No.**	**Number of Chargers Live**	**Simulation vs. Physical**	**Control Flag**	**Average loop time [ms]**	**CPU Load on BBB from `top` [%]**
5	C1	1	Physical - ID4	3	407	24.6
6	C2 - old CC	1	Physical - ID4	3	1513	11.7
7	C3	1	Physical - ID4	3	1460	11.7
8	C4	1	Physical - ID4	3	1537	11.7
9	C2 - new CC	1	Physical - ID4	3	401	20.8

5.4.4 Replacing the ID4 with Nissan Leaf

To assess whether the choice of vehicle model influenced loop time, Test 10 replaced the ID4 with the Nissan Leaf (comparing Tests 5 and 10). The resulting change in average loop time was minor, indicating that the vehicle model has limited impact on control loop duration under these conditions.

Table 16: Displaying the impact of substituting the ID4 with a Nissan Leaf.

Test No.	**Charger No.**	**Number of Chargers Live**	**Simulation vs. Physical**	**Control Flag**	**Average loop time [ms]**	**CPU Load on BBB from `top` [%]**
5	C1	1	Physical - ID4	3	407	24.6
10	C1	1	Physical - Nissan Leaf	3	479	24.9

5.4.5 Assessing Scalability

To assess the system scalability, additional tests were conducted with varying numbers of live chargers publishing data. A simulation was conducted on charger C4 to investigate whether the increase in loop time scales linearly or exponentially with the number of active chargers. In this test, the duration of each run was reduced to 1 minute to allow for a higher number of tests.

Figure 44 illustrates how the average loop time increases as more chargers begin publishing to EDDK. Four control strategies were tested, corresponding to different control flags:

- **Control flag 1 (blue):** The charger receives current setpoints directly from the central server without performing local frequency control.
- **Control flag 2 (yellow):** The charger receives central setpoints but modifies them in response to the frequency deviations shown in Figure 43.
- **Control flag 3 (green):** The charger performs local priority-based current dispatch and does not respond to frequency deviations.
- **Control flag 4 (red):** The charger performs local priority-based current dispatch and responds to frequency deviations.

The results suggest that under control flag 1, the loop time remains relatively constant regardless of the number of chargers, indicating minimal processing overhead from subscribing to and receiving messages from a larger number of chargers. In contrast, both control flags 2, 3 and 4 exhibit a linear increase in loop time with the number of active chargers as outlined in [34].

Control flag 2 adds the computational burden of frequency control on top of the central dispatch, while control flag 3 requires local computation of the current dispatch without frequency control. Both contribute noticeably to increased processing time.

Control flag 4, which combines both local dispatch computation and frequency-based adjustments, results in the highest loop times. These are approximately equal to the loop times under control flag 1 plus the combined overheads of flags 2 and 3. This additive behavior aligns well with the modular design of the control logic, where each feature contributes independently to overall computational load.

It is worth noting that when applying control flag 2—the baseline used in frequency activation time results presented in Figure 38—the system exhibits a loop time increase of approximately 1.97 ms for each additional active charger. This increase directly impacts the latency of frequency-based control responses.

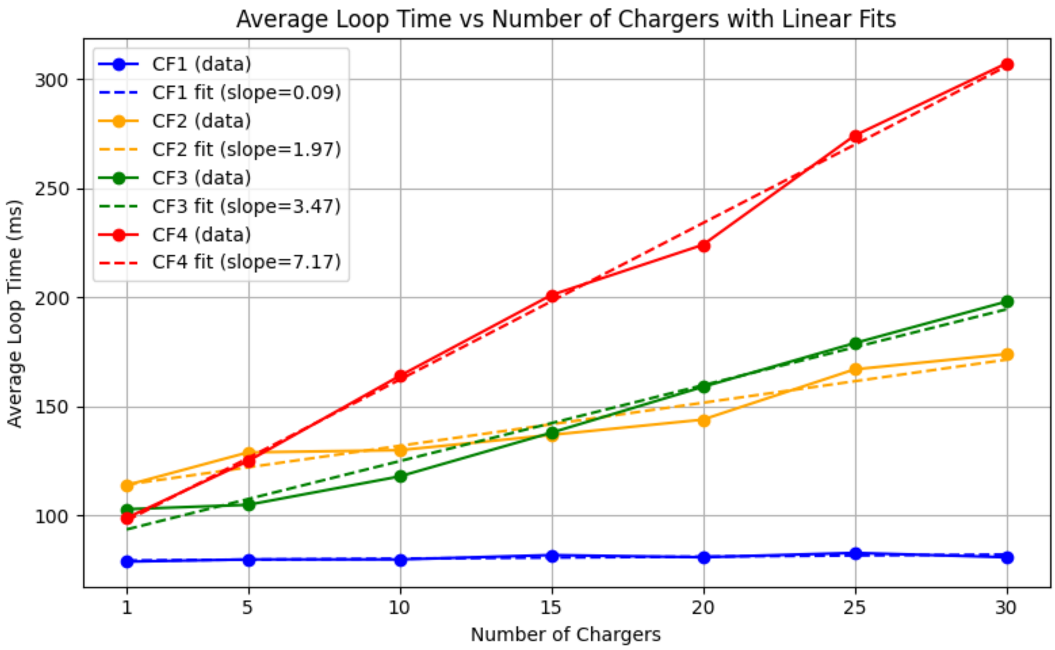

Figure 44: Loop time as a function of the number of publishing chargers. Blue: control flag 1 (central dispatch only), yellow: control flag 2 (central dispatch with frequency control), green: control flag 3 (local dispatch without frequency control), red: control flag 4 (local dispatch with frequency control). Lines display linear trend fits.

5.4.6 CPU Load

Finally, CPU load measurements taken using the Linux `top` command show an increase in processing demand, when scaling from 1 to 30 active chargers (Test 10 to 11) as seen in Table 17. In this case, the CPU usage rises by 20 percentage points (from 24.9% to 44.2%), corresponding to an estimated 0.6 percentage point increase per added charger.

With regards to average loop time vs. CPU load, we observe an inversely proportional relation as displayed in Figure 45.

6 Discussion

This section critically reflects on the developed control architecture and algorithms, addressing both their strengths and limitations. It discusses technical standards that may constrain large-scale deployment, highlights potential improvements to enhance user experience and system resilience, and outlines key directions for future work.

The discussion begins in subsection 6.1 with an analysis of the control protocols and the implications of the priority scheme. In subsection 6.2, attention shifts to technical and regulatory factors, including current standards and hardware constraints. Finally, subsection 6.3 describes areas for continued research and system development, particularly with respect to system expansion, long-term performance, and advanced diagnostic or measurement techniques.

6.1 Control Protocols and Priority Scheme

6.1.1 The Priority as User Input

The efficiency and fairness of the control system in this thesis depend heavily on users understanding kWh values and honestly entering information into the *Control Board*. An improvement would be enabling communication between the EV and charger to share the SOC, allowing users to simply enter a desired SOC and departure time. Alternatively, users could input current SOC, target SOC, battery capacity, and departure time, though this may be inconvenient. This inconvenience could be reduced through a registration system storing user data, but this solution would only accommodate the issue for charging sessions with one's regular supplier. Another solution could involve AI-based license plate recognition to identify the EV model, infer battery capacity, and suggest energy needs based on typical user behavior. For workplace chargers, priorities could shift toward ensuring sufficient charge for the commute home.

A more simple approach could be to replace priority with the amount of time a given EV had been parked, where more hours meant higher priority. The "priority" could then be set to 0, whenever the charge of the EV was complete and the current went to 0 A. While this eliminates user interaction, it risks introducing inefficiencies and may prevent some EVs from reaching their intended SOC, particularly in high-demand scenarios.

6.1.2 Considering Max Power in Priority

Our measure of priority could be altered to also consider the max power of a given EV. In a situation where a user enters a priority that is equal to the amount of power the EV can actually draw, but still insufficient to be allowed to charge, the given EV would fail to charge the amount demanded by the user. Likewise, when performing upregulating frequency control, we do not consider whether ramping down the charge of a given EV would result in it being unable to reach its desired amount of energy charged. A workaround could be that the control algorithm prioritizes the downramping of the charging EVs with the largest gap between their priority and max power per phase.

6.1.3 Synchronized Updates of Priority

The priority of each EV updates every 5 minutes to avoid oscillations between almost equally prioritized EVs. In our case, the 5 minutes start, when the EV is given a user input, which in practice may result in more frequent passing of priority and therefore the right to charge. To avoid this the time of priority updates should be synchronized.

6.1.4 Feedback Loop and Max Current

The method for estimating an EV's maximum charging current assumes an initial value of 32 A, which is then adjusted based on observed deviations between setpoint and actual current. An improvement to this strategy would be to adopt a more dynamic algorithm. For instance, an initial setpoint of 18 A could be selected—this value is sufficient to diagnose most 16 A-limited EVs while remaining within reach of higher-capacity vehicles. The algorithm could then iteratively increase the current in steps (e.g., 2 A), monitoring whether the EV follows the commanded setpoint. Once a divergence is detected, the rounded mean observed current over the last couple of samples could be interpreted as the vehicle's true current limit.

A limitation of the existing approach is the absence of a reset mechanism for the maximum current. This is particularly problematic in scenarios where an EV begins charging with a cold battery, limiting its ability to draw high currents [47]. In such cases, the algorithm may prematurely conclude a lower max current than what the EV could achieve once the battery reaches optimal temperature.

While [10] presents a feedback-loop mechanism based on deviations between aggregated park setpoints and real-time power measurements from the P_IoT device, this project has chosen the approach presented in subsubsection 4.2.5, as knowing the max current of each EV is essential when performing downregulation. However, this method has inherent limitations. Deviations smaller than the defined threshold of 1 A are ignored, which may result in under-utilization of available charging power. For a three phased EV, this could lead to a potential power undershoot of:

$$1\ \text{A} \cdot 3\ \phi \cdot 230\ \text{V} = 0.690\ \text{kW}$$

The undershoot is per EV and could potentially accumulate for a large car park. More critically, persistent current draws more than 1 A below the setpoint will cause the algorithm to continually downscale an EV's maximum current—potentially to 0 A—effectively disabling its ability to charge.

Although these issues did not arise in the experimental setups used in this project, the integration of a dynamic feedback loop could enhance robustness and prevent such edge-case failures.

6.1.5 Imbalance Costs Caused by Ramping Constraints

Our code is designed to always respect the spot bid, meaning that current is dispatched until the next ampere (considering phases) would result in an over-consumption compared to the spot bid. This results in the consumption often undershooting the spot bid. Furthermore, the upramping of EVs is not instantaneous, as seen in the results section or exemplified in Figure 28. Hence, the average power consumed will, although not quantified, most likely be below the spot bid. In Denmark, imbalances are settled at 15-minute resolution [48]. To avoid imbalance costs, the developed code could be enhanced by monitoring the aggregated consumption in 15 minute periods and adjusting the power reference accordingly, so all energy bought on the spot market is utilized.

6.1.6 EDDK as Communication Platform

Implementing distributed control reduces reliance on a central controller, thereby minimizing vulnerability to single points of failure such as cyberattacks. However, in our current setup, all communication still routes through EDDK, effectively creating a new point of failure. A potential mitigation strategy is to enable direct communication between all BBBs, for example via SSH tunnels or VPN connections. While this would improve system resilience, it introduces additional configuration complexity.

A major limitation of using the EDDK as the communication platform is its low update frequency of 1 Hz, which introduces unnecessary system delays. Introducing a higher frequency on updates would most likely reduce the overshoots observed in Figure 41. Faster inter-charger communication could be achieved either by adopting a communication platform that supports higher update rates or by implementing the direct communication approach described above.

6.1.7 Scalability

As shown in Figure 44, the loop time exhibits an approximately linear increase with respect to the number of actively publishing chargers. Specifically, under control flag 2, the slope is approximately 1.97 ms per additional charger. This control strategy, which combines central setpoint calculation with local frequency adjustment, is considered the best compromise between scalability and response speed.

Despite favorable scalability characteristics, the system's responsiveness under control flag 2 is bounded. As shown in Table 12, the 95% confidence interval for the current reaction time extends up to 1211 ms. Given the FFR service requirement of a 1300 ms maximum response time, the remaining margin of 89 ms would be exhausted after the addition of approximately 45 more EVs. Beyond this point, the system would fail to comply with FFR constraints more than 5% of the time.

To address this, a hierarchical control strategy—such as architectures H1 or H2 described in Figure 5—could be employed for larger fleets. In such an architecture, the EV park is partitioned into divisions of manageable size. Each division is associated its own dedicated *Central Control* script, executing the same priority-based dispatch algorithm as outlined in this thesis. Periodically, a higher-level controller aggregates the status of each division and computes optimal fractions of the spot market and frequency regulation bids for each division based on their available capacity and occupancy. This hierarchical approach enables the system to preserve fast response times while accommodating a significantly larger number of EVs.

6.1.8 Dummy Charging

Early in the development process, dummy charging was implemented to maintain charger functionality during internet outages. However, selecting an appropriate current setpoint proved difficult. A low current, such as 6 A, leads to inefficient power conversion, while assigning a higher current to an uncontrollable charger reduces overall system flexibility. This raises the question of whether a charger without internet connectivity should be allowed to charge at all—a decision that ultimately depends on the size of the charger cluster and the frequency of connection losses. In any case, the charger should clearly indicate to the user that it is operating in dummy charging mode.

While the code adjusts the PCC current capacity based on the readings on the DEIF as described in subsubsection 4.2.1, correction of the power reference as a result of dummy charging has not been implemented. The implementation hereof should not be difficult, but is important to address if the system is to implemented in real life.

6.2 Technicalities and Standards

As the code in this project was developed with the available hardware in mind, certain adaptations are necessary to ensure compatibility with current commercial standards.

6.2.1 Setpoint Update Rate

According to IEC 61851-1:2019 [19], Table A.6 specifies: *"In normal operation, during the 5 s allowed adjustment time (t10–t9) the EV supply equipment shall not initiate a new sequence 6 for changing the PWM"*, effectively capping the update rate for current setpoints at 0.2 Hz. Under typical operation, this is not an issue, as a 5-second delay has minimal impact on tracking spot market bids. Additionally, the control logic can be configured to only send setpoint updates to the CC only when a change in setpoint is detected.

However, this limitation poses a significant challenge for frequency control applications, which require response times faster than 5 seconds. On the other hand, one could argue that frequency regulation constitutes a special operating mode and is thus not covered by the "normal operation" constraint in the standard.

A potential workaround is to phase-shift the control loop start times across the fleet, creating a staggered reaction pattern. This could result in a portion of EVs responding quickly enough to meet performance requirements, but the effectiveness of this strategy would require further testing and validation.

6.2.2 Phase Management

As shown in Figure 32, the PCC is fully utilized by one single- and one 3-phased EV without the spot bid being fully met. Had C1 been a 3-phased EV, the spot bid consumption could have been met. However, not all cars are made for 3-phased charging, hence a more realistic alternative could be better phase management. All single-phased EVs in our project draw current from phase L1. Several commercial solutions for better phase management are available. The simple—but suboptimal—solution is manually phase-shifting chargers during installation. A more elegant approach is using chargers with integrated dynamic phase balancing tailored to ongoing charging sessions [49]. With this approach, three single-phased EVs with identical max current would occupy the same space on the PCC and consume the same amount of power as a single 3-phased EV. Studies show that better phase balancing can improve capacity usage rate from 45

Furthermore, phase management is recognized as a crucial part of ensuring operational reliability of the power grid, especially in the low-voltage distribution grid, as described in [50]. In 2022, the Danish Safety Technology Authority published the regulation *Fællesregulativet*, requiring electrical installations drawing more than 16 A to distribute current evenly across phases [51].

In contrast to [8], this thesis chose to divide priority by the number of phases, since priority is defined as the minimum average power per phase over the course of the charging period. This approach is considered fair, as it reflects the urgency of charge more consistently regardless of how many phases an EV uses. However, this design choice may lead owners of 3-phased EVs to feel de-prioritized, effectively favoring the inferior single-phase technology rather than incentivizing users to adopt 3-phased EVs.

6.2.3 Vehicle Diagnostics

As discussed in subsubsection 2.1.2, present EV charging communication standards offer limited access to vehicle-specific diagnostics. However, several parameters critical to control—such as maximum charging current, ramping times (up/down), charging efficiency, battery capacity, number of phases, SOC at arrival, and deviation from setpoint—are mostly linked to the vehicle model, with SOC being the exception. Therefore, a significant amount of diagnostic data could be inferred from a single input, assuming researchers compile a database of vehicle-specific parameters.

For this project's control strategy, the two most essential diagnostics are the number of phases and the maximum current, which are addressed in subsubsection 4.2.3 and subsubsection 4.2.5. The timing of these diagnostics is non-trivial. For instance, in FCR downregulation, knowing each EV's upward power potential is crucial for effective control.

Vehicle diagnosis could be implemented by briefly assigning maximum setpoints to newly arrived vehicles and then turning them off again. This would reveal ramping characteristics, maximum current, number of phases, and deviation from setpoint, improving control efficiency. However, this process would disrupt ongoing charging sessions, reducing overall power utilization. A less intrusive approach would be to diagnose only when a vehicle's priority permits charging.

For longer charging sessions, AC–DC conversion efficiency of the OBC becomes increasingly important. Following the findings of [14], the ACDC project assumes a linear efficiency decline from 90% at maximum power to 80% at minimum [8]. In contrast, this project assumes 100% efficiency, as precise energy tracking was not a primary objective.

Another accuracy improvement lies in finer setpoint granularity—e.g., allowing decimal values. The error between setpoint and actual current could be diagnosed and characterized as either proportional or static, and subsequently compensated for in both priority-based dispatch and frequency regulation. As an example, it was observed that the Nissan Leaf would draw 0.6 A more than the setpoint allowed. If this was diagnosed and accounted for, the PWM signal transmitted each loop iteration could be adjusted by the diagnosed deviation.

6.3 Future Work

This subsection serves to highlight the work that was completed, but is regarded important in the improvement of the control system and utilization of EV flexibility in general.

6.3.1 High-Resolution Activation Testing of Single EV

The results in subsection 5.2 demonstrate the response time of a single EV to frequency changes, measured with a 100 ms temporal resolution. At this resolution, frequency and current setpoint measurement accuracy was limited, as shown in Figure 35. To improve accuracy without reducing temporal resolution, the test could be repeated using an oscilloscope with a higher sample rate.

Further enhancements could include estimating additional component delays, e.g. by overlaying the DEIF's frequency measurements in Figure 35 to distinguish between DEIF measurement delay and delay of the control; and quantifying the delay between changes in the I_{set} value from the BBB via Modbus TCP and the I_{set} value derived from the PWM signal. These tests would assist in determining which component to optimize or replace.

6.3.2 Investigating CPU Load Impact on Control Performance

At this stage, it is not possible to draw firm conclusions about the relationship between average loop time and CPU load. In our measurements, periods of high loop time are generally associated with low CPU load. This suggests that the observed increases in loop time are primarily caused by delays in physical components (e.g., communication or actuator response times), rather than computational bottlenecks. During these physical delays, the CPU is idle or waiting, resulting in a lower observed CPU load. It remains possible that this relationship would change if CPU load approached 100%, at which point computation time could become the limiting factor.

6.3.3 System Expansion

If the control is to be implemented in real applications, extensive tests including more EVs at different locations must be conducted, as the relation between the loop iteration time and adding a new location is under-investigated. Further, the overhead on computational load of the *Central Control* with increasing EV fleet has not been examined. However, it is not expected to be relevant, as the speed of transmission of calculated setpoints is not crucial, when staying within reasonable borders e.g. once a minute.

A consideration when facing increased processing requirements in a larger system is the choice of microprocessor. The BBB offers a cost-effective solution with sufficient processing power and a broad set of functionalities that align well with the use case in this thesis. However, if scaling to a significantly larger EV fleet increases the loop iteration time by more than 89 ms, a more powerful processor or offloading computations should be evaluated, if FFR provision remains a priority.

6.3.4 Long-Term Stability and Performance Testing

While the control system has demonstrated satisfactory performance during short charging sessions, tests over extended periods—spanning several hours or days—have not yet been conducted. Although key long-term functionalities, such as handling new EV connections and completing charging sessions, appear to function correctly, the current testing does not provide sufficient evidence to claim a fully reliable setup for sustained operation.

6.3.5 Synchronized Multi-Charger Measurement Setup

Even though the tests conducted in subsection 5.2 indicate compatibility with all up-regulating services, further tests including synchronized oscilloscopes on all chargers are recommended. This would shed light on the performance of all vehicle simultaneously. The setup including synchronized oscilloscopes across locations proved logistically difficult, hence it was not prioritized during the work of this thesis. Prior to this test it is essential that charger 3 and 4 get their CCs updated to new units to match the control speed of charger 1 and 2.

6.3.6 Extending Control Strategies to Restoration Reserves

Regarding frequency regulation markets, this thesis has focused exclusively on developing and demonstrating compliance with frequency containment reserves, driven by the novelty of integrating a microcontroller in a charger and developing a compatible control scheme to reduce response time. However, compliance with restoration reserves, which relies more on accurate forecasting of EV cluster flexibility, remains equally important. [52] presents a chance-constrained optimization model for EV aggregators to bid in the Nordic FCR-D market under Energinet's P-90 requirement, showing potential annual savings of 6–10% using data from 1,400 Danish charging stations. As shown in section 7, the economic potential of aFRR in DK1 and DK2 significantly surpasses that of other services. It would require more flexibility than bidding on the FCR-D market, but the annual cost could be brought down substantially more than 10 % for users with a lot of flexibility. In summation, maintaining a broad research scope in EV flexibility is essential to fully support and optimize grid stability.

7 Conclusion

This thesis has developed and implemented control algorithms for operation of 2 EV charger parks capable of delivering Fast Frequency Reserve (FFR). This has been achieved by installing and configuring electrical components on the same local network within each charger. The setup included a BeagleBone Black microprocessor for running the control algorithm, a DEIF for reading key electric parameters, a modem to provide public internet connection and a Phoenix Charge Control for sending current setpoint signals to the EV's onboard charger. This setup was replicated on 4 chargers in total, evenly distributed across the DTU departments in Lyngby and Risø.

The work done in this thesis comprises the establishment of communication between the internal components of the charger; development of control algorithm for standard operation of 2 car parks; development of control strategies for frequency up- and down-regulation with specialized algorithms for both services based on the ramp rate properties of EVs; a website for monitoring the key electric parameters of the entire EV fleet; and a website for sending control signals to the EV fleet combined with a graphical user interface.

The standard operation of the EV fleet integrates several core functionalities to ensure effective and flexible control. Users can input charging demands, and charging is scheduled based on a dynamic priority system. The control continuously adjusts maximum current and the number of charging phases as needed, while updating priorities every 5 minutes. It also ensures compliance with the capacity limits at the point of common coupling and aligns consumption with the spot market bid. Additionally, charging is reduced or stopped automatically once the maximum state of charge is reached.

Further, it is indicated that the control algorithm enables the EV fleet to perform up-regulating frequency control. This indication is based on our single EV base case result displaying reactions to frequency deviations within 840 ms on average with a 95% confidence interval extending to 1211 ms, which is below the activation requirement of 1300 ms for Denmark's fastest reserve, FFR. Even though only 1 successful test focusing on the reaction time of downregulation was conducted, the results pointed towards compatibility with a reaction time of 1400 ms—well below the 2 s requirements of FCR. Additionally, the tests conducted using lower time-resolution indicate compatibility with the strictest ramping requirements of downregulation being 86% power activated in 7.5 s as specified for FCR-D Down. Control flag 2 is recommended as the preferred default mode, since *Central Control* handles the core computations, ensuring that local charger CPUs remain lightly loaded, while performing frequency regulation locally. Local dispatch via flags

3 and 4 is intended as a fallback mechanism in the event of *Central Control* failure or cybersecurity incidents.

In addition, a dedicated loop time analysis revealed that physical component delays are the main contributors to increased loop times. The control loop time increased linearly by 2 ms per active charger. This scaling behavior directly affects control latency and should be considered in future larger-scale implementations of the control algorithm.

While much testing has been done on the specific functionalities, more work is required in determining the functionality of the control algorithm during operational tests spanning several hours and days, as well as further testing regarding response time of downregulation.

References

[1] United Nations. *The Paris Agreement*. 2015.

[2] Eriksen, J. V., Franz, S. M., Steensberg, J., Vejstrup, A., Bosack, M., Bramstoft, R., et al. "The future demand of renewable fuels in Germany: Understanding the impact of electrification levels and socio-economic developments". *Heliyon*. 2023 Nov. 9. e22271. DOI: `10.1016/j.heliyon.2023.e22271`.

[3] Ratnam, K. S., Palanisamy, K., and Yang, G. "Future low-inertia power systems: Requirements, issues, and solutions - A review". *Renewable and Sustainable Energy Reviews*. 2020 May. 124. P. 109773. DOI: `10.1016/j.rser.2020.109773`.

[4] Mahmud, I., Medha, M. B., and Hasanuzzaman, M. "Global challenges of electric vehicle charging systems and its future prospects: A review". *Research in Transportation Business & Management*. 2023 Aug. 49. P. 101011. DOI: `10.1016/j.rtbm.2023.101011`.

[5] Mobility Denmark. *Nu er der 400.000 elbiler i Danmark*. 2025. URL: `https://mobility.dk/nyheder/nu-er-der-400-000-elbiler-i-danmark/`.

[6] Secchi, M., Ivanova, A., and Eichman, J. "EV mobility diffusion and future perspectives in the EU: results from the FLOW project". *IET Conference Proceedings*. 2023 Nov. 2023. Pp. 1–8. DOI: `10.1049/icp.2023.2678`.

[7] Fang, T., Jouanne, A. von, Agamloh, E., and Yokochi, A. "Opportunities and Challenges of Fuel Cell Electric Vehicle-to-Grid (V2G) Integration". *Energies*. 2024 Nov. 17. DOI: `10.3390/en17225646`.

[8] Marinelli, M. ; Striani, S. ; Pedersen, K. L., Sevdari, K. ; Hach, M. ; Mikkelsen, O. L., et al., *General rights ACDC project-Autonomously Controlled Distributed Chargers Final report*. Tech. rep. 2023.

[9] Striani, S., Pedersen, K. L., Engelhardt, J., and Marinelli, M. "Experimental Investigation of a Distributed Architecture for EV Chargers Performing Frequency Control". *World Electric Vehicle Journal*. 2024 Aug. 15. P. 361. DOI: `10.3390/wevj15080361`.

[10] Pedersen, K. L., Striani, S., Engelhardt, J., and Marinelli, M., "Implementation of priority-based scheduling for electric vehicles through local distributed control". In: *2024 IEEE PES Innovative Smart Grid Technologies Europe (ISGT EUROPE)*. IEEE, Oct. 2024, pp. 1–5. ISBN: 979-8-3503-9042-1. DOI: `10.1109/ISGTEUROPE62998.2024.10863249`.

[11] Simone Striani, T., *DTU Wind and Energy Systems Department of Wind and Energy Systems EV clustering methods for flexibility services*. Tech. rep. 2024.

[12] Striani, S., Unterluggauer, T., Andersen, P. B., and Marinelli, M. "Flexibility potential quantification of electric vehicle charging clusters". *Sustainable Energy, Grids and Networks*. 2024 Dec. 40. DOI: 10.1016/j.segan.2024.101547.

[13] Sevdari, K., Calearo, L., Andersen, P. B., and Marinelli, M. "Ancillary services and electric vehicles: An overview from charging clusters and chargers technology perspectives". *Renewable and Sustainable Energy Reviews*. 2022 Oct. 167. DOI: 10.1016/j.rser.2022.112666.

[14] Sevdari, K., Calearo, L., Bakken, B. H., Andersen, P. B., and Marinelli, M. "Experimental validation of onboard electric vehicle chargers to improve the efficiency of smart charging operation". *Sustainable Energy Technologies and Assessments*. 2023 Dec. 60. P. 103512. DOI: 10.1016/j.seta.2023.103512.

[15] Zunino, P., Engelhardt, J., Striani, S., Pedersen, K. L., and Marinelli, M., "Frequency Control in EV Clusters: Experimental Validation and Time Response Analysis of Centralized and Distributed Architectures". In: *2024 IEEE PES Innovative Smart Grid Technologies Europe (ISGT EUROPE)*. IEEE, Oct. 2024, pp. 1–5. ISBN: 979-8-3503-9042-1. DOI: 10.1109/ISGTEUROPE62998.2024.10863251.

[16] Unterluggauer, T., Rauma, K., Järventausta, P., and Rehtanz, C. "Short-term load forecasting at electric vehicle charging sites using a multivariate multi-step long short-term memory: A case study from Finland". *IET Electrical Systems in Transportation*. 2021 Dec. 11. Pp. 405–419. DOI: 10.1049/els2.12028.

[17] Amazon. *EV + Replacement Electric Vehicle Charging Cable - Type 2 IEC 62196 - PHEV Charger Plug - 3 Phase 32A (22kW) Black/White*. 2025. URL: https://www.amazon.co.uk/EV-Electric-Vehicle-Charger-Charging/dp/B08QQCKQYQ.

[18] Phoenix Contact GmbH & Co. KG, *User manual EV Charge Control Standard-compliant control of the Control Pilot and Proximity Plug interfaces between the electric vehicle and charging station*. Tech. rep. URL: https://asset.conrad.com/media10/add/160267/c1/-/en/000554895ML02/gebruiksaanwijzing-554895-em-cp-pp-eth-oplader.pdf.

[19] Dansk Standard, *DS/EN IEC 61851-1:2019 Konduktive opladningssystemer til elkøretøjer – Del 1: Generelle krav*. Tech. rep. 2019. URL: www.ds.dk.

[20] Calearo, L., Marinelli, M., and Ziras, C. "A review of data sources for electric vehicle integration studies". *Renewable and Sustainable Energy Reviews*. 2021 Nov. 151. P. 111518. DOI: 10.1016/j.rser.2021.111518.

[21] Knez, M., Zevnik, G. K., and Obrecht, M. "A review of available chargers for electric vehicles: United States of America, European Union, and Asia". *Renewable and Sustainable Energy Reviews*. 2019 July. 109. Pp. 284–293. DOI: 10.1016/j.rser.2019.04.013.

[22] Huang, P. and Ma, Z. "Unveiling electric vehicle (EV) charging patterns and their transformative role in electricity balancing and delivery: Insights from real-world data in Sweden". *Renewable Energy*. 2024 Dec. 236. DOI: 10.1016/j.renene.2024.121511.

[23] Cao, X., Striani, S., Engelhardt, J., Ziras, C., and Marinelli, M. "A semi-distributed charging strategy for electric vehicle clusters". *Energy Reports*. 2023 Nov. 9. Pp. 362–367. DOI: 10.1016/j.egyr.2023.10.014.

[24] Engelhardt, J., Gabderakhmanova, T., Rohde, G., and Marinelli, M., "Reconfigurable Stationary Battery with Adaptive Cell Switching for Electric Vehicle Fast-Charging". In: *2020 55th International Universities Power Engineering Conference (UPEC)*. IEEE, Sept. 2020, pp. 1–6. ISBN: 978-1-7281-1078-3. DOI: 10.1109/UPEC49904.2020.9209774.

[25] Kostopoulos, E. D., Spyropoulos, G. C., and Kaldellis, J. K. "Real-world study for the optimal charging of electric vehicles". *Energy Reports*. 2020 Nov. 6. Pp. 418–426. DOI: 10.1016/j.egyr.2019.12.008.

[26] Marinelli, M., Calearo, L., Engelhardt, J., and Rohde, G., "Electrical Thermal and Degradation Measurements of the LEAF e-plus 62-kWh Battery Pack". In: *2022 International Conference on Renewable Energies and Smart Technologies (REST)*. IEEE, July 2022, pp. 1–5. ISBN: 978-1-6654-0971-1. DOI: 10.1109/REST54687.2022.10023130.

[27] Engelhardt, J., Zepter, J. M., Marinelli, M., and Piegari, L. "Efficiency Characteristic and Operating Area of High-Power Reconfigurable Batteries". *IEEE Transactions on Industry Applications*. 2024 Mar. 60. Pp. 3676–3684. DOI: 10.1109/TIA.2023.3344552.

[28] Engelhardt, J., Zepter, J. M., Gabderakhmanova, T., and Marinelli, M. "Energy management of a multi-battery system for renewable-based high power EV charging". *eTransportation*. 2022 Nov. 14. P. 100198. DOI: 10.1016/j.etran.2022.100198.

[29] Ziras, C., Thingvad, M., Fog, T., Yousefi, G., and Weckesser, T. "An empirical analysis of electric vehicle charging behavior based on real Danish residential charging data". *Electric Power Systems Research*. 2024 Sept. 234. DOI: 10.1016/j.epsr.2024.110556.

[30] Bahamonde-Birke, F. J. and Ernst, D. M. "Am I really willing to use my electric vehicle sustainably? A study on the charging preferences of electric vehicle users". *International Journal of Sustainable Transportation*. 2024 Sept. 18. Pp. 744–750. DOI: 10.1080/15568318.2024.2399783.

[31] Engelhardt, J., *Reconfigurable Batteries in Electric Vehicle Fast Chargers: Towards RenewablePowered Mobility*. Tech. rep. DTU Wind and Energy Systems, 2022. DOI: https://doi.org/10.11581/dtu.00000254.

[32] Aghajan-Eshkevari, S., Azad, S., Nazari-Heris, M., Ameli, M. T., and Asadi, S. "Charging and Discharging of Electric Vehicles in Power Systems: An Updated and Detailed Review of Methods, Control Structures, Objectives, and Optimization Methodologies". *Sustainability*. 2022 Feb. 14. P. 2137. DOI: 10.3390/su14042137.

[33] Han, X., Heussen, K., Gehrke, O., Bindner, H. W., and Kroposki, B. "Taxonomy for Evaluation of Distributed Control Strategies for Distributed Energy Resources". *IEEE Transactions on Smart Grid*. 2018 Sept. 9. Pp. 5185–5195. DOI: 10.1109/TSG.2017.2682924.

[34] Nimalsiri, N., Mediwaththe, C., Ratnam, E., Shaw, M., Smith, D., and Halgamuge, S. *A Survey of Algorithms for Distributed Charging Control of Electric Vehicles in Smart Grid*. 2019.

[35] Liu, J., Xu, W., Liu, Z., Fu, G., Jiang, Y., and Zhao, E., "Optimal Operation of Large-scale Electric Vehicles Based on Improved K-means Clustering Algorithm". In: *2022 the 5th International Conference on Robot Systems and Applications (ICRSA)*. New York, NY, USA: ACM, June 2022, pp. 23–28. ISBN: 9781450396486. DOI: 10.1145/3556267.3556280.

[36] Thingvad, A., Ziras, C., Ray, G. L., Engelhardt, J., Mosbak, R. R., and Marinelli, M., "Economic Value of Multi-Market Bidding in Nordic Frequency Markets". In: *2022 International Conference on Renewable Energies and Smart Technologies (REST)*. IEEE, July 2022, pp. 1–5. ISBN: 978-1-6654-0971-1. DOI: 10.1109/REST54687.2022.10023471.

[37] Aziz, M., Huda, M., and Nandiyanto, A. "Opportunity of frequency regulation using electric vehicles in Denmark". *Journal of Engineering Science and Technology*. 2018 May. 13.

[38] Energinet, *Forskrift, C1 - Vilkår for Balanceansvar*. Tech. rep. Dec. 2023. URL: https://energinet.dk/media/d45j2s2m/forskrift-c1.pdf.

[39] Nord Pool AS. *The Power Market - Market members*. 2025. URL: https://www.nordpoolgroup.com/en/the-power-market/The-market-members/.

[40] Energinet, *Forklarende Dokument - Nyt Design for Ubalanceafregning*. Tech. rep. 2024, pp. 7–10. URL: https://energinet.dk/media/ch3dguzm/forklarende-dokument-nyt-design-for-ubalanceafregning-hoering.pdf.

[41] Energinet, *Systemydelser til levering i Danmark. Udbudsbetingelser*. Danish. Tech. rep. Mar. 2025. URL: https://energinet.dk/media/rkvdymux/21_10162-24-

udbudsbetingelser-for-systemydelser-til-levering-i-danmark-mfrr-eam-10877149_2_1.pdf.

[42] Energinet, *Prequalification of Units and Aggregated Portfolios*. Tech. rep. Aug. 2024. URL: https://energinet.dk/media/ox0gqmvw/gaeldende-prequalification-of-units-and-aggregated-portfolios.pdf.

[43] Engelhardt, J., Thingvad, A., Zepter, J. M., Gabderakhmanova, T., and Marinelli, M. "Energy recovery strategies for batteries providing frequency containment reserve in the Nordic power system". *Sustainable Energy, Grids and Networks*. 2022 Dec. 32. P. 100947. DOI: 10.1016/j.segan.2022.100947.

[44] BeagleBoard.org. *BeagleBone Black*. URL: https://docs.beagle.cc/boards/beaglebone/black/index.html.

[45] DEIF, *MIC-2 MKII Multi-instrument*. Tech. rep. 2025. URL: https://deif-cdn-umbraco.azureedge.net/media/b3jj24la/mic-2-mkii-data-sheet-4921210156-uk.pdf?rnd=133785741555300000&v=9.

[46] EnergyDataDK, D. *API descriptions*. 2025. URL: https://energydata.dk/en/api-descriptions/.

[47] Senol, M., Bayram, I. S., Naderi, Y., and Galloway, S. "Electric Vehicles Under Low Temperatures: A Review on Battery Performance, Charging Needs, and Power Grid Impacts". *IEEE Access*. 2023. 11. Pp. 39879–39912. DOI: 10.1109/ACCESS.2023.3268615.

[48] Energinet, *Forskrift C2 - Balancemarked og Balanceafregning*. Tech. rep. 2025, pp. 3–4. URL: https://energinet.dk/media/pv0lz5jc/forskrift-c2-opdateret-18-marts-2025.pdf.

[49] iocCharger. "Commercial EV Charging Solutions". 2025. URL: https://www.iocharger.com/commercial-ev-charging-solutions/.

[50] Kang Ma, Lurui Fang, and Wangwei Kong. "Review of distribution network phase unbalance: Scale, causes, consequences, solutions, and future research direction". *CSEE Journal of Power and Energy Systems*. 2020. DOI: 10.17775/CSEEJPES.2019.03280.

[51] Sikkerhedsstyrelsen, *Fællesregulativet*. Tech. rep. 2022. URL: https://dinel.dk/globalassets/dinel/regler/fallesregulativet_2022.pdf.

[52] Lunde, G. A., Damm, E. V., Gade, P. A. V., and Kazempour, J. "Aggregator of Electric Vehicles Bidding in Nordic FCR-D Markets: A Chance-Constrained Program". 2024 Apr. DOI: https://doi.org/10.48550/arXiv.2404.12818. URL: http://arxiv.org/abs/2404.12818.

Appendix

Complete Overview of Loop Time Data

As only the relevant fractions regarding each subsection in subsection 5.4 were presented, Table 17 presents all the results obtained in the process of determining which component lead to C2 through C4 being slower than C1 as well as investigating the scalability of our control.

Table 17: Comprehensive overview of the results presented in subsection 5.4

Test No.	**Charger No.**	**Number of Chargers Live**	**Simulation vs. Physical**	**Control Flag***	**Average loop time [ms]**	**CPU Load on BBB from `top`** [%]**
1	C1	1	Simulation	3	105	46.6
2	C2	1	Simulation	3	103	45.6
3	C3	1	Simulation	3	103	47.0
4	C4	1	Simulation	3	108	45.9
5	C1	1	Physical - ID4	3	407	24.6
6	C2	1	Physical - ID4	3	1513	11.7
7	C3	1	Physical - ID4	3	1460	11.7
8	C4	1	Physical - ID4	3	1537	11.7
9	C2 - new CC	1	Physical - ID4	3	401	20.8
10	C1	1	Physical - Nissan Leaf	3	479	24.9
11	C1	30	Physical - Nissan Leaf***	3	598	44.2
12	C1	30	Simulation	3	240	74
13	C1	1	Simulation	1	96	33.2
14	C1	30	Simulation	1	122	64
15	C1	30	Simulation	4	277	75

* Control flags indicating whether or not to respond to frequency deviation and where the priority-based current dispatch is to be calculated. See Table 8.

** Running the command `top` in a Linux system displays the active processes in descending order of CPU usage.

*** One physical Nissan Leaf connected to C1, while the remaining 29 chargers are simulated.

Average Loop Time vs. CPU Load

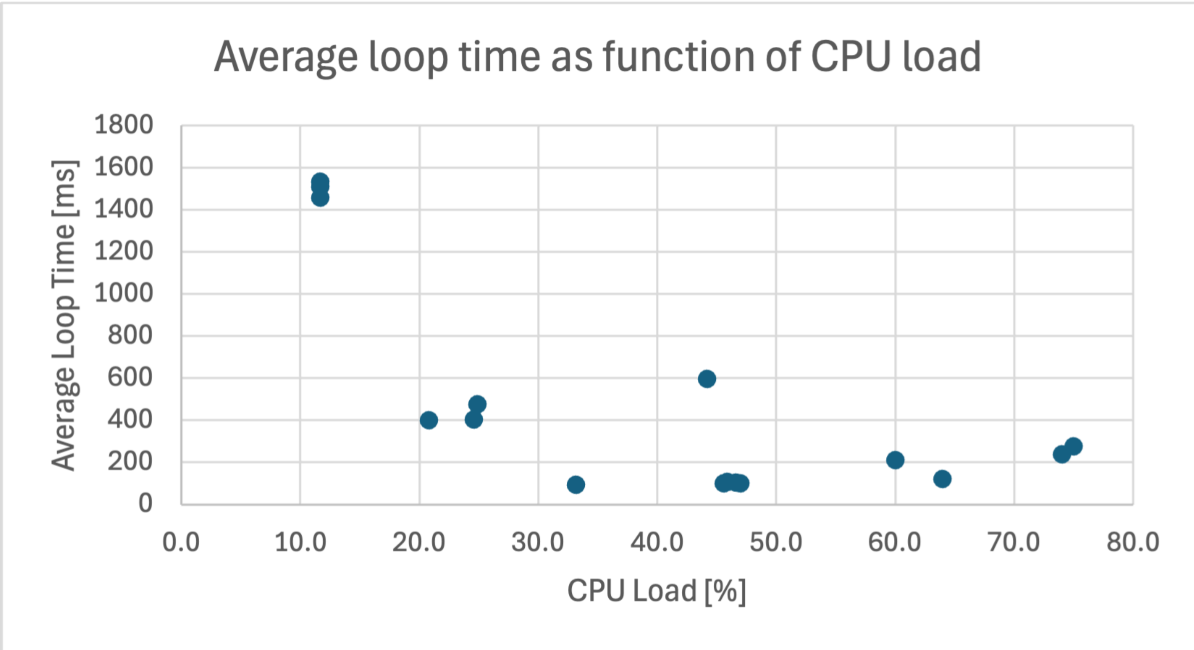

Figure 45: Average loop time vs CPU load. It may seem that with increasing CPU load, the average loop time decreases, but we suspect that the increasing CPU load comes from the shorter delays on physical components, which increases the load factor of the CPU. It should be noted that when the CPU approaches 100%, severe increases in loop time are expected. This is however not investigated.

Daily Earnings for Ancillary Services

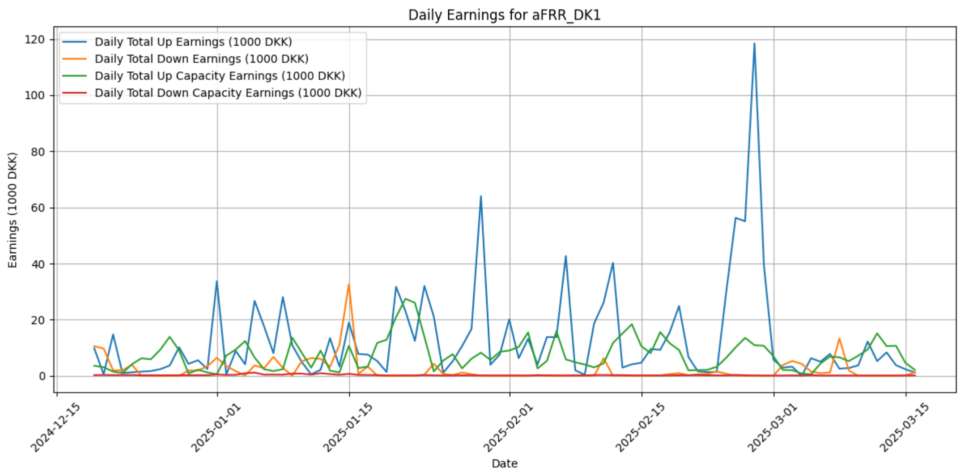

Figure 46: Daily earnings from participation in the aFRR market in DK2. The blue and yellow lines represent earnings from energy activations for up- and down-regulation, respectively, assuming activation during the 20% highest price periods. The green and red lines show earnings from capacity provision (up and down). The time period is limited to three months, as the market structure changed with the integration into the PICASSO platform.

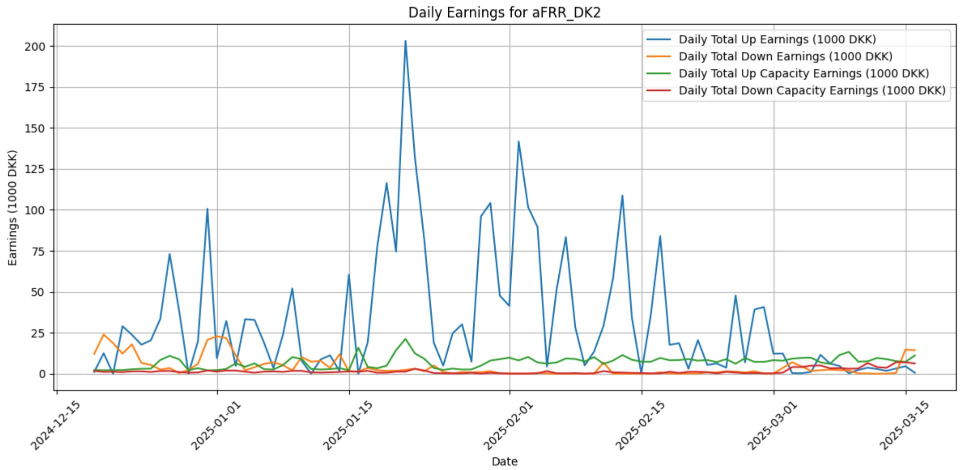

Figure 47: Daily earnings from participation in the aFRR market in DK2. The blue and yellow lines represent earnings from energy activations for up- and down-regulation, respectively, assuming activation during the 35% highest price periods. The green and red lines show earnings from capacity provision (up and down). The time period is limited to three months, as the market structure changed with the integration into the PICASSO platform.

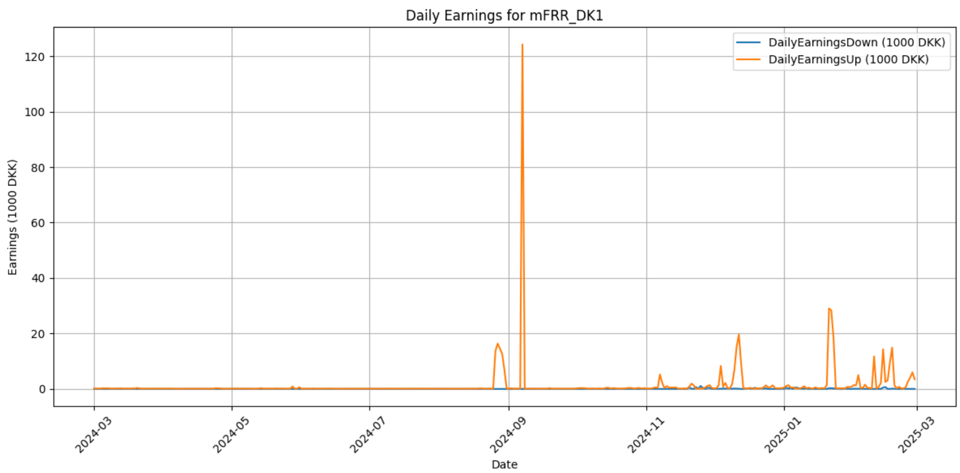

Figure 48: Daily Earnings for mFRR DK1

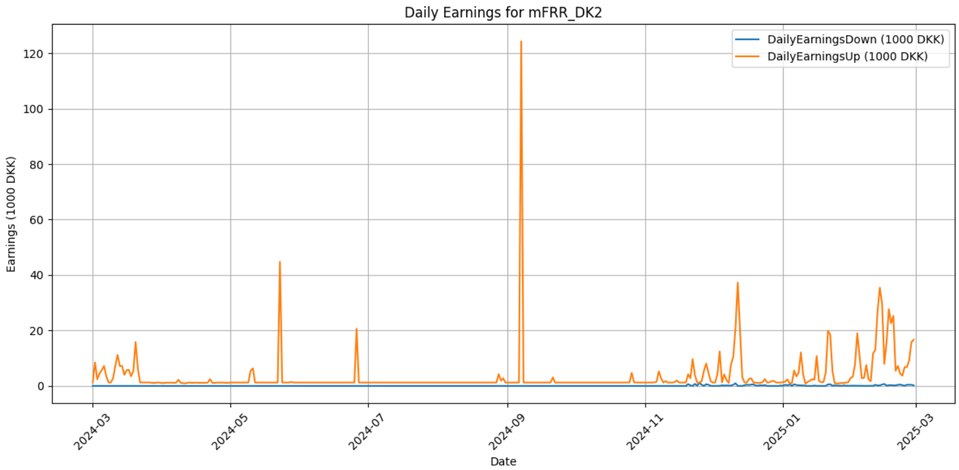

Figure 49: Daily Earnings for mFRR DK2

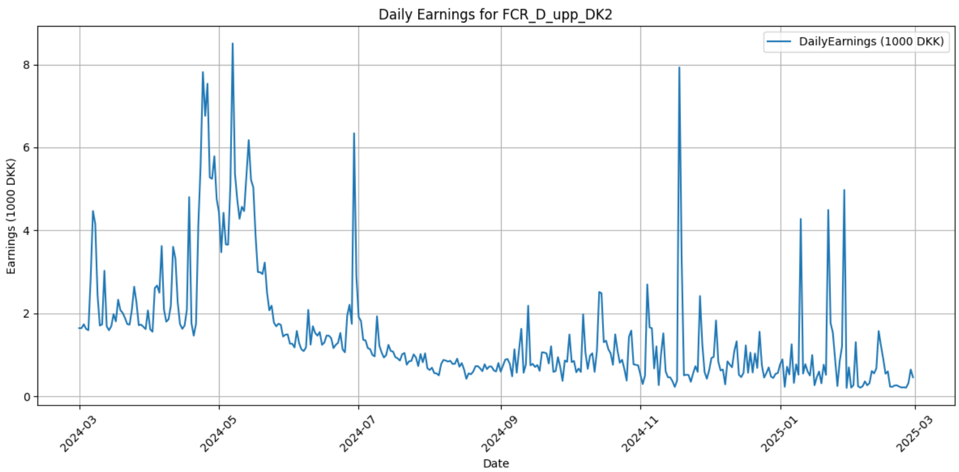

Figure 50: Daily Earnings for FCR-D

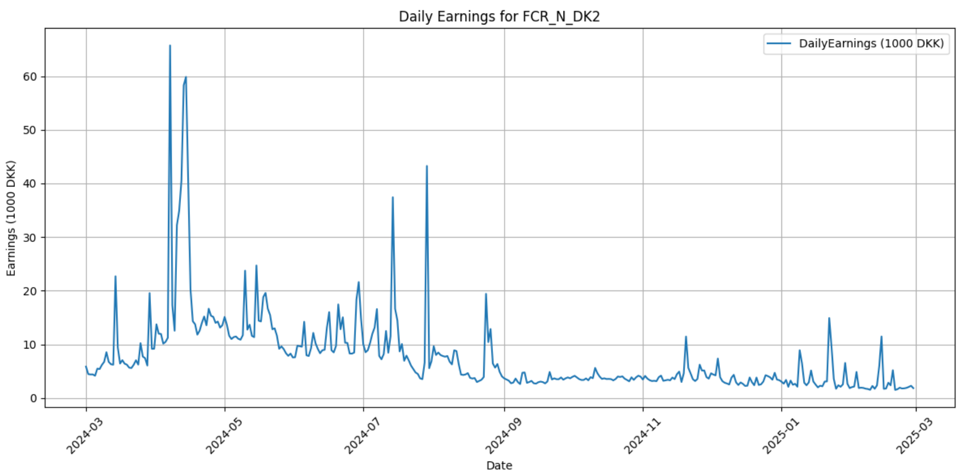

Figure 51: Daily Earnings for FCR-N

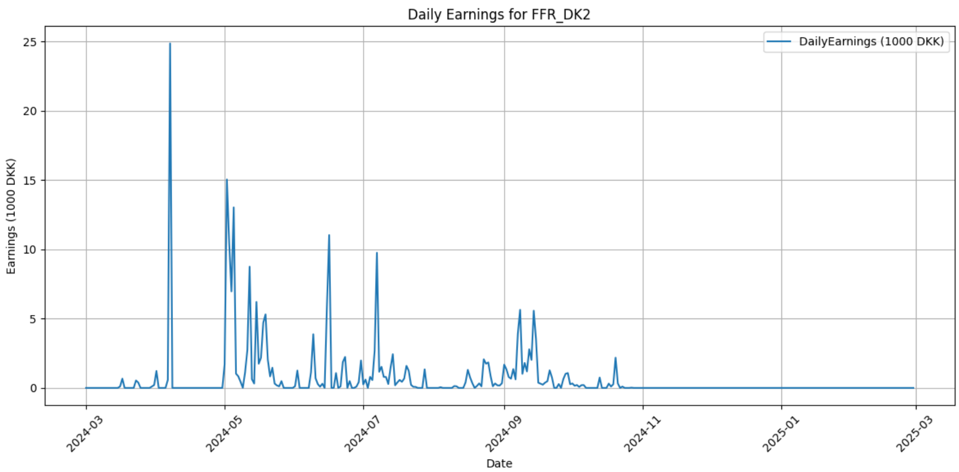

Figure 52: Daily Earnings for FFR